JUNIPER NETWORKS

Junos® OS

Services Interfaces Configuration Guide

Release
13.1



Published: 2013-02-25

Copyright © 2013, Juniper Networks, Inc.

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

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

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

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

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

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

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

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS Services Interfaces Configuration Guide
Release 13.1
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

Revision History
March 2013—R1, Junos OS 13.1

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Abbreviated Table of Contents

About This Guide xlv

Part 1 Overview

Chapter 1 Services Interfaces Overview 3

Chapter 2 Services Interfaces Configuration Statements 5

Part 2 Adaptive Services

Chapter 3 Adaptive Services Overview 37

Chapter 4 Applications Configuration Guidelines 73

Chapter 5 Summary of Applications Configuration Statements 121

Chapter 6 Stateful Firewall Services Configuration Guidelines 131

Chapter 7 Summary of Stateful Firewall Configuration Statements 141

Chapter 8 Stateful Firewall on the Embedded Junos OS Platform Configuration Guidelines 153

Chapter 9 Summary of Stateful Firewall on the Embedded Junos OS Platform Configuration Statements 157

Chapter 10 Load Balancing Configuration Guidelines 167

Chapter 11 Summary of Load Balancing Configuration Statements 173

Chapter 12 Intrusion Detection Service Configuration Guidelines 185

Chapter 13 Summary of Intrusion Detection Service Configuration Statements 197

Chapter 14 IPsec Services Configuration Guidelines 219

Chapter 15 Summary of IPsec Services Configuration Statements 273

Chapter 16 Layer 2 Tunneling Protocol Services Configuration Guidelines 309

Chapter 17 Summary of Layer 2 Tunneling Protocol Configuration Statements 329

Chapter 18 Link Services IQ Interfaces Configuration Guidelines 347

Chapter 19 Summary of Link Services IQ Configuration Statements 409

Chapter 20 Voice Services Configuration Guidelines 421

Chapter 21 Summary of Voice Services Configuration Statements 431

Chapter 22 Class-of-Service Configuration Guidelines 441

Chapter 23 Summary of Class-of-Service Configuration Statements 451

Chapter 24 Service Set Configuration Guidelines 467

Chapter 25 Summary of Service Set Configuration Statements 485

Chapter 26 Service Interface Configuration Guidelines 513

Copyright © 2013, Juniper Networks, Inc.

Chapter 27 Summary of Service Interface Configuration Statements 527
Chapter 28 PGCP Configuration Guidelines for the BGF Feature 547
Chapter 29 Summary of PGCP Configuration Statements 553
Chapter 30 Service Interface Pools Configuration Guidelines 655
Chapter 31 Summary of Service Interface Pools Statements 657
Chapter 32 Border Signaling Gateway Configuration Guidelines 659
Chapter 33 Summary of Border Signaling Gateway Configuration Statements 665
Chapter 34 PTSP Configuration Guidelines 745
Chapter 35 Summary of PTSP Configuration Statements 747

Part 3 Dynamic Application Awareness for Junos OS

Chapter 36 Dynamic Application Awareness for Junos OS Overview 771
Chapter 37 Application Identification Configuration Guidelines 781
Chapter 38 Summary of Application Identification Configuration Statements 799
Chapter 39 Application-Aware Access List Configuration Guidelines 839
Chapter 40 Summary of AACL Configuration Statements 847
Chapter 41 Local Policy Decision Function Configuration Guidelines 859
Chapter 42 Summary of L-PDF Configuration Statements 867

Part 4 Encryption Services

Chapter 43 Encryption Overview 879
Chapter 44 Encryption Interfaces Configuration Guidelines 881
Chapter 45 Summary of Encryption Configuration Statements 891

Part 5 Flow Monitoring and Discard Accounting Services

Chapter 46 Flow Monitoring and Discard Accounting Overview 903
Chapter 47 Flow Monitoring and Discard Accounting Configuration Guidelines 907
Chapter 48 Summary of Flow-Monitoring Configuration Statements 977
Chapter 49 Flow Collection Configuration Guidelines 1053
Chapter 50 Summary of Flow Collection Configuration Statements 1065
Chapter 51 Dynamic Flow Capture Configuration Guidelines 1083
Chapter 52 Flow-Tap Configuration Guidelines 1095
Chapter 53 Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements 1103

Part 6 Link and Multilink Services

Chapter 54 Link and Multilink Services Overview 1123
Chapter 55 Link and Multilink Services Configuration Guidelines 1127
Chapter 56 Summary of Multilink and Link Services Configuration Statements 1189

Copyright © 2013, Juniper Networks, Inc.

Part 7 Real-Time Performance Monitoring Services

Chapter 57 Real-Time Performance Monitoring Services Overview 1215

Chapter 58 Real-Time Performance Monitoring Configuration Guidelines 1217

Chapter 59 Summary of Real-Time Performance Monitoring Configuration Statements 1237

Part 8 Tunnel Services

Chapter 60 Tunnel Services Overview 1269

Chapter 61 Tunnel Interfaces Configuration Guidelines 1273

Chapter 62 Summary of Tunnel Services Configuration Statements 1295

Part 9 Index

Index 1313

Index of Statements and Commands 1337

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Table of Contents

About This Guide xlv

Junos Documentation and Release Notes xlv
Objectives xlvi
Audience xlvi
Supported Platforms xlvi
Using the Indexes xlvii
Using the Examples in This Manual xlvii
Merging a Full Example xlvii
Merging a Snippet xlviii
Documentation Conventions xlviii
Documentation Feedback l
Requesting Technical Support l
Self-Help Online Tools and Resources li
Opening a Case with JTAC li

Part 1 Overview

Chapter 1 Services Interfaces Overview 3

Services PIC Types 3
Supported Platforms 4

Chapter 2 Services Interfaces Configuration Statements 5

[edit applications] Hierarchy Level 5
[edit forwarding-options] Hierarchy Level 6
[edit interfaces] Hierarchy Level 8
[edit logical-systems] Hierarchy Level 12
[edit services] Hierarchy Level 12

Part 2 Adaptive Services

Chapter 3 Adaptive Services Overview 37

Adaptive Services Overview 37
Enabling Service Packages 39
Layer 2 Service Package Capabilities and Interfaces 42
Services Configuration Procedure 43
Packet Flow Through the Adaptive Services or Multiservices PIC 44
Stateful Firewall Overview 45
Stateful Firewall Support for Application Protocols 46
Stateful Firewall Anomaly Checking 46

Copyright © 2013, Juniper Networks, Inc.

Network Address Translation Overview 48
Types of NAT 48
NAT Concept and Facilities Overview 48
IPv4-to-IPv4 Basic NAT 49
NAT-PT 50
Static Destination NAT 50
Twice NAT 50
IPv6 NAT 51
NAT-PT with DNS ALG 51
Dynamic NAT 51
Stateful NAT64 52
Dual-Stack Lite 52
Tunneling Services for IPv4-to-IPv6 Transition Overview 53
6to4 Overview 53
Basic 6to4 53
6to4 Anycast 54
6to4 Provider-Managed Tunnels 54
DS-Lite Softwires—IPv4 over IPv6 55
6rd Softwires—IPv6 over IPv4 56
IPsec Overview 57
IPsec 57
Security Associations 57
IKE 58
Comparison of IPsec Services and ES Interface Configuration 58
Layer 2 Tunneling Protocol Overview 59
Voice Services Overview 59
HTTP URL Tracking and Policy Control for Client Requests 60
Guidelines for Configuring HTTP URL Monitoring for Client Requests 60
Configuring HTTP URL Tracking and Policy Control 61
Class of Service Overview 62
Examples: Services Interfaces Configuration 62
Example: Service Interfaces Configuration 63
Example: VPN Routing and Forwarding (VRF) and Service Configuration . . 65
Example: Dynamic Source NAT as a Next-Hop Service 67
Example: NAT Between VRFs Configuration 68
Example: BOOTP and Broadcast Addresses 71

Chapter 4 Applications Configuration Guidelines 73

Configuring Application Protocol Properties 74
Configuring an Application Protocol 74
Configuring the Network Protocol 76
Configuring the ICMP Code and Type 78
Configuring Source and Destination Ports 79
Configuring the Inactivity Timeout Period 82
Configuring SIP 82
Limitations 83
Configuring an SNMP Command for Packet Matching 84
Configuring an RPC Program Number 84
Configuring the TTL Threshold 84

Copyright © 2013, Juniper Networks, Inc.

Configuring a Universal Unique Identifier . . . 84
Configuring Application Sets . . . 85
ALG Descriptions . . . 85
Supported ALGs . . . 85
ALG Support Details . . . 86
Basic TCP ALG . . . 87
Basic UDP ALG . . . 88
BOOTP . . . 88
DCE RPC Services . . . 88
DNS . . . 88
FTP . . . 89
H323 . . . 89
ICMP . . . 89
IIOP . . . 90
IP . . . 90
NetBIOS . . . 90
NetShow . . . 90
ONC RPC Services . . . 90
PPTP . . . 91
RealAudio . . . 91
Sun RPC and RPC Portmap Services . . . 92
RTSP . . . 93
SIP . . . 93
SNMP . . . 94
SQLNet . . . 94
TFTP . . . 94
Traceroute . . . 94
UNIX Remote-Shell Services . . . 95
Winframe . . . 95
Juniper Networks Defaults . . . 95
Verifying the Output of ALG Sessions . . . 105
FTP Example . . . 106
Sample Output . . . 106
FTP System Log Messages . . . 107
Analysis . . . 107
Troubleshooting Questions . . . 108
RTSP ALG Example . . . 108
Sample Output . . . 108
Analysis . . . 109
Troubleshooting Questions . . . 109
System Log Messages . . . 110
System Log Configuration . . . 110
System Log Output . . . 111
Junos Default Groups . . . 111
Examples: Referencing the Preset Statement from the Junos Default Group . . . 117
Examples: Configuring Application Protocols . . . 118

Copyright © 2013, Juniper Networks, Inc.

Chapter 5 Summary of Applications Configuration Statements 121

application 121
application-protocol 122
application-set 123
applications 123
destination-port 124
icmp-code 124
icmp-type 125
inactivity-timeout 125
learn-sip-register 126
protocol 127
rpc-program-number 128
sip-call-hold-timeout 128
snmp-command 129
source-port 129
ttl-threshold 130
uuid 130

Chapter 6 Stateful Firewall Services Configuration Guidelines 131

Configuring Stateful Firewall Rules 132
Configuring Match Direction for Stateful Firewall Rules 132
Configuring Match Conditions in Stateful Firewall Rules 133
Configuring Actions in Stateful Firewall Rules 134
Configuring IP Option Handling 135
Configuring Stateful Firewall Rule Sets 136
Examples: Configuring Stateful Firewall Rules 136

Chapter 7 Summary of Stateful Firewall Configuration Statements 141

allow-ip-options 142
application-sets 143
applications 143
destination-address 144
destination-address-range 144
destination-prefix-list 145
from 146
match-direction 146
rule 147
rule-set 148
services (stateful-firewall) 148
source-address 149
source-address-range 149
source-prefix-list 150
syslog 150
term 151
then 152

Copyright © 2013, Juniper Networks, Inc.

Chapter 8 Stateful Firewall on the Embedded Junos OS Platform Configuration Guidelines . . . 153

Loading the Stateful Firewall Plug-In . . . 153
Configuring Memory for the Stateful Firewall Plug-In . . . 155
Configuring rsh, rlogin, rexec for Stateful Firewall . . . 155

Chapter 9 Summary of Stateful Firewall on the Embedded Junos OS Platform Configuration Statements . . . 157

control-cores . . . 157
data-cores . . . 158
data-flow-affinity . . . 158
destination (Chassis) . . . 159
extension-provider . . . 160
forwarding-db-size . . . 161
hash-key (Chassis) . . . 162
object-cache-size . . . 163
package (Loading on PIC) . . . 163
policy-db-size . . . 164
syslog (Chassis) . . . 165
wired-process-mem-size . . . 166

Chapter 10 Load Balancing Configuration Guidelines . . . 167

Configuring Load Balancing on AMS Infrastructure . . . 167
Configuring AMS Infrastructure . . . 167
Configuring High Availability . . . 168
Load Balancing Network Address Translation Flows . . . 169
Example: Configuring Static Source Translation on AMS Infrastructure . . . 169

Chapter 11 Summary of Load Balancing Configuration Statements . . . 173

drop-member-traffic (Aggregated Multiservices) . . . 173
enable-rejoin (aggregated Multiservices) . . . 174
family (aggregated Multiservices) . . . 174
high-availability-options (aggregated Multiservices) . . . 175
interfaces (Aggregated Multiservices) . . . 176
load-balancing-options (Aggregated Multiservices) . . . 177
many-to-one (Aggregated Multiservices) . . . 178
member-failure-options (Aggregated Multiservices) . . . 179
member-interface (Aggregated Multiservices) . . . 181
redistribute-all-traffic (Aggregated Multiservices) . . . 182
rejoin-timeout (Aggregated Multiservices) . . . 182
unit (Aggregated Multiservices) . . . 183

Chapter 12 Intrusion Detection Service Configuration Guidelines . . . 185

Configuring IDS Rules . . . 187
Configuring Match Direction for IDS Rules . . . 188
Configuring Match Conditions in IDS Rules . . . 189
Configuring Actions in IDS Rules . . . 190
Configuring IDS Rule Sets . . . 193
Examples: Configuring IDS Rules . . . 193

Copyright © 2013, Juniper Networks, Inc.

Chapter 13 Summary of Intrusion Detection Service Configuration Statements 197

aggregation 197
application-sets (Services IDS) 198
applications (Services IDS) 198
by-destination 199
by-pair 200
by-source 201
destination-address (Services IDS) 202
destination-address-range (Services IDS) 202
destination-prefix (Services IDS) 203
destination-prefix-ipv6 203
destination-prefix-list (Services IDS) 204
force-entry 204
from (Services IDS) 205
ignore-entry 205
logging (Services IDS) 205
match-direction (Services IDS) 206
mss 206
rule (Services IDS) 207
rule-set (Services IDS) 208
services (IDS) 208
session-limit 209
source-address (Services IDS) 210
source-address-range (Services IDS) 210
source-prefix (Services IDS) 211
source-prefix-ipv6 211
source-prefix-list (Services IDS) 212
syn-cookie 212
syslog (Services IDS) 213
term (Services IDS) 214
then (Services IDS) 216
threshold (Services) 217

Chapter 14 IPsec Services Configuration Guidelines 219

Minimum Security Association Configurations 221
Minimum Manual SA Configuration 221
Minimum Dynamic SA Configuration 221
Configuring Security Associations 222
Configuring Manual Security Associations 223
Configuring the Direction for IPsec Processing 224
Configuring the Protocol for a Manual IPsec SA 225
Configuring the Security Parameter Index 225
Configuring the Auxiliary Security Parameter Index 225
Configuring Authentication for a Manual IPsec SA 225
Configuring Encryption for a Manual IPsec SA 226
Configuring Dynamic Security Associations 227
Clearing Security Associations 228

Copyright © 2013, Juniper Networks, Inc.

Configuring IKE Proposals 228
Configuring the Authentication Algorithm for an IKE Proposal 229
Configuring the Authentication Method for an IKE Proposal 229
Configuring the Diffie-Hellman Group for an IKE Proposal 230
Configuring the Encryption Algorithm for an IKE Proposal 230
Configuring the Lifetime for an IKE SA 231
Example: Configuring an IKE Proposal 232
Configuring IKE Policies 232
Configuring the IKE Phase 233
Configuring the Mode for an IKE Policy 233
Configuring the Proposals in an IKE Policy 234
Configuring the Preshared Key for an IKE Policy 234
Configuring the Local Certificate for an IKE Policy 234
Configuring a Certificate Revocation List 235
Configuring the Description for an IKE Policy 236
Configuring Local and Remote IDs for IKE Phase 1 Negotiation 236
Example: Configuring an IKE Policy 236
Configuring IPsec Proposals 237
Configuring the Authentication Algorithm for an IPsec Proposal 238
Configuring the Description for an IPsec Proposal 238
Configuring the Encryption Algorithm for an IPsec Proposal 238
Configuring the Lifetime for an IPsec SA 239
Configuring the Protocol for a Dynamic SA 240
Configuring IPsec Policies 240
Configuring the Description for an IPsec Policy 241
Configuring Perfect Forward Secrecy 241
Configuring the Proposals in an IPsec Policy 241
Example: Configuring an IPsec Policy 242
IPsec Policy for Dynamic Endpoints 242
Configuring IPsec Rules 243
Configuring Match Direction for IPsec Rules 244
Configuring Match Conditions in IPsec Rules 244
Configuring Actions in IPsec Rules 246
Enabling IPsec Packet Fragmentation 247
Configuring Destination Addresses for Dead Peer Detection 247
Configuring or Disabling IPsec Anti-Replay 248
Enabling System Log Messages 249
Specifying the MTU for IPsec Tunnels 249
Configuring IPsec Rule Sets 249
Configuring Dynamic Endpoints for IPsec Tunnels 250
Authentication Process 250
Implicit Dynamic Rules 251
Reverse Route Insertion 251
Configuring an IKE Access Profile 252
Referencing the IKE Access Profile in a Service Set 253
Configuring the Interface Identifier 254
Default IKE and IPsec Proposals 254

Copyright © 2013, Juniper Networks, Inc.

Tracing IPsec Operations . . . 255
Disabling IPsec Tunnel Endpoint in Traceroute . . . 256
Tracing IPsec PKI Operations . . . 256
Configuring IPSec on the Services SDK . . . 257
Examples: Configuring IPsec Services . . . 258
Example: Configuring Statically Assigned Tunnels . . . 258
Example: Configuring Dynamically Assigned Tunnels . . . 261
Multitask Example: Configuring IPsec Services . . . 265
Configuring the IKE Proposal . . . 266
Configuring the IKE Policy (and Referencing the IKE Proposal) . . . 266
Configuring the IPsec Proposal . . . 267
Configuring the IPsec Policy (and Referencing the IPsec Proposal) . . . 268
Configuring the IPsec Rule (and Referencing the IKE and IPsec Policies) . . . 268
Configuring IPsec Trace Options . . . 269
Configuring the Access Profile (and Referencing the IKE and IPsec Policies) . . . 270
Configuring the Service Set (and Referencing the IKE Profile and the IPsec Rule) . . . 271

Chapter 15 Summary of IPsec Services Configuration Statements . . . 273

anti-replay-window-size (Services IPsec VPN) . . . 273
authentication (Services IPsec VPN) . . . 274
authentication-algorithm . . . 275
authentication-algorithm (Services IKE) . . . 275
authentication-algorithm (Services IPsec) . . . 275
authentication-method (Services IPsec VPN) . . . 276
auxiliary-spi (Services IPsec VPN) . . . 276
backup-remote-gateway . . . 277
clear-dont-fragment-bit (Services IPsec VPN) . . . 277
clear-ike-sas-on-pic-restart . . . 278
clear-ipsec-sas-on-pic-restart . . . 278
description (Services IPsec VPN) . . . 279
destination-address (Services IPsec VPN) . . . 279
dh-group . . . 280
direction . . . 281
dynamic . . . 282
encryption . . . 283
encryption-algorithm (Services IPsec VPN) . . . 284
from (Services IPsec VPN) . . . 285
ike . . . 286
initiate-dead-peer-detection . . . 287
ipsec (Services IPsec VPN) . . . 287
ipsec-inside-interface . . . 288
lifetime-seconds (Services IPsec VPN) . . . 288
local-certificate (Services IPsec VPN) . . . 289
local-id . . . 289
manual . . . 290
match-direction (Services IPsec VPN) . . . 290

Copyright © 2013, Juniper Networks, Inc.

mode (Services IPsec VPN) . . . 291
no-anti-replay (Services IPsec VPN) . . . 291
no-ipsec-tunnel-in-traceroute . . . 292
perfect-forward-secrecy (Services IPsec VPN) . . . 292
policy . . . 293
policy (Services IKE) . . . 293
policy (Services IPsec VPN) . . . 294
pre-shared-key (Services IKE) . . . 294
proposal . . . 295
proposal (Services IKE) . . . 295
proposal (Services IPsec VPN) . . . 296
proposals . . . 296
protocol . . . 297
remote-gateway . . . 297
remote-id . . . 298
rule (Services IPsec VPN) . . . 299
rule-set (Services IPsec VPN) . . . 300
services (IPsec VPN) . . . 300
source-address (Services IPsec VPN) . . . 301
spi . . . 301
syslog (Services IPsec VPN) . . . 302
term (Services IPsec VPN) . . . 303
then (Services IPsec VPN) . . . 304
traceoptions (Services IPsec VPN) . . . 305
traceoptions (PKI) . . . 307
tunnel-mtu (Services IPsec VPN) . . . 308
version (IKE) . . . 308

Chapter 16 Layer 2 Tunneling Protocol Services Configuration Guidelines . . . 309

L2TP Services Configuration Overview . . . 311
L2TP Minimum Configuration . . . 312
Configuring L2TP Tunnel Groups . . . 314
Configuring Access Profiles for L2TP Tunnel Groups . . . 315
Configuring the Local Gateway Address and PIC . . . 315
Configuring Window Size for L2TP Tunnels . . . 316
Configuring Timers for L2TP Tunnels . . . 316
Hiding Attribute-Value Pairs for L2TP Tunnels . . . 316
Configuring System Logging of L2TP Tunnel Activity . . . 317
Configuring the Identifier for Logical Interfaces that Provide L2TP Services . . . 318
Example: Configuring Multilink PPP on a Shared Logical Interface . . . 319
AS PIC Redundancy for L2TP Services . . . 320
IP Packet Fragment Reassembly for L2TP Overview . . . 320
Configuring IP Inline Reassembly for L2TP . . . 321
Tracing L2TP Operations . . . 322
Examples: Configuring L2TP Services . . . 324

Chapter 17 Summary of Layer 2 Tunneling Protocol Configuration Statements . . . 329

facility-override . . . 329
hello-interval . . . 330
hide-avps . . . 331

Copyright © 2013, Juniper Networks, Inc.

host (L2TP) . . . 331
ip-reassembly (L2TP) . . . 332
l2tp-access-profile . . . 332
local-gateway address . . . 333
log-prefix (L2TP) . . . 333
maximum-send-window . . . 334
ppp-access-profile . . . 334
receive-window . . . 335
retransmit-interval (Services) . . . 335
service-interface . . . 336
services . . . 337
services (Hierarchy) . . . 337
services (L2TP System Logging) . . . 338
syslog . . . 339
traceoptions (L2TP) . . . 340
tunnel-group . . . 344
tunnel-timeout . . . 345

Chapter 18 **Link Services IQ Interfaces Configuration Guidelines . . . 347**

Layer 2 Service Package Capabilities and Interfaces . . . 348
Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS . . . 350
Configuring the Association between LSQ and SONET Interfaces . . . 350
Configuring SONET APS Interoperability with Cisco Systems FRF.16 . . . 351
Restrictions on APS Redundancy for LSQ Interfaces . . . 352
Configuring LSQ Interface Redundancy in a Single Router Using SONET APS . . 352
Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces . . . 352
Configuring Redundant Paired LSQ Interfaces . . . 353
Restrictions on Redundant LSQ Interfaces . . . 354
Configuring Link State Replication for Redundant Link PICs . . . 356
Examples: Configuring Redundant LSQ Interfaces for Failure Recovery . . . 357
Configuring CoS Scheduling Queues on Logical LSQ Interfaces . . . 361
Configuring Scheduler Buffer Size . . . 363
Configuring Scheduler Priority . . . 363
Configuring Scheduler Shaping Rate . . . 363
Configuring Drop Profiles . . . 364
Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces . . . 365
Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces . . . 367
Configuring Multiclass MLPPP on LSQ Interfaces . . . 367
Oversubscribing Interface Bandwidth on LSQ Interfaces . . . 369
Examples: Oversubscribing an LSQ Interface . . . 372
Configuring Guaranteed Minimum Rate on LSQ Interfaces . . . 374
Example: Configuring Guaranteed Minimum Rate . . . 376
Configuring Link Services and CoS on Services PICs . . . 377
Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP . . . 380
Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP . . 383
Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 . . . 385
Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16 . . 388

Copyright © 2013, Juniper Networks, Inc.

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI . . . 390
Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI . . . 393
Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 . . . 395
Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12 . . . 398
Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 . . . 402
Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP . . . 403
Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 . . . 404
Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP . . . 406

Chapter 19 Summary of Link Services IQ Configuration Statements . . . 409

cisco-interoperability . . . 409
forwarding-class . . . 410
fragment-threshold . . . 411
fragmentation-map . . . 411
fragmentation-maps . . . 412
hot-standby . . . 412
link-layer-overhead . . . 413
lsq-failure-options . . . 413
multilink-class . . . 414
multilink-max-classes . . . 414
no-fragmentation . . . 415
no-per-unit-scheduler . . . 415
no-termination-request . . . 416
per-unit-scheduler . . . 417
preserve-interface . . . 418
primary . . . 418
redundancy-options . . . 419
secondary . . . 419
trigger-link-failure . . . 420
warm-standby . . . 420

Chapter 20 Voice Services Configuration Guidelines . . . 421

Configuring Services Interfaces for Voice Services . . . 422
Configuring the Logical Interface Address for the MLPPP Bundle . . . 422
Configuring Compression of Voice Traffic . . . 423
Configuring Delay-Sensitive Packet Interleaving . . . 424
Example: Configuring Compression of Voice Traffic . . . 424
Configuring Encapsulation for Voice Services . . . 425
Configuring Network Interfaces for Voice Services . . . 425
Configuring Voice Services Bundles with MLPPP Encapsulation . . . 426
Configuring the Compression Interface with PPP Encapsulation . . . 426
Examples: Configuring Voice Services . . . 426

Copyright © 2013, Juniper Networks, Inc.

Chapter 21 Summary of Voice Services Configuration Statements 431
address (Interfaces) 431
bundle 432
compression 432
compression-device (Interfaces) 433
encapsulation 433
f-max-period 434
family (Interfaces) 435
fragment-threshold (Interfaces LSQ) 436
interfaces 436
maximum-contexts 437
port 437
queues 438
rtp 438
unit (Interfaces) 439
Chapter 22 Class-of-Service Configuration Guidelines 441
Restrictions and Cautions for CoS Configuration on Services Interfaces 442
Configuring CoS Rules 443
Configuring Match Direction for CoS Rules 443
Configuring Match Conditions In CoS Rules 444
Configuring Actions in CoS Rules 445
Configuring Application Profiles for Use as CoS Rule Actions 446
Configuring Reflexive and Reverse CoS Rule Actions 446
Example: Configuring CoS Rules 447
Configuring CoS Rule Sets 448
Examples: Configuring CoS on Services Interfaces 448
Chapter 23 Summary of Class-of-Service Configuration Statements 451
application-profile 452
application-sets (Services CoS) 453
applications (Services CoS) 453
data (FTP) 454
destination-address (CoS) 454
destination-prefix-list (Services CoS) 455
dscp 455
forwarding-class 456
from (Services CoS) 456
ftp (Class-Of-Service) 457
match-direction 457
(reflexive | reverse) 458
rule 459
rule-set (Services CoS) 460
services (COS) 460
sip (Application Profile) 461
source-address (Services CoS) 461
source-prefix-list 462
syslog (Services CoS) 462
term (Services CoS) 463

Copyright © 2013, Juniper Networks, Inc.

then . . . 464
video (Application Profile) . . . 464
voice (Application Profile) . . . 465

Chapter 24 Service Set Configuration Guidelines . . . 467

Configuring Service Sets to be Applied to Services Interfaces . . . 468
Configuring Interface Service Sets . . . 468
Configuring Next-Hop Service Sets . . . 470
Determining Traffic Direction . . . 471
Interface Style Service Sets . . . 471
Next-Hop Style Service Sets . . . 472
Configuring Service Rules . . . 472
Configuring IPsec Service Sets . . . 473
Configuring the Local Gateway Address for IPsec Service Sets . . . 474
IKE Addresses in VRF Instances . . . 474
Configuring IKE Access Profiles for IPsec Service Sets . . . 475
Configuring Certification Authorities for IPsec Service Sets . . . 475
Configuring or Disabling Antireplay Service . . . 476
Clearing the Don't-Fragment Bit . . . 476
Configuring Passive-Mode Tunneling . . . 477
Configuring the Tunnel MTU Value . . . 477
Configuring Service Set Limitations . . . 478
Configuring System Logging for Service Sets . . . 479
Enabling Services PICs to Accept Multicast Traffic . . . 480
Tracing Services PIC Operations . . . 480
Configuring the Adaptive Services Log Filename . . . 481
Configuring the Number and Size of Adaptive Services Log Files . . . 482
Configuring Access to the Log File . . . 482
Configuring a Regular Expression for Lines to Be Logged . . . 482
Configuring the Trace Operations . . . 482
Example: Configuring Service Sets . . . 483

Chapter 25 Summary of Service Set Configuration Statements . . . 485

adaptive-services-pics . . . 485
allow-multicast . . . 486
anti-replay-window-size (Services Service Set) . . . 487
bypass-traffic-on-exceeding-flow-limits . . . 488
bypass-traffic-on-pic-failure . . . 488
clear-dont-fragment-bit (Services Service Set) . . . 489
facility-override . . . 490
host (service-set) . . . 491
ids-rules . . . 491
ike-access-profile . . . 492
interface-service . . . 492
ip-reassembly-rules (Service Set) . . . 493
ipsec-vpn-options . . . 494
ipsec-vpn-rules . . . 494
local-gateway . . . 495
log-prefix (Services) . . . 495
logging (Services) . . . 496

Copyright © 2013, Juniper Networks, Inc.

max-flows . . . 496
message-rate-limit . . . 497
nat-options . . . 498
nat-rules . . . 498
next-hop-service . . . 499
no-anti-replay (Services Service Set) . . . 500
passive-mode-tunneling . . . 500
pgcp-rules . . . 501
port (syslog) . . . 501
ptsp-rules . . . 502
service-interface . . . 502
service-set (Services) . . . 503
services . . . 505
services (Hierarchy) . . . 505
services (System Logging) . . . 506
stateful-firewall-rules . . . 507
syslog (Services Service Set) . . . 507
tcp-mss . . . 508
traceoptions (Services Logging) . . . 509
trusted-ca . . . 510
tunnel-mtu (Services Service Set) . . . 511

Chapter 26 Service Interface Configuration Guidelines . . . 513

Services Interface Naming Overview . . . 515
Configuring the Address and Domain for Services Interfaces . . . 516
Configuring Default Timeout Settings for Services Interfaces . . . 516
Configuring System Logging for Services Interfaces . . . 518
Enabling Fragmentation on GRE Tunnels . . . 519
Applying Filters and Services to Interfaces . . . 520
Configuring Service Filters . . . 521
Configuring AS or Multiservices PIC Redundancy . . . 522
Flow Offloading . . . 525
Examples: Configuring Services Interfaces . . . 525

Chapter 27 Summary of Service Interface Configuration Statements . . . 527

address (Interfaces) . . . 527
cgn-pic . . . 528
clear-dont-fragment-bit . . . 528
dial-options . . . 529
facility-override . . . 530
family (Interfaces) . . . 531
host (Interfaces) . . . 532
inactivity-timeout . . . 532
input (Interfaces) . . . 533
interfaces . . . 533
log-prefix (Interfaces) . . . 534
maximum . . . 534
open-timeout . . . 535
output . . . 535
post-service-filter . . . 536

Copyright © 2013, Juniper Networks, Inc.

primary (Interfaces) . . . 536
rate . . . 537
redundancy-options . . . 537
secondary (Interfaces) . . . 538
service . . . 538
service-domain . . . 539
service-filter (Interfaces) . . . 539
service-set (Interfaces) . . . 540
services (Interfaces) . . . 541
services-options . . . 542
session-limit . . . 543
source-address . . . 543
syslog (Interfaces) . . . 544
tcp-tickles . . . 544
trio-flow-offload . . . 545
unit . . . 546

Chapter 28 PGCP Configuration Guidelines for the BGF Feature . . . 547

Chapter 29 Summary of PGCP Configuration Statements . . . 553

administrative . . . 554
administrative (Control Association) . . . 554
administrative (Virtual Interface) . . . 555
algorithm . . . 556
application-data-inactivity-detection . . . 556
audit-observed-events-returns . . . 557
base-root . . . 558
bgf-core . . . 559
cancel-graceful . . . 561
cancel-graceful (Control Association) . . . 561
cancel-graceful (Virtual Interface) . . . 562
cleanup-timeout . . . 562
context-indications . . . 563
control-association-indications . . . 564
controller-address . . . 565
controller-failure . . . 565
controller-port . . . 566
data-inactivity-detection (Services PGCP) . . . 566
default . . . 567
delivery-function . . . 568
destination-address . . . 568
destination-port . . . 569
detect . . . 569
diffserv . . . 570
disable-session-mirroring . . . 570
disconnect . . . 571
down . . . 572
dscp (Services PGCP) . . . 573
encoding . . . 573
event-timestamp-notification . . . 574

Copyright © 2013, Juniper Networks, Inc.

failover-cold . . . 574
failover-warm . . . 575
failure . . . 576
fast-update-filters . . . 577
file (Services PGCP) . . . 578
flag (Services PGCP) . . . 579
gateway (Services PGCP) . . . 580
gateway-address . . . 584
gateway-controller . . . 585
gateway-port . . . 586
graceful . . . 587
graceful (Control Association) . . . 587
graceful (Virtual Interface) . . . 588
graceful-restart . . . 588
h248-options . . . 589
h248-profile . . . 591
h248-properties . . . 592
h248-stack . . . 595
h248-timers . . . 596
hanging-termination-detection . . . 596
inactivity-delay . . . 597
inactivity-duration (Services PGCP) . . . 597
inactivity-timeout . . . 598
inactivity-timer . . . 599
initial-average-ack-delay . . . 599
interim-ah-scheme . . . 600
ip-flow-stop-detection . . . 600
ipsec-transport-security-association . . . 601
latch-deadlock-delay . . . 601
max-burst-size . . . 602
max-burst-size (All Streams) . . . 602
max-burst-size (RTCP Streams) . . . 603
max-concurrent-calls . . . 604
maximum-fuf-percentage . . . 605
maximum-inactivity-time . . . 606
maximum-net-propagation-delay . . . 607
maximum-synchronization-mismatches . . . 607
maximum-terms . . . 608
maximum-waiting-delay . . . 608
media . . . 609
mg-maximum-pdu-size . . . 610
mg-originated-pending-limit . . . 611
mg-provisional-response-timer-value . . . 612
mg-segmentation-timer . . . 613
mgc-maximum-pdu-size . . . 614
mgc-originated-pending-limit . . . 615
mgc-provisional-response-timer-value . . . 616
mgc-segmentation-timer . . . 617
monitor . . . 618

Copyright © 2013, Juniper Networks, Inc.

network-operator-id . . . 618
no-dscp-bit-mirroring . . . 619
no-rtcp-check . . . 619
normal-mg-execution-time . . . 620
normal-mgc-execution-time . . . 621
notification-behavior . . . 622
notification-rate-limit . . . 622
notification-regulation . . . 623
overload-control . . . 623
peak-data-rate . . . 624
    peak-data-rate (All Streams) . . . 624
    peak-data-rate (RTCP) . . . 625
platform . . . 626
profile-name . . . 627
profile-version . . . 627
queue-limit-percentage . . . 628
reconnect . . . 629
reject-all-commands-threshold . . . 629
reject-new-calls-threshold . . . 630
report-service-change . . . 630
request-timestamp . . . 631
routing-instance . . . 631
rtcp . . . 632
rtp . . . 632
rule . . . 633
rule-set . . . 633
sbc-utils (Services PGCP) . . . 634
segmentation . . . 635
send-notification-on-delay . . . 636
service-change . . . 637
service-change-type . . . 638
service-interface . . . 638
service-state . . . 639
    service-state (Virtual BGF) . . . 639
    service-state (Virtual Interface) . . . 640
services (PGCP) . . . 640
session-mirroring . . . 641
source-address . . . 641
source-port (Mirrored BGF Packets) . . . 642
state-loss . . . 643
stop-detection-on-drop . . . 643
sustained-data-rate . . . 644
    sustained-data-rate (All Streams) . . . 644
    sustained-data-rate (RTCP Streams) . . . 645
timerx . . . 646
tmax-retransmission-delay . . . 646
traceoptions (Services PGCP) . . . 647
traffic-management . . . 648
up . . . 649

Copyright © 2013, Juniper Networks, Inc.

use-lower-case . . . 649
use-wildcard-response . . . 650
virtual-interface . . . 651
virtual-interface-down . . . 652
virtual-interface-indications . . . 653
virtual-interface-up . . . 653
warm . . . 654

Chapter 30 Service Interface Pools Configuration Guidelines . . . 655

Configuring Service Interface Pools . . . 655

Chapter 31 Summary of Service Interface Pools Statements . . . 657

interface . . . 657
pool . . . 658
service-interface-pools . . . 658

Chapter 32 Border Signaling Gateway Configuration Guidelines . . . 659

Chapter 33 Summary of Border Signaling Gateway Configuration Statements . . . 665

actions . . . 666
accelerations . . . 667
admission-control . . . 668
admission-control (Border Signaling Gateway) . . . 668
admission-control (New Transaction Policy) . . . 669
availability-check-profiles . . . 670
blacklist-period . . . 671
clusters . . . 672
committed-burst-size . . . 673
committed-information-rate . . . 674
data-inactivity-detection (Services Border Signaling Gateway) . . . 674
datastore . . . 675
default-media-realm . . . 676
dialogs . . . 677
dscp (Services Border Signaling Gateway) . . . 678
egress-service-point . . . 679
embedded-spdf . . . 680
file (Services Border Signaling Gateway) . . . 681
flag (Services Border Signaling Gateway) . . . 682
forward-manipulation . . . 683
framework . . . 684
from . . . 686
from (New Call Usage Policy) . . . 687
from (New Transaction Policy) . . . 688
from (Service Class) . . . 690
gateway (Services Border Signaling Gateway) . . . 691
inactivity-duration (Services Border Signaling Gateway) . . . 696
manipulation-rule . . . 697
media-policy . . . 698
media-type . . . 699
message-manipulation . . . 700
maximum-records-in-cache . . . 701

Copyright © 2013, Juniper Networks, Inc.

maximum-time-in-cache . . . 701
message-manipulation-rules . . . 702
minimum . . . 703
name-resolution-cache . . . 704
new-call-usage-input-policies . . . 704
new-call-usage-output-policies . . . 705
new-call-usage-policy . . . 706
new-call-usage-policy-set . . . 707
new-transaction-input-policies . . . 707
new-transaction-output-policies . . . 708
new-transaction-policy . . . 709
new-transaction-policy-set . . . 711
next-hop (Services) . . . 712
on-3xx-response . . . 713
request-uri . . . 714
reverse-manipulation . . . 715
route (Services) . . . 716
routing-destinations . . . 717
sbc-utils (Services Border Signaling Gateway) . . . 718
servers . . . 719
service-class . . . 720
service-interface . . . 721
    service-interface (Gateway) . . . 721
    service-interface (Service Point) . . . 721
service-point . . . 722
service-point-type . . . 723
service-policies . . . 723
services (Border Signaling Gateway) . . . 724
session-trace . . . 725
signaling . . . 726
signaling-realms . . . 727
sip . . . 728
sip-header . . . 731
sip-stack . . . 733
term . . . 735
    term (New Call Usage Policy) . . . 735
    term (New Transaction Policy) . . . 736
    term (Service Class) . . . 737
then . . . 738
    then (New Call Usage Policy) . . . 738
    then (New Transaction Policy) . . . 739
    then (Service Class) . . . 740
timer-c . . . 741
timers . . . 741
traceoptions (Services Border Gateway Signaling) . . . 742
transactions . . . 743
transport-details . . . 744

Copyright © 2013, Juniper Networks, Inc.

Chapter 34 PTSP Configuration Guidelines 745
Chapter 35 Summary of PTSP Configuration Statements 747
application-group-any 747
application-groups 747
applications (Services PTSP) 748
count-type 748
demux 749
forward-rule (Configuring) 750
forward-rule (Including in Rule) 751
from (Forward Rule) 751
from (Rule) 752
local-address 753
local-address-range 754
local-port-range 754
local-ports 755
local-prefix-list 755
match-direction (Services PTSP) 756
protocol 756
remote-address 757
remote-address-range 758
remote-port-range 758
remote-ports 759
remote-prefix-list 759
rule (Configuring) 760
rule (Including in Rule Set) 761
rule-set (Services PTSP) 761
services (PTSP) 762
term (Forward Rule) 763
term (Rule) 764
then (Forward Rule) 765
then (Rule) 766
Part 3 Dynamic Application Awareness for Junos OS
Chapter 36 Dynamic Application Awareness for Junos OS Overview 771
IDP Overview 772
APPID Overview 774
AACL Overview 775
L-PDF Overview 776
Configuring Multiple IDP Detectors 777
Best-Effort Application Identification of DPI-Serviced Flows 777
Features that Support Application-Level Filtering 777
Best-Effort Application Determination 778
APPID, AACL, and L-PDF Processing in Preconvergence Scenarios 778
Prior to a Final or Best-Effort Application Identification 778
Upon Best-Effort Application Identification 778
While Application Identification Is on a Best-Effort Basis 779
If a Flow Ends Before an Application Identification Is Made 779

Copyright © 2013, Juniper Networks, Inc.

If a Flow Ends While Application Identification on a Best-Effort Basis . . . 779

Chapter 37 Application Identification Configuration Guidelines . . . 781

Defining an Application Identification . . . 783
Configuring APPID Rules . . . 785
Using Stateful Firewall Rules to Identify Data Sessions . . . 786
Configuring Application Profiles . . . 788
Configuring Application Groups . . . 788
Application Identification for Nested Applications . . . 789
Disabling Application Identification for Nested Applications . . . 790
Configuring Global APPID Properties . . . 791
Configuring Automatic Download of Application Package Updates . . . 792
Configuring APPID Support for Heuristics . . . 792
Configuring APPID Support for Unidirectional Traffic . . . 793
Tracing APPID Operations . . . 794
Configuring the APPID Log Filename . . . 794
Configuring the Number and Size of APPID Log Files . . . 795
Configuring Access to the Log File . . . 795
Configuring a Regular Expression for Lines to Be Logged . . . 795
Configuring the Tracing Flags . . . 795
Examples: Configuring Application Identification Properties . . . 796

Chapter 38 Summary of Application Identification Configuration Statements . . . 799

address . . . 800
application . . . 801
application (Defining) . . . 801
application (Including in Rule) . . . 802
application-group . . . 802
application-groups (Services Application Identification) . . . 803
application-system-cache-timeout . . . 803
applications (Services Application Identification) . . . 804
automatic . . . 804
chain-order . . . 805
context . . . 805
destination (Services) . . . 806
direction . . . 806
disable . . . 807
disable (APPID Application) . . . 807
disable (APPID Application Group) . . . 807
disable (APPID Port Mapping) . . . 808
disable-global-timeout-override . . . 808
download . . . 809
enable-heuristics . . . 809
enable-asymmetic-traffic-processing . . . 810
enable-heuristics . . . 810
idle-timeout . . . 811
ignore-errors . . . 811
inactivity-non-tcp-timeout . . . 812
inactivity-tcp-timeout . . . 812

Copyright © 2013, Juniper Networks, Inc.

index (Nested Applications) . . . 813
index . . . 813
ip . . . 814
max-checked-bytes . . . 814
maximum-transactions . . . 815
member . . . 815
min-checked-bytes . . . 816
nested-application . . . 817
nested-application-settings . . . 818
no-application-identification . . . 818
no-application-system-cache . . . 819
no-clear-application-system-cache . . . 819
no-nested-application . . . 820
no-protocol-method . . . 820
no-signature-based . . . 821
order (Services Application Identification) . . . 821
pattern . . . 822
port-mapping . . . 822
port-range . . . 823
profile . . . 823
protocol . . . 824
hcm . . . 825
from . . . 825
host . . . 826
request-url . . . 826
term . . . 827
then . . . 827
url . . . 828
url-list . . . 828
url-rule . . . 829
url-rule-set . . . 829
rule . . . 830
rule (Configuring) . . . 830
rule (Including in Rule Set) . . . 831
rule-set (Services Application Identification) . . . 831
services (Application Identification) . . . 832
session-timeout . . . 833
session-timeout (Interfaces) . . . 833
session-timeout (Application Identification) . . . 833
signature . . . 834
signature-method-all-ports . . . 834
source . . . 835
support-uni-directional-traffic . . . 835
traceoptions (Application Identification) . . . 836
type . . . 837
type-of-service . . . 837
url . . . 838

Copyright © 2013, Juniper Networks, Inc.

Chapter 39 Application-Aware Access List Configuration Guidelines 839
Configuring AACL Rules 840
Configuring Match Direction for AACL Rules 841
Configuring Match Conditions in AACL Rules 841
Configuring Actions in AACL Rules 842
Logging AACL Flows Based on Application 843
Configuring AACL Rule Sets 844
Configuring Logging of AACL Flows 845
Example: Configuring AACL Rules 845
Chapter 40 Summary of AACL Configuration Statements 847
applications (Services AACL) 847
application-groups (Services AACL) 848
application-group-any 848
application-unknown 848
destination-address 849
destination-address-range 849
destination-prefix-list 850
from 850
match-direction 851
rule 852
rule-set (Services AACL) 853
services (AACL) 853
source-address (AACL) 854
source-address-range 854
source-prefix-list 855
term 856
then 857
Chapter 41 Local Policy Decision Function Configuration Guidelines 859
Configuring Statistics Profiles 859
Configuring an L-PDF Statistics Profile 860
Configuring an AACL Statistics Profile 861
Applying L-PDF Profiles to Service Sets 863
Tracing L-PDF Operations 864
Chapter 42 Summary of L-PDF Configuration Statements 867
aacl-fields 868
aacl-statistics-profile 869
application-aware-access-list-fields 870
file 871
local-policy-decision-function 872
policy-decision-statistics-profile 873
statistics (System Services) 874
traceoptions (Services Local Policy Decision Function) 875
Part 4 Encryption Services
Chapter 43 Encryption Overview 879
Encryption Overview 879

Copyright © 2013, Juniper Networks, Inc.

Chapter 44 Encryption Interfaces Configuration Guidelines 881
Configuring Encryption Interfaces 881
Specifying the Security Association Name for Encryption Interfaces 882
Configuring the MTU for Encryption Interfaces 882
Example: Configuring an Encryption Interface 882
Configuring Filters for Traffic Transiting the ES PIC 883
Traffic Overview 883
Configuring the Security Association 884
Configuring an Outbound Traffic Filter 885
Example: Configuring an Outbound Traffic Filter 885
Applying the Outbound Traffic Filter 886
Example: Applying the Outbound Traffic Filter 886
Configuring an Inbound Traffic Filter 886
Example: Configuring an Inbound Traffic Filter 887
Applying the Inbound Traffic Filter to the Encryption Interface 887
Example: Applying the Inbound Traffic Filter to the Encryption Interface 887
Configuring an ES Tunnel Interface for a Layer 3 VPN 888
Configuring ES PIC Redundancy 888
Example: Configuring ES PIC Redundancy 889
Configuring IPsec Tunnel Redundancy 889
Chapter 45 Summary of Encryption Configuration Statements 891
address (Interfaces) 891
backup-destination 891
backup-interface 892
destination (Interfaces) 893
es-options 894
family 895
filter 896
interfaces 896
ipsec-sa 897
source 897
tunnel 898
unit (Interfaces) 899
Part 5 Flow Monitoring and Discard Accounting Services
Chapter 46 Flow Monitoring and Discard Accounting Overview 903
Passive Flow Monitoring Overview 903
Active Flow Monitoring Overview 904
Chapter 47 Flow Monitoring and Discard Accounting Configuration Guidelines 907
Configuring Traffic Sampling 912
Minimum Configuration for Traffic Sampling 912
Configuring Traffic Sampling 913
Disabling Traffic Sampling 915
Sampling Once 915
Configuring Traffic Sampling Output 915
Traffic Sampling Output Format 916

Copyright © 2013, Juniper Networks, Inc.

Tracing Traffic Sampling Operations . . . 917
Traffic Sampling Examples . . . 917
Example: Sampling a Single SONET/SDH Interface . . . 917
Example: Sampling All Traffic from a Single IP Address . . . 918
Example: Sampling All FTP Traffic . . . 919
Configuring Flow Monitoring . . . 920
Configuring Flow-Monitoring Interfaces . . . 920
Configuring Flow-Monitoring Properties . . . 922
Directing Traffic to Flow-Monitoring Interfaces . . . 923
Exporting Flows . . . 923
Configuring Time Periods when Flow Monitoring is Active and Inactive . . . 923
Example: Configuring Flow Monitoring . . . 924
Example: Configuring Active Monitoring on Logical Systems . . . 925
Enabling Flow Aggregation . . . 928
Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd . . . 929
Configuring Flow Aggregation to Use Version 9 Flow Templates . . . 932
Configuring the Traffic to Be Sampled . . . 933
Configuring the Version 9 Template Properties . . . 933
Restrictions . . . 935
Fields Included in Each Template Type . . . 935
MPLS Sampling Behavior . . . 936
Verification . . . 937
Examples: Configuring Version 9 Flow Templates . . . 937
Configuring Sampling Instances . . . 941
Configuring Inline Sampling . . . 942
Configuring Inline Sampling on MX80 Routers . . . 946
Directing Replicated Flows to Multiple Flow Servers . . . 947
Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers . . . 947
Directing Replicated Version 9 Flow Aggregates to Multiple Servers . . . 948
Logging cflowd Flows Before Export . . . 949
Configuring Port Mirroring . . . 950
Configuring Tunnels . . . 952
Port Mirroring with Next-Hop Groups . . . 953
Configuring Inline Port Mirroring . . . 954
Filter-Based Forwarding with Multiple Monitoring Interfaces . . . 955
Restrictions . . . 955
Configuring Port Mirroring on Services Interfaces . . . 956
Examples: Configuring Port Mirroring . . . 957
Load Balancing Among Multiple Monitoring Interfaces . . . 963
Configuring Discard Accounting . . . 967
Enabling Passive Flow Monitoring . . . 968
Passive Flow Monitoring for MPLS Encapsulated Packets . . . 970
Removing MPLS Labels from Incoming Packets . . . 970
Example: Enabling IPv4 Passive Flow Monitoring . . . 972
Example: Enabling IPv6 Passive Flow Monitoring . . . 974
Configuring Services Interface Redundancy with Flow Monitoring . . . 975

Copyright © 2013, Juniper Networks, Inc.

Chapter 48 Summary of Flow-Monitoring Configuration Statements 977

accounting 978
address (Interfaces) 979
aggregate-export-interval 979
aggregation 980
autonomous-system-type 981
cflowd 982
cflowd (Discard Accounting) 982
cflowd (Flow Monitoring) 983
core-dump 984
destination (Tunnel Remote End) 984
disable (Forwarding Options) 985
disable-all-instances 985
engine-id (Forwarding Options) 986
engine-type 987
extension-service 988
export-format 989
family 990
family (Interfaces) 990
family (Monitoring) 991
family (Port Mirroring) 992
family (Sampling) 993
file 995
file (Sampling) 995
file (Trace Options) 995
filename 996
files 996
filter 997
flow-active-timeout 998
flow-export-rate 999
flow-control-options 999
flow-export-destination 1000
flow-inactive-timeout 1001
flow-monitoring 1002
flow-server 1003
forwarding-options 1004
inline-jflow 1005
input 1006
input (Port Mirroring) 1006
input (Sampling) 1006
input-interface-index 1007
instance 1008
instance (Port Mirroring) 1008
instance (Sampling) 1009
interface 1011
interface (Accounting or Sampling) 1011
interface (Monitoring) 1012
interface (Port Mirroring) 1012
interfaces 1013

Copyright © 2013, Juniper Networks, Inc.

ipv4-template . . . 1013
ipv6-template . . . 1013
label-position . . . 1014
local-dump . . . 1014
match . . . 1015
maximum-packet-length . . . 1016
max-packets-per-second . . . 1017
monitoring . . . 1018
mpls-ipv4-template . . . 1019
mpls-template . . . 1019
multiservice-options . . . 1020
next-hop (Forwarding Options) . . . 1020
next-hop-group . . . 1021
    next-hop-group (Forwarding Options) . . . 1021
    next-hop-group (Port Mirroring) . . . 1022
no-core-dump . . . 1022
no-filter-check . . . 1022
no-local-dump . . . 1022
no-remote-trace (Trace Options) . . . 1023
no-stamp . . . 1023
no-syslog . . . 1023
no-world-readable . . . 1023
option-refresh-rate . . . 1024
output . . . 1025
    output (Accounting) . . . 1025
    output (Monitoring) . . . 1026
    output (Port Mirroring) . . . 1026
    output (Sampling) . . . 1027
output-interface-index . . . 1028
passive-monitor-mode . . . 1029
pop-all-labels . . . 1030
port . . . 1031
port-mirroring . . . 1032
rate . . . 1033
receive-options-packets . . . 1033
receive-ttl-exceeded . . . 1034
required-depth . . . 1034
run-length . . . 1035
sample-once . . . 1035
sampling . . . 1036
    sampling (Forwarding Options) . . . 1037
    sampling (Interfaces) . . . 1039
services . . . 1040
size . . . 1040
source-address (Forwarding Options) . . . 1041
stamp . . . 1041
syslog . . . 1042

Copyright © 2013, Juniper Networks, Inc.

template . . . 1043
template (Forwarding Options) . . . 1043
template (Services) . . . 1044
template-refresh-rate . . . 1045
traceoptions (Forwarding Options) . . . 1045
unit . . . 1046
version . . . 1047
version9 . . . 1048
version9 (Forwarding Options) . . . 1048
version9 (Services) . . . 1049
version-ipfix . . . 1050
version-ipfix (Forwarding Options) . . . 1050
version-ipfix (Services) . . . 1051
world-readable . . . 1051

Chapter 49 Flow Collection Configuration Guidelines . . . 1053

Configuring Flow Collection . . . 1055
Configuring Destination FTP Servers for Flow Records . . . 1055
Configuring a Packet Analyzer . . . 1055
Configuring File Formats . . . 1056
Configuring Interface Mappings . . . 1056
Configuring Transfer Logs . . . 1057
Configuring Retry Attempts . . . 1057
Sending cflowd Records to Flow Collector Interfaces . . . 1058
Configuring Flow Collection Mode and Interfaces on Services PICs . . . 1058
Example: Configuring Flow Collection . . . 1058

Chapter 50 Summary of Flow Collection Configuration Statements . . . 1065

analyzer-address . . . 1065
analyzer-id . . . 1066
archive-sites . . . 1066
collector . . . 1067
data-format . . . 1067
destinations . . . 1068
filename-prefix . . . 1068
file-specification . . . 1069
file-specification (File Format) . . . 1069
file-specification (Interface Mapping) . . . 1069
flow-collector . . . 1070
ftp . . . 1072
ftp (Flow Collector Files) . . . 1073
ftp (Transfer Log Files) . . . 1074
interface-map . . . 1074
maximum-age . . . 1075
name-format . . . 1076
password . . . 1078
password (Flow Collector File Servers) . . . 1078
password (Transfer Log File Servers) . . . 1078
retry (Services Flow Collector) . . . 1079
retry-delay . . . 1079

Copyright © 2013, Juniper Networks, Inc.

transfer . . . 1080
transfer-log-archive . . . 1080
username (Services) . . . 1081
variant . . . 1081

Chapter 51 Dynamic Flow Capture Configuration Guidelines . . . 1083

Dynamic Flow Capture Architecture . . . 1083
Liberal Sequence Windowing . . . 1084
Intercepting IPv6 Flows . . . 1084
Configuring Dynamic Flow Capture . . . 1085
Configuring the Capture Group . . . 1086
Configuring the Content Destination . . . 1086
Configuring the Control Source . . . 1087
Configuring the DFC PIC Interface . . . 1088
Configuring the Firewall Filter . . . 1089
Configuring System Logging . . . 1090
Configuring Tracing Options for Dynamic Flow Capture Events . . . 1090
Configuring Thresholds . . . 1090
Limiting the Number of Duplicates of a Packet . . . 1091
Example: Configuring Dynamic Flow Capture . . . 1092

Chapter 52 Flow-Tap Configuration Guidelines . . . 1095

Flow-Tap Architecture . . . 1096
Configuring the Flow-Tap Service . . . 1097
Configuring the Flow-Tap Interface . . . 1097
Strengthening Flow-Tap Security . . . 1098
Restrictions on Flow-Tap Services . . . 1099
Configuring FlowTapLite . . . 1100
Examples: Configuring Flow-Tap Services . . . 1101

Chapter 53 Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements . . . 1103

address (Services Dynamic Flow Capture) . . . 1103
allowed-destinations . . . 1104
capture-group . . . 1105
content-destination . . . 1106
control-source . . . 1107
duplicates-dropped-periodicity . . . 1107
dynamic-flow-capture . . . 1108
flow-tap . . . 1109
g-duplicates-dropped-periodicity . . . 1110
g-max-duplicates . . . 1111
hard-limit . . . 1111
hard-limit-target . . . 1112
input-packet-rate-threshold . . . 1112
interface (Services Flow Tap) . . . 1113
interfaces (Services Dynamic Flow Capture) . . . 1113
max-duplicates . . . 1114
minimum-priority . . . 1114
no-syslog . . . 1115

Copyright © 2013, Juniper Networks, Inc.

notification-targets . . . 1115
pic-memory-threshold . . . 1116
service-port . . . 1116
services (Dynamic Flow Capture) . . . 1117
shared-key . . . 1117
soft-limit . . . 1118
soft-limit-clear . . . 1118
source-addresses . . . 1119
traceoptions (Dynamic Flow Capture) . . . 1120
ttl . . . 1120

Part 6 Link and Multilink Services

Chapter 54 Link and Multilink Services Overview . . . 1123

Link and Multilink Services Overview . . . 1123

Chapter 55 Link and Multilink Services Configuration Guidelines . . . 1127

Multilink and Link Services PICs Overview . . . 1129
Multilink Interfaces on Channelized MICs Overview . . . 1131
Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links . . . 1133
Configuring the Number of Bundles on Link Services PICs . . . 1134
Configuring the Links in a Multilink or Link Services Bundle . . . 1135
Multilink and Link Services Logical Interface Configuration Overview . . . 1136
Default Settings for Multilink and Link Services Logical Interfaces . . . 1137
Configuring Encapsulation for Multilink and Link Services Logical Interfaces . . . 1138
Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces . . . 1139
Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces . . . 1140
Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces . . . 1141
Configuring MRRU on Multilink and Link Services Logical Interfaces . . . 1141
Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces . . . 1142
Configuring DLCIs on Link Services Logical Interfaces . . . 1143
Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles . . . 1143
Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles . . . 1143
Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces . . . 1144
Configuring LFI with DLCI Scheduling . . . 1145
Example: Configuring LFI with DLCI Scheduling . . . 1145
Configuring Link Services Physical Interfaces . . . 1147
Default Settings for Link Services Interfaces . . . 1147
Configuring Encapsulation for Link Services Physical Interfaces . . . 1148
Configuring Acknowledgment Timers on Link Services Physical Interfaces . . . 1148
Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 . . . 1149
Configuring Keepalives on Link Services Physical Interfaces . . . 1150

Copyright © 2013, Juniper Networks, Inc.

Configuring CoS on Link Services Interfaces 1151
CoS for Link Services Interfaces on M Series and T Series Routers 1151
Example: Configuring CoS on Link Services Interfaces 1152
Examples: Configuring Multilink Interfaces 1156
Example: Configuring a Multilink Interface with MLPPP 1156
Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces 1157
Example: Configuring a Multilink Interface with MLFR FRF.15 1158
Examples: Configuring Link Interfaces 1159
Example: Configuring a Link Services Interface with Two Links 1160
Example: Configuring a Link Services Interface with MLPPP 1161
Example: Configuring a Link Services Interface with MLFR FRF.15 1162
Example: Configuring a Link Services PIC with MLFR FRF.16 1162
Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types 1163
Example: Configuring an MLPPP Bundle 1169
Example: Configuring Multilink Frame Relay FRF.15 1172
Example: Configuring Multilink Frame Relay FRF.16 1175
Example: Configuring Link Interfaces on Channelized MICs 1179

Chapter 56 Summary of Multilink and Link Services Configuration Statements 1189

acknowledge-retries 1189
acknowledge-timer 1190
action-red-differential-delay 1191
address (Interfaces) 1192
bundle 1192
destination (Interfaces) 1193
disable-mlppp-inner-ppp-pfc 1194
dlci 1194
drop-timeout 1195
encapsulation 1196
encapsulation (Logical Interface) 1196
encapsulation (Physical Interface) 1197
family 1198
fragment-threshold 1199
hello-timer 1200
interfaces 1200
interleave-fragments 1201
lmi-type 1201
minimum-links 1202
mlfr-uni-nni-bundle-options 1203
mrru 1204
mtu 1205
multicast-dlci 1205
n391 1206
n392 1206
n393 1207
red-differential-delay 1207
short-sequence 1208

Copyright © 2013, Juniper Networks, Inc.

t391 . . . 1208
t392 . . . 1209
unit (Interfaces) . . . 1210
yellow-differential-delay . . . 1211

Part 7 Real-Time Performance Monitoring Services

Chapter 57 Real-Time Performance Monitoring Services Overview . . . 1215

Real-Time Performance Monitoring Services Overview . . . 1215

Chapter 58 Real-Time Performance Monitoring Configuration Guidelines . . . 1217

Configuring BGP Neighbor Discovery Through RPM . . . 1218
Configuring Real-Time Performance Monitoring . . . 1220
Configuring RPM Probes . . . 1221
Configuring RPM Receiver Servers . . . 1224
Limiting the Number of Concurrent RPM Probes . . . 1225
Configuring RPM Timestamping . . . 1225
Configuring TWAMP . . . 1228
Configuring TWAMP Interfaces . . . 1229
Configuring TWAMP Servers . . . 1229
Enabling RPM for the Services SDK . . . 1230
Examples: Configuring BGP Neighbor Discovery Through RPM . . . 1231
Examples: Configuring Real-Time Performance Monitoring . . . 1232

Chapter 59 Summary of Real-Time Performance Monitoring Configuration Statements . . . 1237

authentication-mode . . . 1237
bgp . . . 1238
client-list . . . 1239
data-fill . . . 1239
data-size . . . 1240
destination-interface . . . 1241
destination-port . . . 1242
dscp-code-point . . . 1243
hardware-timestamp . . . 1244
history-size . . . 1244
inactivity-timeout (Services RPM) . . . 1245
logical-system . . . 1245
max-connection-duration . . . 1246
maximum-connections . . . 1246
maximum-connections-per-client . . . 1247
maximum-sessions . . . 1247
maximum-sessions-per-connection . . . 1248
moving-average-size . . . 1248
one-way-hardware-timestamp . . . 1249
port . . . 1250
port (RPM) . . . 1250
port (TWAMP) . . . 1250
probe . . . 1251
probe-count . . . 1252

Copyright © 2013, Juniper Networks, Inc.

probe-interval . . . 1252
probe-limit . . . 1253
probe-server . . . 1253
probe-type . . . 1254
routing-instance . . . 1255
routing-instances . . . 1255
rpm . . . 1256
server . . . 1257
server-inactivity-timeout . . . 1257
services (RPM) . . . 1258
source-address . . . 1258
target (Services RPM) . . . 1259
tcp . . . 1259
test . . . 1260
test-interval . . . 1261
thresholds . . . 1262
traps . . . 1263
twamp . . . 1264
twamp-server . . . 1265
udp . . . 1265

Part 8 Tunnel Services

Chapter 60 Tunnel Services Overview . . . 1269

Tunnel Services Overview . . . 1269
GRE Keepalive Time Overview . . . 1271

Chapter 61 Tunnel Interfaces Configuration Guidelines . . . 1273

Configuring Unicast Tunnels . . . 1273
Configuring a Key Number on GRE Tunnels . . . 1276
Enabling Fragmentation on GRE Tunnels . . . 1276
Specifying an MTU Setting for the Tunnel . . . 1277
Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header . . . 1277
Configuring Packet Reassembly . . . 1278
Configuring GRE Keepalive Time . . . 1279
Restricting Tunnels to Multicast Traffic . . . 1280
Configuring Logical Tunnel Interfaces . . . 1281
Connecting Logical Systems . . . 1281
Configuring Tunnel Interfaces for Routing Table Lookup . . . 1282
Configuring Virtual Loopback Tunnels for VRF Table Lookup . . . 1283
Configuring PIM Tunnels . . . 1284
Configuring IPv6-over-IPv4 Tunnels . . . 1285
Configuring IPv4-over-IPv6 Tunnels . . . 1285
Configuring Dynamic Tunnels . . . 1286
Configuring Tunnel Interfaces on MX Series Routers . . . 1286
Configuring Tunnel Interfaces on T4000 Routers . . . 1287
Examples: Configuring Unicast Tunnels . . . 1288
Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup . . . 1289
Example: Configuring an IPv6-over-IPv4 Tunnel . . . 1290
Example: Configuring an IPv4-over-IPv6 Tunnel . . . 1291

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring Logical Tunnels 1292
Example: Configuring Keepalive for a GRE Interface 1293

Chapter 62 Summary of Tunnel Services Configuration Statements 1295

allow-fragmentation 1295
backup-destination 1296
copy-tos-to-outer-ip-header 1296
destination 1297
destination (Tunnel Remote End) 1297
destination (Routing Instance) 1297
destination-networks 1298
do-not-fragment 1299
dynamic-tunnels 1300
hold-time (OAM) 1301
interfaces 1301
keepalive-time 1302
key 1303
multicast-only 1303
peer-unit 1304
reassemble-packets 1304
routing-instance 1305
routing-instances 1305
routing-options 1306
source 1306
source-address 1307
ttl 1307
tunnel 1308
unit (Interfaces) 1309

Part 9 Index

Index 1313
Index of Statements and Commands 1337

Copyright © 2013, Juniper Networks, Inc.

List of Figures

Part 2 Adaptive Services

Chapter 3 **Adaptive Services Overview 37**

Figure 1: Packet Flow Through the Adaptive Services or MultiServices PIC 45
Figure 2: Dynamic NAT Flow 51
Figure 3: Stateful NAT64 Flow 52
Figure 4: DS-Lite Flow 52
Figure 5: 6rd Softwire Flow 56

Chapter 14 **IPsec Services Configuration Guidelines 219**

Figure 6: IPsec Dynamic Endpoint Tunneling Topology 262

Part 4 Encryption Services

Chapter 44 **Encryption Interfaces Configuration Guidelines 881**

Figure 7: Example: IPsec Tunnel Connecting Security Gateways 883
Figure 8: IPsec Tunnel Redundancy 889

Part 5 Flow Monitoring and Discard Accounting Services

Chapter 46 **Flow Monitoring and Discard Accounting Overview 903**

Figure 9: Passive Monitoring Application Topology 904
Figure 10: Active Monitoring Configuration Topology 906

Chapter 47 **Flow Monitoring and Discard Accounting Configuration Guidelines 907**

Figure 11: Configure Sampling Rate 914

Chapter 49 **Flow Collection Configuration Guidelines 1053**

Figure 12: Flow Collector Interface Topology Diagram 1059

Chapter 51 **Dynamic Flow Capture Configuration Guidelines 1083**

Figure 13: Dynamic Flow Capture Topology 1084

Chapter 52 **Flow-Tap Configuration Guidelines 1095**

Figure 14: Flow-Tap Topology 1097

Part 6 Link and Multilink Services

Chapter 55 **Link and Multilink Services Configuration Guidelines 1127**

Figure 15: Multilink Interface Configuration 1135
Figure 16: Configuring MLPPP and LFI on Serial Links 1169

Copyright © 2013, Juniper Networks, Inc.

Part 8 Tunnel Services

Chapter 61 Tunnel Interfaces Configuration Guidelines 1273

Figure 17: IPv6 Tunnel Connecting Two IPv4 Networks Across an IPv6 Network 1291

Copyright © 2013, Juniper Networks, Inc.

List of Tables

About This Guide xlv

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

Part 2 Adaptive Services

Chapter 3 Adaptive Services Overview 37

Table 3: AS and Multiservices PIC Services by Service Package, PIC, and Platform 40
Table 4: Statement Equivalents for ES and AS Interfaces 58

Chapter 4 Applications Configuration Guidelines 73

Table 5: Application Protocols Supported by Services Interfaces 75
Table 6: Network Protocols Supported by Services Interfaces 77
Table 7: ICMP Codes and Types Supported by Services Interfaces 78
Table 8: Port Names Supported by Services Interfaces 79
Table 9: ALGs Supported by the Junos OS 85
Table 10: RealAudio Product Port Usage 91
Table 11: Supported RPC Services 92

Chapter 6 Stateful Firewall Services Configuration Guidelines 131

Table 12: IP Option Values 135

Chapter 11 Summary of Load Balancing Configuration Statements 173

Table 13: Behavior of Member Interface After One Multiservices PIC Fails 179
Table 14: Behavior of Member Interface After Two Multiservices PICs Fail 180

Chapter 14 IPsec Services Configuration Guidelines 219

Table 15: Default IKE and IPsec Proposals for Dynamic Negotiations 255

Chapter 16 Layer 2 Tunneling Protocol Services Configuration Guidelines 309

Table 16: System Log Message Severity Levels 317

Chapter 24 Service Set Configuration Guidelines 467

Table 17: System Log Message Severity Levels 479
Table 18: Adaptive Services Tracing Flags 483

Chapter 26 Service Interface Configuration Guidelines 513

Table 19: System Log Message Severity Levels 518

Part 6 Link and Multilink Services

Chapter 55 Link and Multilink Services Configuration Guidelines 1127

Table 20: Multilink and Link Services PIC Capacities 1129

Copyright © 2013, Juniper Networks, Inc.

Table 21: Multilink and Link Services Logical Interface Statements 1137
Table 22: Link Services Physical Interface Statements for MLFR FRF.16 1147
Table 23: Link Services CoS Queues 1151
Table 24: Link Services Bundle 1160

Part 8 Tunnel Services

Chapter 60 Tunnel Services Overview 1269

Table 25: Tunnel Interface Types 1269

Chapter 61 Tunnel Interfaces Configuration Guidelines 1273

Table 26: Methods for Configuring Egress Filtering 1283

Copyright © 2013, Juniper Networks, Inc.

About This Guide

This preface provides the following guidelines for using the *Junos® OS Services Interfaces Configuration Guide*:

- Junos Documentation and Release Notes on page xlv
- Objectives on page xlvi
- Audience on page xlvi
- Supported Platforms on page xlvi
- Using the Indexes on page xlvii
- Using the Examples in This Manual on page xlvii
- Documentation Conventions on page xlviii
- Documentation Feedback on page l
- Requesting Technical Support on page l

Junos Documentation and Release Notes

For a list of related Junos documentation, see http://www.juniper.net/techpubs/software/junos/.

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

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

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at http://www.juniper.net/books.

Copyright © 2013, Juniper Networks, Inc.

Objectives

This guide provides an overview of the services interfaces provided by Junos OS and describes how to configure these properties on the router.



NOTE: For additional information about the Junos OS—either corrections to or information that might have been omitted from this guide—see the software release notes at http://www.juniper.net/ .

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Supported Platforms

For the features described in this manual, the Junos OS currently supports the following platforms:

- J Series
- M Series

Copyright © 2013, Juniper Networks, Inc.

- MX Series
- SRX Series
- T Series
- EX Series

Using the Indexes

This reference contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines,* refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   ```

Copyright © 2013, Juniper Networks, Inc.

```
interfaces {
   fxp0 {
      disable;
      unit 0 {
         family inet {
            address 10.0.0.1/24;
         }
      }
   }
}
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xlix defines notice icons used in this guide.

Copyright © 2013, Juniper Networks, Inc.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the**configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the**[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub <default-metric** ***metric*****>;**

Copyright © 2013, Juniper Networks, Inc.

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

Convention	Description	Examples
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

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

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to **techpubs-comments@juniper.net**, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract,

Copyright © 2013, Juniper Networks, Inc.

or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 1

Overview

- Services Interfaces Overview on page 3
- Services Interfaces Configuration Statements on page 5

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 1

Services Interfaces Overview

Interfaces used in router networks fall into two categories:

- Networking interfaces, such as Ethernet and SONET interfaces, that primarily provide traffic connectivity. For more information on these interfaces, see the Junos® OS Network Interfaces.
- Services interfaces that provide specific capabilities for manipulating traffic before it is delivered to its destination.

This chapter includes the following sections:

- Services PIC Types on page 3
- Supported Platforms on page 4

Services PIC Types

Services interfaces enable you to add services to your network incrementally. The Juniper Networks Junos OS supports the following services PICs:

- Adaptive services interfaces (Adaptive Services [AS] PICs and Multiservices PICs)—Enable you to perform multiple services on the same PIC by configuring a set of services and applications. The AS and Multiservices PICs offer a special range of services you configure in one or more service sets: stateful firewalls, Network Address Translation (NAT), intrusion detection service (IDS), class-of-service functionality, and IP Security (IPsec). You can also configure voice services and Layer 2 Tunneling Protocol (L2TP) services. For more information about these services, see "Adaptive Services Overview" on page 37.

  

  NOTE: On Juniper Networks MX Series 3D Universal Edge Routers, the Multiservices DPC provides essentially the same capabilities as the Multiservices PIC. The interfaces on both platforms are configured in the same way.

- ES PIC—Provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations,

Copyright © 2013, Juniper Networks, Inc.

and support for digital certificates. For more information about encryption interfaces, see "Configuring Encryption Interfaces" on page 881.

- Monitoring Services PICs—Enable you to monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to perform the following tasks:
  - Gather and export detailed information about IPv4 traffic flows between source and destination nodes in your network.
  - Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
  - Perform discard accounting on an incoming traffic flow.
  - Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
  - Direct filtered traffic to different packet analyzers and present the data in its original format.

  For more information about flow monitoring interfaces, see Flow Monitoring.

- Multilink Services and Link Services PICs—Enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members. The Junos OS supports two services PICs based on the Multilink Protocol: the Multilink Services PIC and the Link Services PIC. For more information about multilink and link services interfaces, see Link and Multilink Properties.
- Tunnel Services PIC—By encapsulating arbitrary packets inside a transport protocol, provides a private, secure path through an otherwise public network. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and MPLS. For more information about tunnel interfaces, see Tunnel Properties.

Supported Platforms

For information about which platforms support Adaptive Services and MultiServices PICs and their features, see "Enabling Service Packages" on page 39.

For information about PIC support on a specific Juniper Networks M Series Multiservice Edge Router or T Series Core Router, see the appropriate *PIC Guide* for the platform.

For information about MS-DPC support on a specific MX Series router, see the appropriate *DPC Guide* for the platform.

For information about services supported on Juniper Networks SRX Series Services Gateways and J Series Services Routers, see the *Junos OS Feature Support Reference for SRX Series and J Series Devices*.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 2

Services Interfaces Configuration Statements

This chapter shows the complete configuration statement hierarchies for configuring services interfaces. It lists all the statements that pertain to configuring services and shows their level in the configuration hierarchy. When you are configuring the Junos OS, your current hierarchy level is shown in the banner on the line preceding the **user@host#** prompt.

For a complete list of the Junos configuration statements, see the *Junos OS Hierarchy and RFC Reference*.

This chapter is organized as follows:

- [edit applications] Hierarchy Level on page 5
- [edit forwarding-options] Hierarchy Level on page 6
- [edit interfaces] Hierarchy Level on page 8
- [edit logical-systems] Hierarchy Level on page 12
- [edit services] Hierarchy Level on page 12

[edit applications] Hierarchy Level

To configure application protocols, include the following statements at the **[edit applications]** hierarchy level of the configuration:

```
application application-name {
    application-protocol protocol-name;
    destination-port port-number;
    icmp-code value;
    icmp-type value;
    inactivity-timeout value;
    learn-sip-register;
    protocol type;
    rpc-program-number number;
    sip-call-hold-timeout seconds;
    snmp-command command;
    source-port port-number;
    ttl-threshold value;
    uuid hex-value;
```

Copyright © 2013, Juniper Networks, Inc.

```
}
application-set application-set-name {
  application application-name;
}
```

[edit forwarding-options] Hierarchy Level

To configure flow monitoring and accounting properties, include the following statements at the **[edit forwarding-options]** hierarchy level:



NOTE: For the complete **[edit forwarding-options]** hierarchy, see the Routing Policy Configuration Guide. This listing includes only the statements used in flow monitoring and accounting services.

```
accounting name {
  output {
    aggregate-export-interval seconds;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface interface-name {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
monitoring name {
  family inet {
    output {
      cflowd hostname port port-number;
      export-format format;
      flow-active-timeout seconds;
      flow-export-destination {
        collector-pic;
      }
      flow-inactive-timeout seconds;
      interface interface-name {
        engine-id number;
        engine-type number;
```


Copyright © 2013, Juniper Networks, Inc.

```
            input-interface-index number;
            output-interface-index number;
            source-address address;
          }
        }
      }
    next-hop-group group-name {
      interface interface-name {
        next-hop address;
      }
    }
    port-mirroring {
      input {
        rate rate;
        run-length number;
      }
      family (inet | inet6) {
        input {
          rate rate;
          run-length number;
          }
          output {
          interface interface-name {
            next-hop address;
          }
          no-filter-check;
        }
      }
      traceoptions {
        file filename {
          files number;
          size bytes;
          (world-readable | no-world-readable);
        }
      }
    }
    sampling {
      disable;
      family (inet | inet6 | mpls) {
        max-packets-per-second number;
        rate number;
        run-length number;
      }
      input {
        rate number;
        run-length number;
      }
      output {
        aggregate-export-interval seconds;
        cflowd hostname {
          aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
              caida-compliant;
```

Copyright © 2013, Juniper Networks, Inc.

```
            }
            source-prefix;
          }
          autonomous-system-type (origin | peer);
          version9 {
            template template-name;
          }
          (local-dump | no-local-dump);
          port port-number;
          source-address address;
          version format;
        }
        file {
          disable;
          filename filename;
          files number;
          size bytes;
          (stamp | no-stamp);
          (world-readable | no-world-readable);
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface interface-name {
          engine-id number;
          engine-type number;
          source-address address;
        }
      }
      traceoptions (Forwarding Options) {
        file filename {
          files number;
          size bytes;
          (world-readable | no-world-readable);
        }
      }
    }
```

[edit interfaces] Hierarchy Level

To configure services interfaces, include the following statements at the **[edit interfaces]** hierarchy level of the configuration. The statements can also be configured at the **[edit logical-systems *logical-system-name* interfaces]** hierarchy level.



NOTE: For the complete **[edit interfaces]** hierarchy, see the Junos® OS Network Interfaces. This listing includes only the statements used in configuring services.

```
[edit interfaces]
interface-name {
  (atm-options | fastether-options | gigether-options | sonet-options) {
    mpls {
      pop-all-labels {
        required-depth number;
```


Copyright © 2013, Juniper Networks, Inc.

```
      }
    }
  }
  encapsulation type;
  lsq-failure-options {
    no-termination-request;
    trigger-link-failure interface-name;
  }
  mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    cisco-interoperability send-lip-remove-link-for-link-reject;
    hello-timer milliseconds;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 number;
    t392 number;
    yellow-differential-delay milliseconds;
    encapsulation type;
  }
  passive-monitor-mode;
  unit logical-unit-number {
    clear-dont-fragment-bit;
    compression {
      rtp {
        f-max-period number;
        maximum-contexts number <force>;
        port {
          minimum port-number;
          maximum port-number;
        }
        queues [ queue-numbers ];
      }
    }
    compression-device interface-name;
    copy-tos-to-outer-ip-header;
    disable-mlppp-inner-ppp-pfc;
    dlci dlci-identifier;
    drop-timeout milliseconds;
    dial-options {
      ipsec-interface-id name;
      l2tp-interface-id name;
      (dedicated | shared);
    }
    encapsulation type;
    family family {
      accounting {
        destination-class-usage;
```

Copyright © 2013, Juniper Networks, Inc.

```
        source-class-usage direction;
      }
      address address {
        destination address;
      }
      bundle (ml-fpc/pic/port | ls-fpc/pic/port);
      ipsec-sa ipsec-sa;
      multicast-only;
      receive-options-packets;
      receive-ttl-exceeded;
      sampling direction;
      service {
        input {
          service-set service-set-name <service-filter filter-name>;
          post-service-filter filter-name;
        }
        output {
          service-set service-set-names <service-filter filter-name>;
        }
      }
    }
    fragment-threshold bytes;
    interleave-fragments;
    minimum-links number;
    mrru bytes;
    multicast-dlci dlci-identifier;
    peer-unit unit-number;
    reassemble-packets;
    rpm ;
    service-domain (inside | outside);
    short-sequence;
    tunnel {
      allow-fragmentation;
      backup-destination address;
      destination destination-address;
      do-not-fragment;
      key number;
      routing-instance {
        destination routing-instance-name;
      }
      source-address address;
      ttl number;
    }
    twamp-server;
  }
  multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
    flow-control-options {
      down-on-flow-control;
      dump-on-flow-control;
      reset-on-flow-control;
    }
  }
  services-options {
    cgn-pic;
```


Copyright © 2013, Juniper Networks, Inc.

```
        disable-global-timeout-override;
        ignore-errors <alg> <tcp>;
        inactivity-non-tcp-timeout seconds;
        inactivity-tcp-timeout seconds;
        inactivity-timeout seconds;
        open-timeout seconds;
        session-limit {
          maximum number;
          rate new-sessions-per-second;
        }
        session-timeout seconds;
        syslog {
          host hostname {
            facility-override facility-name;
            log-prefix prefix-value;
            port port-number;
            services severity-level;
          }
          message-rate-limit messages-per-second;
        }
        tcp-tickles tcp-tickles;
      }
    }
    rlsqnumber {
      redundancy-options {
        hot-standby | warm-standby;
        primary lsq-fpc/pic/port;
        secondary lsq-fpc/pic/port;
      }
    }
    rlsqnumber:number {
      redundancy-options {
        hot-standby | warm-standby;
        primary lsq-fpc/pic/port;
        secondary lsq-fpc/pic/port;
        }
      }
      encapsulation multilink-frame-relay-uni-nni;
        unit logical-unit-number {
        encapsulation multilink-frame-relay-end-to-end ;
        }
      }
    }
    rspnumber {
      redundancy-options {
        primary sp-fpc/pic/port;
        secondary sp-fpc/pic/port;
      }
    }
    so-fpc/pic/port {
      unit logical-unit-number {
        passive-monitor-mode;
      }
    }
```

Copyright © 2013, Juniper Networks, Inc.

[edit logical-systems] Hierarchy Level

The following lists the statements that can be configured at the **[edit logical-systems]** hierarchy level that are documented in this manual. For more information about logical systems, see the Junos OS Routing Protocols Configuration Guide.

```
logical-system-name {
    interfaces interface-name {
        interface-configuration;
    }
}
```

[edit services] Hierarchy Level

To configure services, include the following statements at the **[edit services]** hierarchy level of the configuration:



NOTE: For the complete **[edit services]** hierarchy, see the *Junos OS Hierarchy and RFC Reference*. This listing includes only the statements documented in this manual; additional statements are documented in the Junos OS Subscriber Management, Release 13.1.

```
aacl {
    rule rule-name {
        match-direction (input | output | input-output);
        term term-name {
            from {
                application-group-any;
                application-groups [ application-group-names ];
                application-unknown;
                applications [ application-names ];
                destination-address address <any-unicast>;
                destination-address-range low minimum-value high maximum-value;
                destination-prefix-list list-name;
                source-address address <any-unicast>;
                source-address-range low minimum-value high maximum-value;
                source-prefix-list list-name;
            }
            then {
                (accept | discard);
                count (application | application-group | application-group-any | none);
                forwarding-class class-name;
                policer policer-name;
            }
        }
    }
    rule-set rule-set-name {
        [ rule rule-names ];
    }
}
adaptive-services-pics {
```


Copyright © 2013, Juniper Networks, Inc.

```
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable> <match
        regex>;
      flag flag;
      no-remote-trace;
    }
  }
  application-identification {
    application application-name {
      disable;
      enable-heuristics;
      idle-timeout seconds;
      index number;
      session-timeout seconds;
      type type;
      type-of-service service-type;
      port-mapping {
        port-range {
          tcp (port | range);
          udp (port | range);
        }
        disable;
      }
    }
    application-group group-name {
      disable;
      application-groups {
        name [application-group-name];
      }
      applications {
        name [application-name];
      }
      index number;
    }
    application-system-cache-timeout seconds;
    max-checked-bytes bytes;
    min-checked-bytes bytes;
    nested-application
    nested-application-settings
    no-application-identification;
    no-application-system-cache;
    no-clear-application-system-cache;
    no-signature-based;
    profile profile-name {
      [ rule-set rule-set-name ];
    }
    rule rule-name {
      disable;
      address address-name {
        destination {
          ip address</prefix-length>;
          port-range {
            tcp [ ports-and-port-ranges ];
            udp [ ports-and-port-ranges ];
          }
        }
```

```
            source {
              ip address</prefix-length>;
              port-range {
                tcp [ ports-and-port-ranges ];
                udp [ ports-and-port-ranges ];
              }
            }
            order number;
          }
          application application-name;
        }
        rule-set rule-set-name {
          rule application-rule-name;
        }
        traceoptions {
          file filename <files number> <match regex> <size size> <world-readable |
            no-world-readable>;
          flag flag;
          no-remote-trace;
        }
      }
      border-signaling-gateway {
        gateway gateway-name {
          admission-control admission-control-profile {
            dialogs {
              maximum-concurrent number;
              committed-attempts-rate dialogs-per-second;
              committed-burst-size number-of-dialogs;
            }
            transactions {
              maximum-concurrent number;
              committed-attempts-rate transactions-per-second;
              committed-burst-size number-of-transactions;
            }
          }
          embedded-spdf {
            service-class service-class-name {
              term term-name {
                from {
                  media-type (any-media | audio | video);
                }
                then {
                  committed-burst-size bytes;
                  committed-information-rate bytes-per-second;
                  dscp (alias | do-not-change | dscp-value);
                  reject;
                }
              }
            }
          }
          service-point service-point-name {
            default-media-realm
            service-interface interface-name.unit-number;
            service-point-type service-point-type;
            service-policies {
              new-call-usage-input-policies [policy-and-policy-set-names];
```


Copyright © 2013, Juniper Networks, Inc.

```
    new-call-usage-output-policies [policy-and-policy-set-names];
    new-transaction-input-policies [policy-and-policy-set-names];
    new-transaction-output-policies[policy-and-policy-set-names];
  }
  transport-details <port port-number> <ip-address ip-address> <tcp> <udp>;
}
sip {
  message-manipulation-rules {
    manipulation-rule rule-name {
      actions {
        sip-header header-field-name {
          field-value {
            modify-regular-expression regular-expression with field-value;
            add field-value;
            add-missing field-value;
            add-overwrite field-value;
            remove-regular-expression regular-expression;
            remove-all;
            reject-regular-expression regular-expression;
          }
        }
        request-uri request-uri {
          field-value {
            modify-regular-expression regular-expression with field-value;
          }
        }
      }
    }
  }
  new-call-usage-policy policy-name {
  term term-name {
    from {
      contact [ contact-fields ];
      method {
        method-invite;
      }
      request-uri [ uri-fields ];
      source-address [ ip-addresses ];
    }
    then {
      media-policy {
        data-inactivity-detection (Services PGCP) {
          inactivity-duration (Services Border Signaling Gateway) seconds;
        }
        no-anchoring;
        service-class service-class-name;
      }
      trace;
    }
  }
  new-call-usage-policy-set policy-set-name {
    policy-name [ policy-names ];
  }
  new-transaction-policy policy-name {
    term term-name {
      from {
```

Copyright © 2013, Juniper Networks, Inc.

```
            contact {
              registration-state [ registered | not-registered ];
              regular-expression [ regular-expression ];
              uri-hiding [ hidden-uri | not-hidden-uri ];
            }
            method {
              method-invite;
              method-message;
              method-options;
              method-publish;
              method-refer;
              method-register;
              method-subscribe;
            }
            request-uri {
              registration-state [ registered | not-registered ];
              regular-expression [ regular-expression ];
              uri-hiding [ hidden-uri | not-hidden-uri ];
            }
          source-address [ ip-addresses ];
          }
          then {
            (accept | reject);
            admission-control admission-control-profile;
            message-manipulation {
              forward-manipulation {
                manipulation-rule-name;
              }
              reverse-manipulation {
                manipulation-rule-name;
              }
            }
            on-3xx-response{
              recursion-limit number;
              }
            }
            route (Services) {
              egress-service-point service-point-name;
              next-hop (Services) (request-uri | address ipv4-address | <port port-number>
                <transport-protocol (udp | tcp)>);
              server-cluster cluster-name;
            }
            signaling-realm signaling-realm;
            trace;
          }
        }
      }
      new-transaction-policy-set policy-set-name {
        policy-name [ policy-names ];
      }
      routing-destinations {
        availability-check-profiles {
          profile-name;
          keepalive-interval {
            available-server seconds;
            unavailable-server seconds;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
        keepalive-method sip-options;
        keepalive-strategy (do-not-send <blackout-period seconds> | send-always <
          failures-before-unavailable number> < successes-before-available number |
          send-when-unavailable < successes-before-available number);
        transaction-timeout seconds;
      }
      clusters [
        cluster-name;
        server server-name {
          priority priority-level;
          weight weight-level;
        }
      }
      default-availability-check-profile profile-name;
      }
      servers {
        server-name {
        address ip4-address <port port-number> <transport (udp | tcp)>;
        admission-control profile-name;
        availability-check-profile profile-name;
        service-point service-point-name;
      }
    timers {
      inactive-callseconds;
      timer-c seconds;
    }
}
traceoptions {
    file {
      filename filename;
      files number;
      match regex;
      size size;
    }
    flag {
      datastore {
        data trace-level;
        db trace-level;
        handle trace-level;
        minimum trace-level;
      }
      framework {
        action trace-level;
        event trace-level;
        executor trace-level;
        freezer trace-level;
        minimum trace-level;
        memory-pool trace-level;
      }
      minimum trace-level;
      sbc-utils {
        common trace-level;
        configuration trace-level;
        device-monitor trace-level;
        ipc trace-level;
```

Copyright © 2013, Juniper Networks, Inc.

```
            memory-management trace-level;
            message trace-level;
            minimum trace-level;
            user-interface trace-level;
          }
          session-trace trace-level;
          signaling {
            b2b trace-level;
            b2b-wrapper trace-level;
            minimum trace-level;
            policy trace-level;
            sip-stack-wrapper trace-level;
            topology-hiding trace-level;
            ua trace-level;
          }
          sip-stack {
            dev-logging;
            event-tracing;
            ips-tracing;
            pd-log-detail (full | summary);
            pd-log-level (audit | exception | problem);
            per-tracing;
            verbose-logging;
          }
        }
      }
    }
  }
  cos {
    application-profile profile-name {
      ftp {
        data {
          dscp (alias | bits);
          forwarding-class class-name;
        }
      }
      sip {
        video {
          dscp (alias | bits);
          forwarding-class class-name;
        }
        voice {
          dscp (alias | bits);
          forwarding-class class-name;
        }
      }
    }
    rule rule-name {
      match-direction (input | output | input-output);
      term term-name {
        from {
          application-sets set-name;
          applications [ application-names ];
          destination-address (CoS) address;
          destination-prefix-list list-name <except>;
          source-address address;
```


Copyright © 2013, Juniper Networks, Inc.

```
            source-prefix-list list-name <except>;
          }
          then {
            application-profile profile-name;
            dscp (alias | bits);
            forwarding-class class-name;
            (reflexive | reverse) {
              application-profile profile-name;
              dscp (alias | bits);
              forwarding-class class-name;
              syslog;
            }
            syslog;
          }
        }
      }
      rule-set rule-set-name {
        rule rule-name;
      }
    }
    dynamic-flow-capture {
      capture-group client-name {
        content-destination identifier {
          address address;
          hard-limit bandwidth;
          hard-limit-target bandwidth;
          soft-limit bandwidth;
          soft-limit-clear bandwidth;
          ttl hops;
        }
        control-source identifier {
          allowed-destinations [ destination ];
          minimum-priority value;
          no-syslog;
          notification-targets address port port-number;
          service-port port-number;
          shared-key value;
          source-addresses [ address ];
        }
        duplicates-dropped-periodicity seconds;
        input-packet-rate-threshold rate;
        interfaces interface-name;
        max-duplicates number;
        pic-memory-threshold percentage percentage;
      }
      g-max-duplicates number;
      g-duplicates-dropped-periodicity seconds;
    }
    ip-reassembly{
      profile profile-name
      rule rule-name{
      match-direction direction;
    }
    flow-collector {
      analyzer-address address;
      analyzer-id name;
```

```
destinations {
  ftp:url {
    password "password";
  }
  file-specification {
    variant variant-number {
      data-format format;
      name-format format;
      transfer {
        record-level number;
        timeout seconds;
      }
    }
  }
  interface-map {
    collector interface-name;
    file-specification variant-number;
    interface-name {
      collector interface-name;
      file-specification variant-number;
    }
  }
  retry number;
  retry-delay seconds;
  transfer-log-archive {
    archive-sites {
      ftp:url {
        password "password";
        username username;
      }
    }
    filename-prefix prefix;
    maximum-age minutes;
  }
}
flow-monitoring {
  version9 {
    template template-name {
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      ipv4-template;
      ipv6-template;
      mpls-template {
        label-position [ positions ];
      }
      mpls-ipv4-template {
        label-position [ positions ];
      }
      option-refresh-rate packets packets seconds seconds;
      template-refresh-rate packets packets seconds seconds;
    }
  }
}
flow-tap {
  (interface interface-name | tunnel-interface interface-name);
}
```


Copyright © 2013, Juniper Networks, Inc.

```
ids {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        application-sets set-name;
        applications [ application-names ];
        destination-address (address | any-unicast) <except>;
        destination-address-range low minimum-value high maximum-value<except>;
        destination-prefix-list list-name <except>;
        source-address (address | any-unicast) <except>;
        source-address-range low minimum-value high maximum-value <except>;
        source-prefix-list list-name <except>;
      }
      then {
        aggregation {
          destination-prefix prefix-number | destination-prefix-ipv6 prefix-number;
          source-prefix prefix-number | source-prefix-ipv6 prefix-number;
        }
        (force-entry | ignore-entry);
        logging {
          syslog;
          threshold rate;
        }
        session-limit {
          by-destination {
            hold-time seconds;
            maximum number;
            packets number;
            rate number;
          }
          by-pair {
            maximum number;
            packets number;
            rate number;
          }
          by-source {
            hold-time seconds;
            maximum number;
            packets number;
            rate number;
          }
        }
        syn-cookie {
          mss value;
          threshold rate;
        }
      }
    }
  }
  rule-set rule-set-name {
    rule rule-name;
  }
}
ipsec-vpn {
  clear-ike-sas-on-pic-restart;
```

Copyright © 2013, Juniper Networks, Inc.

```
clear-ipsec-sas-on-pic-restart;
ike {
  proposal proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256);
    authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2 | group5 |group14);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
  }
  policy policy-name {
    description description;
    local-certificate identifier;
    local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
    version (1 | 2);
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
    remote-id {
      ipv4_addr [ values ];
      ipv6_addr [ values ];
      key_id [ values ];
    }
  }
}
ipsec {
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    description description;
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
  }
  policy policy-name {
    description description;
    perfect-forward-secrecy {
      keys (group1 | group2);
    }
    proposals [ proposal-names ];
  }
}
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      destination-address address;
      ipsec-inside-interface interface-name;
      source-address address;
    }
    then {
      anti-replay-window-size bits;
      backup-remote-gateway address;
      clear-dont-fragment-bit;
      dynamic {
        ike-policy policy-name;
        ipsec-policy policy-name;
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
            initiate-dead-peer-detection;
            manual {
              direction (inbound | outbound | bidirectional) {
                authentication {
                  algorithm (hmac-md5-96 | hmac-sha1-96);
                  key (ascii-text key | hexadecimal key );
                }
                auxiliary-spi spi-value;
                encryption {
                  algorithm algorithm;
                  key (ascii-text key | hexadecimal key );
                }
                protocol (ah | bundle | esp);
                spi spi-value;
              }
            }
            no-anti-replay;
            remote-gateway address;
            syslog;
            tunnel-mtu bytes;
          }
        }
      }
      rule-set rule-set-name {
        rule rule-name;
      }
      no-ipsec-tunnel-in-traceroute;
      traceoptions {
        file {
          files number;
          size bytes;
        }
        flag flag;
        level level;
      }
    }
    l2tp {
      tunnel-group name {
        hello-interval seconds;
        hide-avps;
        l2tp-access-profile profile-name;
        local-gateway address address;
        maximum-send-window packets;
        ppp-access-profile profile-name;
        receive-window packets;
        retransmit-interval seconds;
        service-interface interface-name;
        syslog {
          host hostname {
            services severity-level;
            facility-override facility-name;
            log-prefix prefix-value;
          }
        }
        tunnel-timeout seconds;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    traceoptions {
      debug-level level;
      filter {
        protocol name;
      }
      flag flag;
      interfaces interface-name {
        debug-level level;
        flag flag;
      }
    }
}
logging {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag;
    }
}
nat {
    ipv6-multicast-interfaces (all | interface-name) {
      disable;
    }
    pool nat-pool-name {
      address ip-prefix</prefix-length>;
      address-range low minimum-value high maximum-value;
      pgcp {
        hint [ hint-strings ];
        ports-per-session ports;
        remotely-controlled;
        transport;
      }
      port (automatic | range low minimum-value high maximum-value) {
        random-allocation;
      }
    }
    rule rule-name {
      match-direction (input | output);
      term term-name {
        from {
          application-sets set-name;
          applications [ application-names ];
          destination-address (address | any-unicast) <except>;
          destination-address-range low minimum-value high maximum-value <except>;
          destination-prefix-list list-name <except>;
          source-address (address | any-unicast) <except>;
          source-address-range low minimum-value high maximum-value <except>;
          source-prefix-list list-name <except>;
        }
        then {
          syslog;
          translated {
            destination-pool nat-pool-name;
            destination-prefix destination-prefix;
            dns-alg-pool dns-alg-pool;
```


Copyright © 2013, Juniper Networks, Inc.

```
                dns-alg-prefix dns-alg-prefix;
                overload-pool overload-pool-name;
                overload-prefix overload-prefix;
                source-pool nat-pool-name;
                source-prefix source-prefix;
                translation-type {
                  (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44 |
                    napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44
                    |twice-dynamic-nat-44 |twice-napt-44);
                }
                use-dns-map-for-destination-translation;
              }
            }
          }
        }
        rule-set rule-set-name {
          rule rule-name;
        }
      }
      pgcp {
        gateway gateway-name {
          cleanup-timeout seconds;
          gateway-address gateway-address;
          fast-update-filters {
            maximum-terms number-of-terms;
            maximum-fuf-percentage percentage;
          }
          gateway-controller gateway-controller-name {
            controller-address ip-address;
            controller-port port-number;
            interim-ah-scheme {
              algorithm algorithm;
            }
          }
          gateway-port gateway-port;
          graceful-restart {
            maximum-synchronization-mismatches number-of-mismatches;
            seconds;
          }
          data-inactivity-detection {
            inactivity-delay seconds;
            latch-deadlock-delay seconds;
            send-notification-on-delay;
            inactivity-duration seconds;
            no-rtcp-check
            stop-detection-on-drop;
            report-service-change {
              service-change-type (forced-906) | forced-910);
            }
          }
          h248-properties {
            application-data-inactivity-detection {
              ip-flow-stop-detection (regulated-notify | immediate-notify);
            }
            base-root {
              mg-provisional-response-timer-value {
```

Copyright © 2013, Juniper Networks, Inc.

```
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mgc-originated-pending-limit {
      default number-of-messages;
      maximum number-of-messages;
      minimum number-of-messages;
    }
    normal-mg-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    normal-mgc-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
  }
  diffserv {
    dscp {
      default (dscp-value | alias | do-not-change);
    }
  }
  event-timestamp-notification {
    request-timestamp (requested | suppressed | autonomous);
  {
  hanging-termination-detection {
    timerx seconds;
  }
  notification-behavior {
    notification-regulation default (once | 0 - 100);
  }
  segmentation {
    mg-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mgc-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mg-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
    mgc-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
  }
```


Copyright © 2013, Juniper Networks, Inc.

```
    traffic-management {
      max-burst-size {
        default bytes-per-second;
        maximum bytes-per-second;
        minimum bytes-per-second;
        rtcp {
          (fixed-value bytes-per-second | percentage percentage);
        }
      }
      peak-data-rate {
        default bytes-per-second;
        maximum bytes-per-second;
        minimum bytes-per-second;
        rtcp {
          (fixed-value bytes-per-second | percentage percentage);
        }
      }
      sustained-data-rate {
        default bytes-per-second;
        maximum bytes-per-second;
        minimum bytes-per-second;
        rtcp {
          (fixed-value bytes-per-second | percentage percentage);
        }
        rtcp-include;
      }
    }
    inactivity-timer {
      inactivity-timeout {
        detect;
        maximum-inactivity-time {
          default 10-millisecond-units;
          maximum 10-millisecond-units;
          minimum 10-millisecond-units;
        }
      }
    }
  }
  h248-options {
    audit-observed-events-returns;
    encoding {
      no-dscp-bit-mirroring;
      use-lower-case
    }
    service-change {
      context-indications {
        state-loss (forced-910 | forced-915 | none);
      }
      control-association-indications {
        disconnect {
          controller-failure (failover-909 | restart-902);
          reconnect (disconnected-900 | restart-902);
        }
        down {
          administrative (forced-905 | forced-908 | none);
          failure (forced-904 | forced-908 | none);
```

Copyright © 2013, Juniper Networks, Inc.

```
            graceful (graceful-905 | none);
          }
          up {
            cancel-graceful (none | restart-918);
            failover-cold (failover-920 | restart-901);
            failover-warm (failover-919 | restart-902);
          }
        }
        virtual-interface-indications {
          virtual-interface-down {
            administrative (forced-905 | forced-906 | none);
            failure (forced-904 | forced-906 | none);
            graceful (graceful-905 | none);
          }
          use-wildcard-response;
          virtual-interface-up {
            cancel-graceful (none | restart-918);
            warm (none | restart-900);
          }
        }
      }
    }
    h248-timers {
      initial-average-ack-delay milliseconds;
      maximum-net-propagation-delay milliseconds;
      maximum-waiting-delay milliseconds;
      tmax-retransmission-delay milliseconds;
    }
    max-concurrent-calls number-of-calls;
    monitor {
      media {
        rtcp;
        rtp;
      }
    }
    service-state (in-service | out-of-service-forced | out-of-service-graceful);
    session-mirroring {
      delivery-function delivery-function-name {
        destination-address destination-address;
        destination-port destination-port;
        network-operator-id network-operator-id;
        source-address source-address;
        source-port (Mirrored BGF Packets) source-port;
      }
      disable-session-mirroring;
    }
  }
  nat-pool nat-pool-name;
  rule rule-name {
    gateway gateway-name;
    nat-pool nat-pool-name;
  }
  rule-set rule-set-name {
    rule rule-name;
  }
  traceoptions {
```


Copyright © 2013, Juniper Networks, Inc.

```
        file <filename filename> <files number> <match regex> <size size> <world-readable
          | no-world-readable>;
        flag {
          bgf-core {
            common trace-level;
            default trace-level;
            firewall trace-level;
            gate-logic trace-level;
            pic-broker trace-level;
            policy trace-level;
            statistics trace-level;
          }
          default trace-level;
          h248-stack {
            control-association trace-level;
            default trace-level;
            messages;
            media-gateway trace-level;
          }
          sbc-utils {
            common trace-level;
            configuration trace-level;
            default trace-level;
            device-monitor trace-level;
            ipc trace-level;
            memory-management trace-level;
            messaging trace-level;
            user-interface trace-level;
          }
        }
      }
      virtual-interface interface-number {
        nat-pool nat-pool-name;
        service-interface interface-identifier;
        routing-instance instance-name {
          service-interface interface-name.unit-number;
        }
        service-state (in-service | out-of-service-forced | out-of-service-graceful);
      }
      session-mirroring {
        delivery-function delivery-function-name {
          destination-address destination-address;
          destination-port destination-port;
          network-operator-id network-operator-id;
          source-address source-address;
          source-port (Mirrored BGF Packets) source-port;
        }
        disable-session-mirroring;
      }
    }
    ptsp {
      forward-rule rule-name {
        term precedence {
          from {
            application-groups [ application-group-name ];
            applications [ application-name ];
```

Copyright © 2013, Juniper Networks, Inc.

```
            local-address address <except>;
            local-address-range low low-value high high-value <except >;
            local-prefix-list prefix-list-name <except >;
          }
          then {
            forwarding-instance forwarding-instance unit-number unit-number;
          }
        }
      }
      rule rule-name {
        count-type (application | rule);
        demux (destination-address | source-address);
        forward-rule forward-rule-name;
        match-direction (input | input-output | output);
        term precedence {
          from {
            application-group-any;
            application-groups [ application-group-name ];
            applications [ application-name ];
            local-port-range low low-value high high-value;
            local-ports [ value-list ];
            protocol protocol-number;
            remote-address address <except>;
            remote-address-range low low-value high high-value <except>;
            remote-port-range low low-value high high-value;
            remote-ports [ value-list ];
            remote-prefix-list prefix-list-name <except>;
          }
          then {
            (accept | discard);
            count (application | application-group | application-group-any | rule | none);
            forwarding-class forwarding-class;
            police policer-name;
          }
        }
      }
      rule-set rule-set-name {
        rule rule-name;
      }
    }
    rpm {
      bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name <routing-instances routing-instance-name>;
        moving-average-size number;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instances instance-name;
        test-interval interval;
      }
      probe owner {
        test test-name {
```


Copyright © 2013, Juniper Networks, Inc.

```
            data-fill data;
            data-size size;
            destination-interface interface-name;
            destination-port port;
            dscp-code-point dscp-bits;
            hardware-timestamp;
            history-size size;
            moving-average-size number;
            one-way-hardware-timestamp;
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance instance-name;
            source-address address;
            target (url | address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
        }
    }
    probe-limit limit;
    probe-server {
        tcp {
            destination-interface interface-name;
            port number;
        }
        udp {
            destination-interface interface-name;
            port number;
        }
    }
    twamp {
        server {
            authentication-mode (authenticated | encrypted | none);
            client-list list-name {
                address address;
            }
            inactivity-timeout seconds;
            maximum-connections count;
            maximum-connections-per-client count;
            maximum-sessions count;
            maximum-sessions-per-connection count;
            port number;
        }
    }
}
service-set service-set-name {
    aacl-rules rule-name;
    policy-decision-statistics-profile profile-name;
    (ids-rules rule-names | ids-rule-sets rule-set-name);
    (ipsec-vpn-rules rule-names | ipsec-vpn-rule-sets rule-set-name);
    (nat-rules rule-names | nat-rule-sets rule-set-name);
    (pgcp-rules rule-names | pgcp-rule-sets rule-set-name);
    (ptsp-rules rule-names | ptsp-rule-sets rule-set-name);
    (stateful-firewall-rules rule-names | stateful-firewall-rule-sets rule-set-name);
    allow-multicast;
```

Copyright © 2013, Juniper Networks, Inc.

```
    extension-service service-name {
      provider-specific rules;
    }
    interface-service {
      service-interface interface-name;
    }
    ip-reassembly-rules
    ipsec-vpn-options {
      anti-replay-window-size bits;
      clear-dont-fragment-bit;
      ike-access-profile profile-name;
      local-gateway address;
      no-anti-replay;
      passive-mode-tunneling;
      trusted-ca [ ca-profile-names ];
      tunnel-mtu bytes;
    }
    max-flows number;
    next-hop-service {
      inside-service-interface interface-name.unit-number;
      outside-service-interface interface-name.unit-number;
      outside-service-interface-type interface-type;
      service-interface-pool name;
    }
    service-order {
      forward-flow [ service-name1 service-name2 ];
      reverse-flow [ service-name1 service-name2 ];
    }
    syslog {
      host hostname {
        services severity-level;
        facility-override facility-name;
        port port-number;
      }
    }
  }
  softwire {
    softwire-concentrator {
      ds-lite ds-lite-softwire-concentrator  {
        auto-update-mtu;
        copy-dscp;
        flow-limit flow-limit;
        mtu-v6 mtu-v6;
        softwire-address address;
      }
      v6rdv6rd-softwire-concentator{
        ipv4-prefix ipv4-prefix;
        v6rd-prefix ipv6-prefix;
        mtu-v4 mtu-v4;
      }
    }
    rulerule-name {
      match-direction (input | output);
      term term-name{
        then {
          ds-lite name;
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
        ipv6-multicast-filters
      }
      stateful-firewall {
        rule rule-name {
          match-direction (input | output | input-output);
          term term-name {
            from {
              application-sets set-name;
              applications [ application-names ];
              destination-address (address | any-unicast) <except>;
              destination-address-range low minimum-value high maximum-value <except>;
              destination-prefix-list list-name <except>;
              source-address (address | any-unicast) <except>;
              source-address-range low minimum-value high maximum-value<except>;
              source-prefix-list list-name <except>;
            }
            then {
              (accept | discard | reject);
              allow-ip-options [ values ];
              syslog;
            }
          }
        }
        rule-set rule-set-name {
          rule rule-name;
        }
      }
    }
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 2

Adaptive Services

- Adaptive Services Overview on page 37
- Applications Configuration Guidelines on page 73
- Summary of Applications Configuration Statements on page 121
- Stateful Firewall Services Configuration Guidelines on page 131
- Summary of Stateful Firewall Configuration Statements on page 141
- Stateful Firewall on the Embedded Junos OS Platform Configuration Guidelines on page 153
- Summary of Stateful Firewall on the Embedded Junos OS Platform Configuration Statements on page 157
- Load Balancing Configuration Guidelines on page 167
- Summary of Load Balancing Configuration Statements on page 173
- Intrusion Detection Service Configuration Guidelines on page 185
- Summary of Intrusion Detection Service Configuration Statements on page 197
- IPsec Services Configuration Guidelines on page 219
- Summary of IPsec Services Configuration Statements on page 273
- Layer 2 Tunneling Protocol Services Configuration Guidelines on page 309
- Summary of Layer 2 Tunneling Protocol Configuration Statements on page 329
- Link Services IQ Interfaces Configuration Guidelines on page 347
- Summary of Link Services IQ Configuration Statements on page 409
- Voice Services Configuration Guidelines on page 421
- Summary of Voice Services Configuration Statements on page 431
- Class-of-Service Configuration Guidelines on page 441
- Summary of Class-of-Service Configuration Statements on page 451
- Service Set Configuration Guidelines on page 467
- Summary of Service Set Configuration Statements on page 485
- Service Interface Configuration Guidelines on page 513
- Summary of Service Interface Configuration Statements on page 527
- PGCP Configuration Guidelines for the BGF Feature on page 547
- Summary of PGCP Configuration Statements on page 553

Copyright © 2013, Juniper Networks, Inc.

- Service Interface Pools Configuration Guidelines on page 655
- Summary of Service Interface Pools Statements on page 657
- Border Signaling Gateway Configuration Guidelines on page 659
- Summary of Border Signaling Gateway Configuration Statements on page 665
- PTSP Configuration Guidelines on page 745
- Summary of PTSP Configuration Statements on page 747

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 3

Adaptive Services Overview

This chapter discusses the following topics:

- Adaptive Services Overview on page 37
- Enabling Service Packages on page 39
- Services Configuration Procedure on page 43
- Packet Flow Through the Adaptive Services or Multiservices PIC on page 44
- Stateful Firewall Overview on page 45
- Network Address Translation Overview on page 48
- Tunneling Services for IPv4-to-IPv6 Transition Overview on page 53
- IPsec Overview on page 57
- Layer 2 Tunneling Protocol Overview on page 59
- Voice Services Overview on page 59
- HTTP URL Tracking and Policy Control for Client Requests on page 60
- Configuring HTTP URL Tracking and Policy Control on page 61
- Class of Service Overview on page 62
- Examples: Services Interfaces Configuration on page 62

Adaptive Services Overview

MultiServices PICs and MultiServices Dense Port Concentrators (MS-DPCs) provide *adaptive services interfaces*, which allow you to coordinate multiple services on a single PIC by configuring a set of services and applications. MultiServices PICs and MS-DPCs offer a special range of services you configure in one or more service sets.

The MultiServices PIC is available in three versions, the MultiServices 100, the MultiServices 400, and the MultiServices 500, which differ in memory size and performance. All versions offer enhanced performance in comparison with AS PICs. MultiServices PICs are supported on M Series and T Series routers except M20 routers.

The MultiServices DPC is available for MX Series routers; it includes a subset of the functionality supported on the MultiServices PIC. Currently the MultiServices DPC supports the following Layer 3 services: stateful firewall, NAT, IDS, IPsec, active flow monitoring, RPM, and generic routing encapsulation (GRE) tunnels (including GRE key and

Copyright © 2013, Juniper Networks, Inc.

fragmentation); it also supports graceful Routing Engine switchover (GRES) and Dynamic Applicaton Awareness for Junos OS. For more information about supported packages, see "Enabling Service Packages" on page 39.

It is also possible to group several Multiservices PICs into an aggregated Multiservices (AMS) system. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs. Starting with Junos OS 11.4, all MX Series routers will support high availability (HA) and Network Address Translation (NAT) on AMS infrastructure. See "Configuring Load Balancing on AMS Infrastructure" on page 167 for more information.



NOTE: **The MultiServices PICs are polling based and not interrupt based; as a result, a high value in the show chassis pic "Interrupt load average" field may not mean that the PIC has reached its maximum limit of processing.**

The following services are configured within a service set and are available only on adaptive services interfaces:

- Stateful firewall—A type of firewall filter that considers state information derived from previous communications and other applications when evaluating traffic.
- Network Address Translation (NAT)—A security procedure for concealing host addresses on a private network behind a pool of public addresses.
- Intrusion detection service (IDS)—A set of tools for detecting, redirecting, and preventing certain kinds of network attack and intrusion.
- IP Security (IPsec)—A set of tools for configuring manual or dynamic security associations (SAs) for encryption of data traffic.
- Class of service (CoS)—A subset of CoS functionality for services interfaces, limited to DiffServ code point (DSCP) marking and forwarding-class assignment. CoS BA classification is not supported on services interfaces.

The configuration for these services comprises a series of rules that you can arrange in order of precedence as a *rule set*. Each rule follows the structure of a firewall filter, with a **from** statement containing input or match conditions and a **then** statement containing actions to be taken if the match conditions are met.

The following services are also configured on the MultiServices PICs and MS-DPCs, but do not use the rule set definition:

- Layer 2 Tunneling Protocol (L2TP)—A tool for setting up secure tunnels using Point-to-Point Protocol (PPP) encapsulation across Layer 2 networks.
- Link Services Intelligent Queuing (LSQ)—Interfaces that support Junos OS class-of-service (CoS) components, link fragmentation and interleaving (LFI) (FRF.12), Multilink Frame Relay (MLFR) user-to-network interface (UNI) network-to-network interface (NNI) (FRF.16), and Multilink PPP (MLPPP).

 Copyright © 2013, Juniper Networks, Inc.

- Voice services—A feature that uses the Compressed Real-Time Transport Protocol (CRTP) to enable voice over IP traffic to use low-speed links more effectively.

In addition, Junos OS includes the following tools for configuring services:

- Application protocols definition—Allows you to configure properties of application protocols that are subject to processing by router services, and group the application definitions into application sets.
- Service-set definition—Allows you to configure combinations of directional rules and default settings that control the behavior of each service in the service set.



NOTE: Logging of adaptive services interfaces messages to an external server by means of the fxp0 port is not supported on M Series routers. The architecture does not support system logging traffic out of a management interface. Instead, access to an external server is supported on a Packet Forwarding Engine interface.

Enabling Service Packages

For AS PICs, Multiservices PICs, Multiservices DPCs, and the internal Adaptive Services Module (ASM) in the M7i router, there are two service packages: Layer 2 and Layer 3. Both service packages are supported on all adaptive services interfaces, but you can enable only one service package per PIC, with the exception of a combined package supported on the ASM. On a single router, you can enable both service packages by installing two or more PICs on the platform.



NOTE: Graceful Routing Engine switchover (GRES) is automatically enabled on all services PICs and DPCs except the ES PIC. It is supported on all M Series, MX Series, and T Series routers except for TX Matrix routers. Layer 3 services should retain state after switchover, but Layer 2 services will restart. For IPsec services, Internet Key Exchange (IKE) negotiations are not stored and must be restarted after switchover. For more information about GRES, see the Junos OS High Availability Configuration Guide.

You enable service packages per PIC, not per port. For example, if you configure the Layer 2 service package, the entire PIC uses the configured package. To enable a service package, include the service-package statement at the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services]** hierarchy level, and specify **layer-2** or **layer-3**:

```
[edit chassis fpc slot-number pic pic-number adaptive-services]
service-package (layer-2 | layer-3);
```

To determine which package an AS PIC supports, issue the **show chassis hardware** command: if the PIC supports the Layer 2 package, it is listed as **Link Services II**, and if it supports the Layer 3 package, it is listed as **Adaptive Services II**. To determine which

Copyright © 2013, Juniper Networks, Inc.

package a Multiservices PIC supports, issue the **show chassis pic fpc-slot *slot-number* pic-slot *slot-number*** command. The **Package** field displays the value **Layer-2** or **Layer-3**.

NOTE: The ASM has a default option (layer-2-3) that combines the features available in the Layer 2 and Layer 3 service packages.

After you commit a change in the service package, the PIC is taken offline and then brought back online immediately. You do not need to manually take the PIC offline and online.

NOTE: Changing the service package causes all state information associated with the previous service package to be lost. You should change the service package only when there is no active traffic going to the PIC.

The services supported in each package differ by PIC and platform type. Table 3 on page 40 lists the services supported within each service package for each PIC and platform. For information about services supported on SRX Series Services Gateways and J Series Services Routers, see the *Junos OS Feature Support Reference for SRX Series and J Series Devices*.

On the AS and Multiservices PICs, *link services* support includes Junos OS CoS components, LFI (FRF.12), MLFR end-to-end (FRF.15), MLFR UNI NNI (FRF.16), MLPPP (RFC 1990), and multiclass MLPPP. For more information, see "Layer 2 Service Package Capabilities and Interfaces" on page 42 and "Layer 2 Service Package Capabilities and Interfaces" on page 348.

NOTE: The AS PIC II for Layer 2 Service is dedicated to supporting the Layer 2 service package only.

For additional information about Layer 3 services, see the Junos OS Feature Guides.

Table 3: AS and Multiservices PIC Services by Service Package, PIC, and Platform

Services	ASM	AS/AS2 PICs and Multiservices PICs	AS/AS2 and Multiservices PICs	AS2 and Multiservices PICs	AS2 and Multiservices PICs
Layer 2 Service Package (Only)	M7i	M7i, M10i, and M20	M40e and M120	M320, T320, and T640	TX Matrix
Link Services:					
• Link services	Yes	Yes	Yes	Yes	No
• Multiclass MLPPP	Yes	Yes	Yes	Yes	No

Copyright © 2013, Juniper Networks, Inc.

Table 3: AS and Multiservices PIC Services by Service Package, PIC, and Platform *(continued)*

Services	ASM	AS/AS2 PICs and Multiservices PICs	AS/AS2 and Multiservices PICs	AS2 and Multiservices PICs	AS2 and Multiservices PICs
Voice Services:					
• CRTP and LFI	Yes	Yes	Yes	Yes	No
• CRTP and MLPPP	Yes	Yes	Yes	Yes	No
• CRTP over PPP (without MLPPP)	Yes	Yes	Yes	Yes	No
Layer 3 Service Package (Only)	**M7i**	**M7i, M10i, and M20**	**M40e and M120**	**M320, T320, and T640**	**TX Matrix**
Security Services:					
• CoS	Yes	Yes	Yes	Yes	No
• Intrusion detection system (IDS)	Yes	Yes	Yes	Yes	No
• IPsec	Yes	Yes	Yes	Yes	No
• NAT	Yes	Yes	Yes	Yes	No
• Stateful firewall	Yes	Yes	Yes	Yes	No
Accounting Services:					
• Active monitoring	Yes	Yes	Yes	Yes	Yes
• Dynamic flow capture (Multiservices 400 PIC only)	No	No	No	Yes	No
• Flow-tap	Yes	Yes	Yes (M40e only)	Yes	No
• Passive monitoring (Multiservices 400 PIC only)	No	Yes	Yes (M40e only)	Yes	No
• Port mirroring	Yes	Yes	Yes	Yes	Yes
LNS Services:					
• L2TP LNS	Yes	Yes (M7i and M10i only)	Yes (M120 only)	No	No
Voice Services:					

Copyright © 2013, Juniper Networks, Inc.

Table 3: AS and Multiservices PIC Services by Service Package, PIC, and Platform *(continued)*

Services	ASM	AS/AS2 PICs and Multiservices PICs	AS/AS2 and Multiservices PICs	AS2 and Multiservices PICs	AS2 and Multiservices PICs
• BGF	Yes	Yes	Yes	Yes	No
Layer 2 and Layer 3 Service Package (Common Features)	**M7i**	**M7i, M10i, and M20**	**M40e and M120**	**M320, T320, and T640**	**TX Matrix**
RPM Services:					
• RPM probe timestamping	Yes	Yes	Yes	Yes	No
Tunnel Services:					
• GRE (**gr-*fpc/pic/port***)	Yes	Yes	Yes	Yes	Yes
• GRE fragmentation (**clear-dont-fragment-bit**)	Yes	Yes	Yes	No	No
• GRE key	Yes	Yes	Yes	Yes	No
• IP-IP tunnels (**ip-*fpc/pic/port***)	Yes	Yes	Yes	Yes	Yes
• Logical tunnels (**lt-*fpc/pic/port***)	No	No	No	No	No
• Multicast tunnels (**mt-*fpc/pic/port***)	Yes	Yes	Yes	Yes	Yes
• PIM de-encapsulation (**pd-*fpc/pic/port***)	Yes	Yes	Yes	Yes	Yes
• PIM encapsulation (**pe-*fpc/pic/port***)	Yes	Yes	Yes	Yes	Yes
• Virtual tunnels (**vt-*fpc/pic/port***)	Yes	Yes	Yes	Yes	Yes

Layer 2 Service Package Capabilities and Interfaces

When you enable the Layer 2 service package, you can configure link services. On the AS and Multiservices PICs and the ASM, link services include support for the following:

- Junos CoS components—"Layer 2 Service Package Capabilities and Interfaces" on page 348 describes how the Junos CoS components work on link services IQ (**lsq**) interfaces. For detailed information about Junos CoS components, see the Junos OS Class of Service Configuration Guide.
- LFI on Frame Relay links using FRF.12 end-to-end fragmentation—The standard for FRF.12 is defined in the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*.
- LFI on MLPPP links.

Copyright © 2013, Juniper Networks, Inc.

- MLFR UNI NNI (FRF.16)—The standard for FRF.16 is defined in the specification FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*.
- MLPPP (RFC 1990)
- MLFR end-to-end (FRF.15)

For the LSQ interface on the AS and Multiservices PICs, the configuration syntax is almost the same as for Multilink and Link Services PICs. The primary difference is the use of the interface-type descriptor **lsq** instead of **ml** or **ls**. When you enable the Layer 2 service package, the following interfaces are automatically created:

gr-*fpc/pic/port*
ip-*fpc/pic/port*
lsq-*fpc/pic/port*
lsq-*fpc/pic/port*:0
...
lsq-*fpc/pic/port*:*N*
mt-*fpc/pic/port*
pd-*fpc/pic/port*
pe-*fpc/pic/port*
sp-*fpc/pic/port*
vt-*fpc/pic/port*

Interface types **gr**, **ip**, **mt**, **pd**, **pe**, and **vt** are standard tunnel interfaces that are available on the AS and Multiservices PICs whether you enable the Layer 2 or the Layer 3 service package. These tunnel interfaces function the same way for both service packages, except that the Layer 2 service package does not support some tunnel functions, as shown in Table 3 on page 40.

Interface type **lsq-*fpc/pic/port*** is the physical link services IQ (**lsq**) interface. Interface types **lsq-*fpc/pic/port*:0** through **lsq-*fpc/pic/port*:*N*** represent FRF.16 bundles. These interface types are created when you include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** option. For more information, see "Layer 2 Service Package Capabilities and Interfaces" on page 348 and Link and Multilink Properties.



NOTE: Interface type **sp** is created because it is needed by the Junos OS. For the Layer 2 service package, the **sp** interface is not configurable, but you should not disable it.

Services Configuration Procedure

You follow these general steps to configure services:

1. Define application objects by configuring statements at the **[edit applications]** hierarchy level.
2. Define service rules by configuring statements at the **[edit services (ids | ipsec-vpn | nat | stateful-firewall) rule]** hierarchy level.
3. Group the service rules by configuring the **rule-set** statement at the **[edit services (ids | ipsec-vpn | nat | stateful-firewall)]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

4. Group service rule sets under a service-set definition by configuring the **service-set** statement at the **[edit services]** hierarchy level.
5. Apply the service set on an interface by including the **service-set** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output)]** hierarchy level. Alternatively, you can configure logical interfaces as a next-hop destination by including the **next-hop-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level.



NOTE: You can configure IDS, NAT, and stateful firewall service rules within the same service set. You must configure IPsec services in a separate service set, although you can apply both service sets to the same PIC.

Packet Flow Through the Adaptive Services or Multiservices PIC

You can optionally configure service sets to be applied at one of three points while the packets transit the router:

- An interface service set applied at the inbound interface.
- A next-hop service set applied at the forwarding table.
- An interface service set applied at the outbound interface.

The packet flow is as follows, graphically displayed in Figure 1 on page 45. (You can configure a service set as either an interface service set or a next-hop service set.)

1. Packets enter the router on the inbound interface.
2. A policer, filter, service filter, service set, postservice filter, and input forwarding-table filter are applied sequentially to the traffic; these are all optional items in the configuration. If an interface service set is applied, the packets are forwarded to the AS or MultiServices PIC for services processing and then sent back to the Packet Forwarding Engine; if a service filter is also applied, only packets matching the service filter are sent to the PIC. The optional postservice filter is applied and postprocessing takes place.
3. A next-hop service set can be applied to the VPN routing and forwarding (VRF) table or to **inet.0**. If it is applied, packets are sent to the PIC for services processing and sent back to the Packet Forwarding Engine.



NOTE: For NAT, the next-hop service set can only be applied to the VRF table. For all other services, the next-hop service set can be applied to either the VRF table or to inet.0.

4. On the output interface, an output filter, output policer, and interface service set can be applied sequentially to the traffic if you have configured any of these items. If an

Copyright © 2013, Juniper Networks, Inc.

interface service set is applied, the traffic is forwarded to the PIC for processing and sent back to the Packet Forwarding Engine, which then forwards the traffic.

5. Packets exit the router.

Figure 1: Packet Flow Through the Adaptive Services or MultiServices PIC

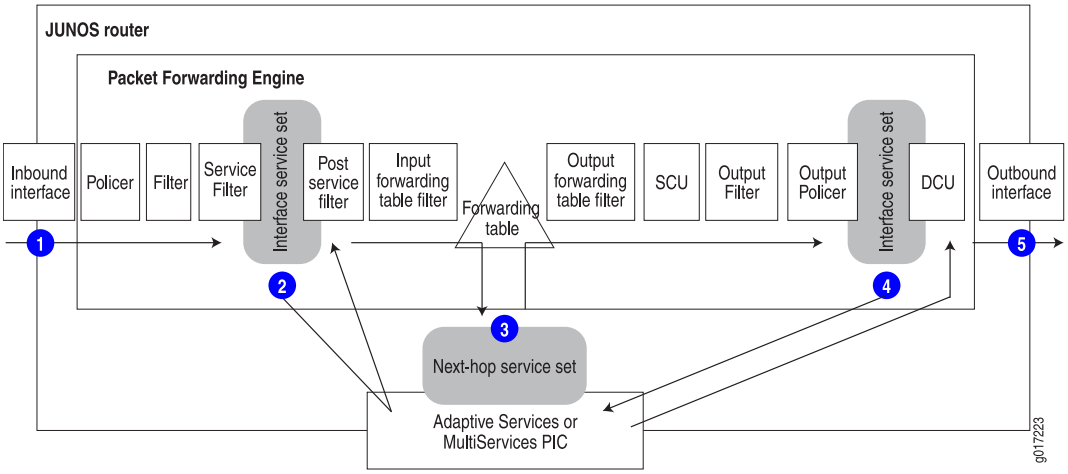

NOTE: When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG_ERR is generated. This mechanism applies to both Layer 2 and Layer 3 service packages.

Stateful Firewall Overview

Routers use firewalls to track and control the flow of traffic. Adaptive Services and MultiServices PICs employ a type of firewall called a *stateful firewall*. Contrasted with a *stateless* firewall that inspects packets in isolation, a stateful firewall provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions for new communication attempts.

Stateful firewalls group relevant *flows* into *conversations*. A flow is identified by the following five properties:

- Source address
- Source port
- Destination address
- Destination port
- Protocol

A typical Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) conversation consists of two flows: the initiation flow and the responder flow. However,

Copyright © 2013, Juniper Networks, Inc.

some conversations, such as an FTP conversation, might consist of two control flows and many data flows.

Firewall rules govern whether the conversation is allowed to be established. If a conversation is allowed, all flows within the conversation are permitted, including flows that are created during the life cycle of the conversation.

You configure stateful firewalls using a powerful rule-driven conversation handling path. A *rule* consists of direction, source address, source port, destination address, destination port, IP protocol value, and application protocol or service. In addition to the specific values you configure, you can assign the value **any** to rule objects, addresses, or ports, which allows them to match any input value. Finally, you can optionally negate the rule objects, which negates the result of the type-specific match.

Firewall rules are directional. For each new conversation, the router software checks the initiation flow matching the direction specified by the rule.

Firewall rules are ordered. The software checks the rules in the order in which you include them in the configuration. The first time the firewall discovers a match, the router implements the action specified by that rule. Rules still unchecked are ignored.

For more information, see "Configuring Stateful Firewall Rules" on page 132.

Stateful Firewall Support for Application Protocols

By inspecting the application protocol data, the AS or MultiServices PIC firewall can intelligently enforce security policies and allow only the minimal required packet traffic to flow through the firewall.

The firewall rules are configured in relation to an interface. By default, the stateful firewall allows all sessions initiated from the hosts behind the interface to pass through the router.

Stateful Firewall Anomaly Checking

The stateful firewall recognizes the following events as anomalies and sends them to the IDS software for processing:

- IP anomalies:
  - IP version is not correct.
  - IP header length field is too small.
  - IP header length is set larger than the entire packet.
  - Bad header checksum.
  - IP total length field is shorter than header length.
  - Packet has incorrect IP options.

Copyright © 2013, Juniper Networks, Inc.

- Internet Control Message Protocol (ICMP) packet length error.
- Time-to-live (TTL) equals 0.

- IP address anomalies:
  - IP packet source is a broadcast or multicast.
  - Land attack (source IP equals destination IP).
- IP fragmentation anomalies:
  - IP fragment overlap.
  - IP fragment missed.
  - IP fragment length error.
  - IP packet length is more than 64 kilobytes (KB).
  - Tiny fragment attack.
- TCP anomalies:
  - TCP port 0.
  - TCP sequence number 0 and flags 0.
  - TCP sequence number 0 and FIN/PSH/RST flags set.
  - TCP flags with wrong combination (TCP FIN/RST or SYN/(URG|FIN|RST).
  - Bad TCP checksum.
- UDP anomalies:
  - UDP source or destination port 0.
  - UDP header length check failed.
  - Bad UDP checksum.
- Anomalies found through stateful TCP or UDP checks:
  - SYN followed by SYN-ACK packets without ACK from initiator.
  - SYN followed by RST packets.
  - SYN without SYN-ACK.
  - Non-SYN first flow packet.
  - ICMP unreachable errors for SYN packets.
  - ICMP unreachable errors for UDP packets.
- Packets dropped according to stateful firewall rules.

If you employ stateful anomaly detection in conjunction with stateless detection, IDS can provide early warning for a wide range of attacks, including these:

- TCP or UDP network probes and port scanning

Copyright © 2013, Juniper Networks, Inc.

- SYN flood attacks
- IP fragmentation-based attacks such as teardrop, bonk, and boink

Network Address Translation Overview

- Types of NAT on page 48

Types of NAT

The types of Network Address Translation (NAT) supported by the Junos OS are described in the following sections:

- NAT Concept and Facilities Overview on page 48
- IPv4-to-IPv4 Basic NAT on page 49
- NAT-PT on page 50
- Static Destination NAT on page 50
- Twice NAT on page 50
- IPv6 NAT on page 51
- NAT-PT with DNS ALG on page 51
- Dynamic NAT on page 51
- Stateful NAT64 on page 52
- Dual-Stack Lite on page 52

NAT Concept and Facilities Overview

NAT is a mechanism for translating IP addresses. NAT provides the technology used to support a wide range of networking goals, including:

- Concealing a set of host addresses on a private network behind a pool of public addresses.
- Providing a security measure to protect the host addresses from direct targeting in network attacks.
- Providing a tool set for coping with IPv4 address depletion and IPV6 transition issues.

The Junos OS provides carrier-grade NAT (CGN) for IPv4 and IPv6 networks, and facilitates the transit of traffic between different types of networks.

The multiservices Dense Port Concentrator (DPC) and multiservices PIC interfaces support the following types of traditional CGN:

- Static-source translation—Allows you to hide a private network. It features a one-to-one mapping between the original address and the translated address; the mapping is configured statically. For more information, see "Basic NAT" on page 49.
- Dynamic-source translation—Includes two options: dynamic address-only source translation and Network Address Port Translation (NAPT):

 Copyright © 2013, Juniper Networks, Inc.

    - Dynamic address-only source translation—A NAT address is picked up dynamically from a source NAT pool and the mapping from the original source address to the translated address is maintained as long as there is at least one active flow that uses this mapping. For more information, see "Dynamic NAT" on page 51.
    - NAPT—Both the original source address and the source port are translated. The translated address and port are picked up from the corresponding NAT pool. For more information, see "NAPT" on page 49.

- Static destination translation—Allows you to make selected private servers accessible. It features a one-to-one mapping between the translated address and the destination address; the mapping is configured statically. For more information, see "Static Destination NAT" on page 50.
- Protocol translation—Allows you to assign addresses from a pool on a static or dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries. For more information, see "NAT-PT" on page 50, "NAT-PT with DNS ALG" on page 51, and "Stateful NAT64" on page 52.
- Encapsulation of IPv4 packets into IPv6 packets using softwires—Enables packets to travel over softwires to a carrier-grade NAT endpoint where they undergo source-NAT processing to hide the original source address. For more information, see "Tunneling Services for IPv4-to-IPv6 Transition Overview" on page 53.

The Junos OS supports NAT functionality described in IETF RFCs and Internet drafts, as shown in "Supported NAT and SIP Standards" in *Standards Supported in Junos OS 11.4*.

IPv4-to-IPv4 Basic NAT

Basic Network Address Translation or Basic NAT is a method by which IP addresses are mapped from one group to another, transparent to end users. Network Address Port Translation or NAPT is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. Together, these two operations, referred to as traditional NAT, provide a mechanism to connect a realm with private addresses to an external realm with globally unique registered addresses.

Traditional NAT, specified in RFC 3022, *Traditional IP Network Address Translator*, is fully supported by the Junos OS. In addition, NAPT is supported for source addresses.

Basic NAT

With Basic NAT, a block of external addresses is set aside for translating addresses of hosts in a private domain as they originate sessions to the external domain. For packets outbound from the private network, Basic NAT translates source IP addresses and related fields such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, Basic NAT translates the destination IP address and the checksums listed above.

NAPT

Use NAPT to enable the components of the private network to share a single external address. NAPT translates the transport identifier (for example, TCP port number, UDP port number, or ICMP query ID) of the private network into a single external address.

Copyright © 2013, Juniper Networks, Inc.

NAPT can be combined with Basic NAT to use a pool of external addresses in conjunction with port translation.

For packets outbound from the private network, NAPT translates the source IP address, source transport identifier (TCP/UDP port or ICMP query ID), and related fields, such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, NAPT translates the destination IP address, the destination transport identifier, and the IP and transport header checksums.

NAT-PT

NAT-Protocol Translation (NAT-PT) is an obsolete IPv4-to-IPv6 transition mechanism and is no longer recommended. NAT64 is the newer, recommended solution. Using a pool of IPv4 addresses, NAT-PT assigns addresses from that pool to IPv6 nodes on a dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries. Inbound and outbound sessions must traverse the same NAT-PT router so that it can track those sessions. RFC 2766, *Network Address Translation - Protocol Translation (NAT-PT)*, recommends the use of NAT-PT for translation between IPv6-only nodes and IPv4-only nodes, and *not* for IPv6-to-IPv6 translation between IPv6 nodes or IPv4-to-IPv4 translation between IPv4 nodes.

NAT-PT, specified in RFC 2766, *Network Address Translation - Protocol Translation (NAT-PT)* and obsoleted by RFC 2766, *Reasons to Move Network Address Translator - Protocol Translator (NAT-PT) to Historic Status*, is still supported by the Junos OS.

Static Destination NAT

Use static destination NAT to translate the destination address for external traffic to an address specified in a destination pool. The destination pool contains one address and no port configuration.

For more information about static destination NAT, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations.*

Twice NAT

In Twice NAT, both the source and destination addresses are subject to translation as packets traverse the NAT router. The source information to be translated can be either address only or address and port. For example, you would use Twice NAT when you are connecting two networks in which all or some addresses in one network overlap with addresses in another network (whether the network is private or public). In traditional NAT, only one of the addresses is translated.

To configure Twice NAT, you must specify both a destination address and a source address for the match direction, pool or prefix, and translation type.

You can configure application-level gateways (ALGs) for ICMP and traceroute under stateful firewall, NAT, or class-of-service (CoS) rules when Twice NAT is configured in the same service set. These ALGs cannot be applied to flows created by the Packet Gateway Control Protocol (PGCP). Twice NAT does not support other ALGs. By default, the Twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages.

Copyright © 2013, Juniper Networks, Inc.

Twice NAT, specified in RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, is fully supported by the Junos OS.

IPv6 NAT

IPv6-to-IPv6 NAT (NAT66), defined in Internet draft draft-mrw-behave-nat66-01, *IPv6-to-IPv6 Network Address Translation (NAT66)*, is fully supported by the Junos OS.

NAT-PT with DNS ALG

NAT-PT and Domain Name System (DNS) ALG are used to facilitate communication between IPv6 hosts and IPv4 hosts. Using a pool of IPv4 addresses, NAT-PT assigns addresses from that pool to IPv6 nodes on a dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries. Inbound and outbound sessions must traverse the same NAT-PT router so that it can track those sessions. RFC 2766, *Network Address Translation - Protocol Translation (NAT-PT)*, recommends the use of NAT-PT for translation between IPv6-only nodes and IPv4-only nodes, and *not* for IPv6-to-IPv6 translation between IPv6 nodes or IPv4-to-IPv4 translation between IPv4 nodes.

DNS is a distributed hierarchical naming system for computers, services, or any resource connected to the Internet or a private network. The DNS ALG is an application-specific agent that allows an IPv6 node to communicate with an IPv4 node and vice versa.

When DNS ALG is employed with NAT-PT, the DNS ALG translates IPv6 addresses in DNS queries and responses to the corresponding IPv4 addresses and vice versa. IPv4 name-to-address mappings are held in the DNS with “A” queries. IPv6 name-to-address mappings are held in the DNS with “AAAA” queries.

NOTE: For IPv6 DNS queries, use the **do-not-translate-AAAA-query-to-A-query** statement at the **[edit applications application *application-name*]** hierarchy level.

Related Documentation

- Configuring NAT Rules
- Configuring NAT-PT
- Example: Configuring NAT-PT

Dynamic NAT

Dynamic NAT flow is shown in Figure 2 on page 51.

Figure 2: Dynamic NAT Flow

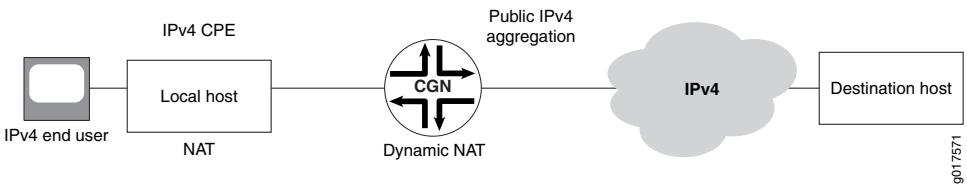

With dynamic NAT, you can map a private IP address (source) to a public IP address drawing from a pool of registered (public) IP addresses. NAT addresses from the pool are assigned dynamically. Assigning addresses dynamically also allows a few public IP addresses to be used by several private hosts, in contrast with an equal-sized pool required by source static NAT.

For more information about dynamic address translation, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

Stateful NAT64

Stateful NAT64 flow is shown in Figure 3 on page 52.

Figure 3: Stateful NAT64 Flow

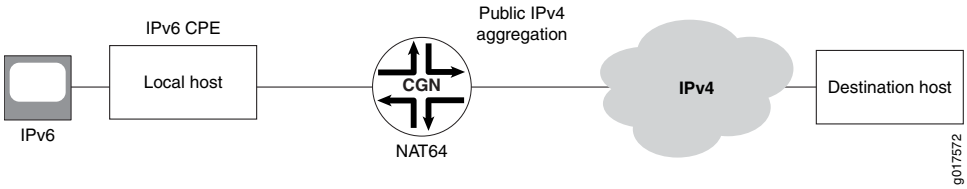

Stateful NAT64 is a mechanism to move to an IPv6 network and at the same time deal with IPv4 address depletion. By allowing IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP, several IPv6-only clients can share the same public IPv4 server address. To allow sharing of the IPv4 server address, NAT64 translates incoming IPv6 packets into IPv4 (and vice versa).

When stateful NAT64 is used in conjunction with DNS64, no changes are usually required in the IPv6 client or the IPv4 server. DNS64 is out of scope of this document because it is normally implemented as an enhancement to currently deployed DNS servers.

Stateful NAT64, specified in RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*, is fully supported by the Junos OS.

Dual-Stack Lite

Dual-stack lite (DS-Lite) flow is shown in Figure 4 on page 52.

Figure 4: DS-Lite Flow

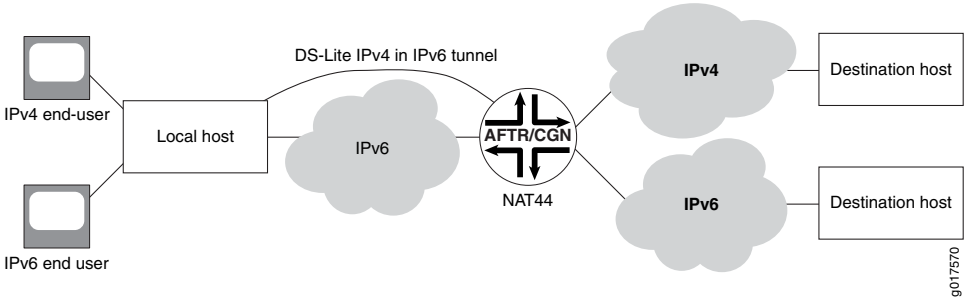

DS-Lite employs IPv4-over-IPv6 tunnels to cross an IPv6 access network to reach a carrier-grade IPv4-IPv4 NAT. This facilitates the phased introduction of IPv6 on the Internet by providing backward compatibility with IPv4.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring a DS-Lite Softwire Concentrator

Tunneling Services for IPv4-to-IPv6 Transition Overview

The Junos OS enables service providers to transition to IPv6 by using softwire encapsulation and decapsulation techniques. A softwire is a tunnel that is created between softwire Customer Premises Equipment (CPE). A softwire CPE can share a unique common internal state for multiple softwires, making it a very light and scalable solution. When you use softwires, you need not maintain an interface infrastructure for each softwire, unlike a typical mesh of generic routing encapsulation (GRE) tunnels that would require you to do so. A softwire initiator at the customer end encapsulates native packets and tunnels them to a softwire concentrator at the service provider. The softwire concentrator decapsulates the packets and sends them to their destination. A softwire is created when a softwire concentrator receives the first tunneled packet of a flow and prepares for flow processing. The softwire exists as long as the softwire concentrator is providing flows for routing. A flow counter is maintained; when the number of active flows is 0, the softwire is deleted. Statistics are kept for both flows and softwires.

Softwire addresses are not specifically configured under any physical or virtual interface. Therefore, the number of established softwires does not affect throughput, and scalability is independent of the number of interfaces. The scalability is only limited to the number of flows that the platform (services DPC or PIC) can support.

This topic contains the following sections:

- 6to4 Overview on page 53
- DS-Lite Softwires—IPv4 over IPv6 on page 55
- 6rd Softwires—IPv6 over IPv4 on page 56

6to4 Overview

- Basic 6to4 on page 53
- 6to4 Anycast on page 54
- 6to4 Provider-Managed Tunnels on page 54

Basic 6to4

6to4 is an Internet transition mechanism for migrating from IPv4 to IPv6, a system that allows IPv6 packets to be transmitted over an IPv4 network (generally the IPv4 Internet) without the need to configure explicit tunnels. 6to4 is described in RFC 3056, *Connection of IPv6 Domains via IPv4 Clouds*. 6to4 is especially relevant during the initial phases of deployment to full, native IPv6 connectivity, since IPv6 is not required on nodes between the host and the destination. However, it is intended only as a transition mechanism and is not meant to be used permanently.

6to4 can be used by an individual host, or by a local IPv6 network. When used by a host, it must have a global IPv4 address connected, and the host is responsible for the encapsulation of outgoing IPv6 packets and the decapsulation of incoming 6to4 packets.

Copyright © 2013, Juniper Networks, Inc.

If the host is configured to forward packets for other clients, often a local network, it is then a router.

There are two kinds of 6to4 virtual routers: border routers and relay routers. A 6to4 border router is an IPv6 router supporting a 6to4 pseudointerface, and is normally the border router between an IPv6 site and a wide-area IPv4 network. A relay router is a 6to4 router configured to support transit routing between 6to4 addresses and pure native IPv6 addresses.

In order for a 6to4 host to communicate with the native IPv6 Internet, its IPv6 default gateway must be set to a 6to4 address which contains the IPv4 address of a 6to4 relay router. To avoid the need for users to set this up manually, the Anycast address of 192.88.99.1 has been allocated to send packets to a 6to4 relay router. Note that when wrapped in 6to4 with the subnet and hosts fields set to zero, this IPv4 address (192.88.99.1) becomes the IPv6 address 2002:c058:6301::. To ensure BGP routing propagation, a short prefix of 192.88.99.0/24 has been allocated for routes pointed at 6to4 relay routers that use this Anycast IP address. Providers willing to provide 6to4 service to their clients or peers should advertise the Anycast prefix like any other IP prefix, and route the prefix to their 6to4 relay.

Packets from the IPv6 Internet to 6to4 systems must be sent to a 6to4 relay router by normal IPv6 routing methods. The specification states that such relay routers must only advertise 2002::/16 and not subdivisions of it to prevent IPv4 routes from polluting the routing tables of IPv6 routers. From there they can then be sent over the IPv4 Internet to the destination.

6to4 Anycast

Router 6to4 assumes that 6to4 routers and relays are managed and configured cooperatively. In particular, 6to4 sites must configure a relay router to carry the outbound traffic, which becomes the default IPv6 router (except for 2002::/16). The objective of the Anycast variant, defined in RFC 3068, *An Anycast Prefix for 6to4 Relay Routers*, is to avoid the need for such configuration. This makes the solution available for small or domestic users, even those with a single host or simple home gateway instead of a border router. This is achieved by defining 192.88.99.1 as the default IPv4 address for a 6to4 relay, and 2002:c058:6301:: as the default IPv6 router prefix (“well-known prefix”) for a 6to4 site.

RFC 6343, *Advisory Guidelines for 6to4 Deployment*, published in August 2011, identifies a wide range of problems associated with the use of unmanaged 6to4 Anycast relay routers.

6to4 Provider-Managed Tunnels

A solution to many problems associated with unmanaged Anycast 6to4 is presented in IETF informational draft draft-kuarsingh-v6ops-6to4-provider-managed-tunnel-02, *6to 4 Provider-Managed Tunnels (PMT)*. That document, a “work in progress,” proposes a solution that allows providers to exercise greater control over the routing of 6to4 traffic.

Anycast 6to4 implies a default configuration for the user site. It does not require any particular user action. It does require an IPv4 Anycast route to be in place to a relay at

Copyright © 2013, Juniper Networks, Inc.

192.88.99.1. Traffic does not necessarily return to the same 6to4 gateway because of the the “well-known” 6to4 prefix used and advertised by all 6to4 traffic.

6to4 provider-managed tunnels (PMTs) facilitate the management of 6to4 tunnels using an Anycast configuration. 6to4 PMT enables service providers to improve 6to4 operation when network conditions provide suboptimal performance or break normal 6to4 operation. 6to4 PMT provides a stable provider prefix and forwarding environment by utilizing existing 6to4 relays with an added function of IPv6 prefix translation that controls the flow of return traffic.

The 6to4 managed tunnel model behaves like a standard 6to4 service between the customer IPv6 host or gateway and the 6to4 PMT relay (within the provider domain). The 6to4-PMT relay shares properties with 6rd (RFC5969) by decapsulating and forwarding embedded IPv6 flows, within an IPv4 packet, to the IPv6 Internet. The model provides an additional function which translates the source 6to4 prefix to a provider assigned prefix which is not found in 6rd (RFC5969) or traditional 6to4 operation. The 6to4-PMT relay provides a stateless (or stateful) mapping of the 6to4 prefix to a provider-supplied prefix by mapping the embedded IPv4 address in the 6to4 prefix to the provider prefix.

DS-Lite Softwires—IPv4 over IPv6

When an Internet service provider (ISP) begins to allocate new subscriber home IPv6 addresses and IPv6-capable equipment, dual-stack lite (DS-Lite) provides a method for the private IPv4 addresses behind the IPv6 customer edge (CE) WAN equipment to reach the IPv4 network. DS-Lite enables IPv4 customers to continue to access the Internet using their current hardware by using a softwire initiator, referred to as a Basic Bridging Broadband (B4), at the customer edge to encapsulate IPv4 packets into IPv6 packets and tunnel them over an IPv6 network to a softwire concentrator, referred to as an Address Family Transition Router (AFTR), for decapsulation. DS-Lite creates the IPv6 softwires that terminate on the services PIC. Packets coming out of the softwire can then have other services such as NAT applied on them.

DS-Lite is supported on Multiservices 100, 400, and 500 PICs on M Series routers and on MX Series routers equipped with Multiservices Dense Port Concentrator (DPCs).



NOTE: **IPv6 Provider Edge (6PE), or MPLS-enabled IPv6, is available for ISPs with MPLS-enabled networks. These networks now can use Multiprotocol BGP (MP-BGP) to provide connectivity between the DS-Lite B4 and AFTR (or any two IPv6 nodes). DS-Lite properly handles encapsulation and decapsulation despite the presence of additional MPLS header information.**

For more information on DS-Lite softwires, see the IETF draft *Dual Stack Lite Broadband Deployments Following IPv4 Exhaustion*.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The most recent IETF draft documentation for DS-Lite uses new terminology:

- The term *softwire initiator* has been replaced by *B4*.
- The term *softwire concentrator* has been replaced by *AFTR*.

The Junos OS documentation generally uses the original terms when discussing configuration in order to be consistent with the command-line interface (CLI) statements used to configure DS-Lite.

6rd Softwires—IPv6 over IPv4

6rd softwire flow is shown in Figure 5 on page 56.

Figure 5: 6rd Softwire Flow

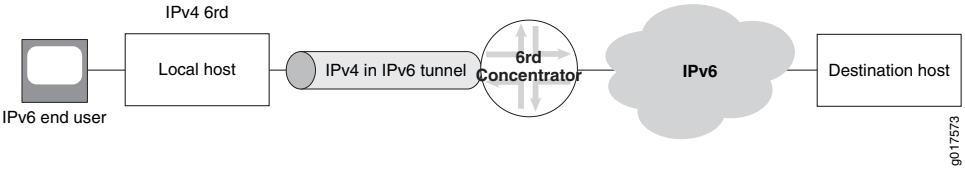

The Junos OS supports a 6rd softwire concentrator on a services DPC or PIC to facilitate rapid deployment of IPv6 service to subscribers on native IPv4 CE WANs. IPv6 packets are encapsulated in IPv4 packets by a softwire initiator at the CE WAN. These packets are tunneled to a softwire concentrator residing on a multiservices DPC (branch relay). A softwire is created when IPv4 packets containing IPv6 destination information are received at the softwire concentrator, which decapsulates IPv6 packets and forwards them for IPv6 routing. All of these functions are performed in a single pass of the services PIC.

In the reverse path, IPv6 packets are sent to the Services DPC where they are encapsulated in IPv4 packets corresponding to the proper softwire and sent to the CE WAN.

The softwire concentrator creates softwires as the IPv4 packets are received from the CE WAN side or IPV6 packets are received from the Internet. A 6rd softwire on the Services DPC is identified by the 3-tuple containing the service set ID, CE softwire initiator IPv4 address, and softwire concentrator IPv4 address. IPv6 flows are also created for the encapsulated IPv6 payload, and are associated with the specific softwire that carried them in the first place. When the last IPv6 flow associated with a softwire ends, the softwire is deleted. This simplifies configuration and there is no need to create or manage tunnel interfaces.

6rd is supported on Multiservices 100, 400, and 500 PICs on M Series and T Series routers, and on MX Series platforms equipped with Multiservices DPCs.

For more information on 6rd softwires, see RFC 5969, *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification*.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Network Address Translation Overview on page 48

IPsec Overview

The Juniper Networks Junos OS supports IPsec. This section discusses the following topics, which provide background information about configuring IPsec.

For a list of the IPsec and IKE standards supported by the Junos OS, see the *Junos OS Hierarchy and RFC Reference*.

- IPsec on page 57
- Security Associations on page 57
- IKE on page 58
- Comparison of IPsec Services and ES Interface Configuration on page 58

IPsec

The IPsec architecture provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. In addition to IPsec, the Junos OS also supports the Internet Key Exchange (IKE), which defines mechanisms for key generation and exchange, and manages security associations (SAs).

IPsec also defines a security association and key management framework that can be used with any network layer protocol. The SA specifies what protection policy to apply to traffic between two IP-layer entities. IPsec provides secure tunnels between two peers.

Security Associations

To use IPsec security services, you create SAs between hosts. An SA is a simplex connection that allows two hosts to communicate with each other securely by means of IPsec. There are two types of SAs:

- Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. Manual SAs statically define the security parameter index (SPI) values, algorithms, and keys to be used, and require matching configurations on both ends of the tunnel. Each peer must have the same configured options for communication to take place.
- Dynamic SAs require additional configuration. With dynamic SAs, you configure IKE first and then the SA. IKE creates dynamic security associations; it negotiates SAs for IPsec. The IKE configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway. This connection is then used to dynamically agree upon keys and other data used by the dynamic IPsec SA. The IKE SA is negotiated first and then used to protect the negotiations that determine the dynamic IPsec SAs.

Copyright © 2013, Juniper Networks, Inc.

IKE

IKE is a key management protocol that creates dynamic SAs; it negotiates SAs for IPsec. An IKE configuration defines the algorithms and keys used to establish a secure connection with a peer security gateway.

IKE performs the following tasks:

- Negotiates and manages IKE and IPsec parameters.
- Authenticates secure key exchange.
- Provides mutual peer authentication by means of shared secrets (not passwords) and public keys.
- Provides identity protection (in main mode).

Two versions of the IKE protocol (IKEv1 and IKEv2) are supported now. IKE negotiates security attributes and establishes shared secrets to form the bidirectional IKE SA. In IKE, inbound and outbound IPsec SAs are established and the IKE SA secures the exchanges. Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. IKE also generates keying material, provides Perfect Forward Secrecy, and exchanges identities.

Comparison of IPsec Services and ES Interface Configuration

Table 4 on page 58 compares the top-level configuration of IPsec features on the ES PIC interfaces and on the AS or MultiServices PIC interfaces.

Table 4: Statement Equivalents for ES and AS Interfaces

ES PIC Configuration	AS and MultiServices PIC IPsec Configuration
[edit security ipsec] proposal {...}	[edit services ipsec-vpn ipsec] proposal {...}
[edit security ipsec] policy {...}	[edit services ipsec-vpn ipsec] policy {...}
[edit security ipsec] security-association sa-dynamic {...}	[edit services ipsec-vpn rule *rule-name*] term *term-name* match-conditions {...} then dynamic {...}]
[edit security ipsec] security-association sa-manual {...}	[edit services ipsec-vpn rule *rule-name*] term *term-name* match-conditions {...} then manual {...}]
[edit security ike] proposal {...}	[edit services ipsec-vpn ike] proposal {...}
[edit security ike] policy {...}	[edit services ipsec-vpn ike] policy {...}
Not available	[edit services ipsec-vpn] rule-set {...}

Copyright © 2013, Juniper Networks, Inc.

Table 4: Statement Equivalents for ES and AS Interfaces *(continued)*

ES PIC Configuration	AS and MultiServices PIC IPsec Configuration
Not available	[edit services ipsec-vpn] service-set {...}
[edit interfaces es-*fpc/pic/port*] tunnel source *address*	[edit services ipsec-vpn service-set *set-name* ipsec-vpn local-gateway *address*]
[edit interfaces es-*fpc/pic/port*] tunnel destination *address*	[edit services ipsec-vpn rule *rule-name*] remote-gateway *address*

For more information about configuring IPsec services on an AS or MultiServices PIC, see IPsec Properties. For more information about configuring encryption services on an ES PIC, see "Configuring Encryption Interfaces" on page 881.

NOTE: Although many of the same statements and properties are valid on both platforms, the configurations are not interchangeable. You must commit a complete configuration for the PIC type that is installed in your router.

Layer 2 Tunneling Protocol Overview

L2TP is defined in RFC 2661,*Layer Two Tunneling Protocol (L2TP)*.

L2TP facilitates the tunneling of PPP packets across an intervening network in a way that is as transparent as possible to both end users and applications. It employs access profiles for group and individual user access, and uses authentication to establish secure connections between the two ends of each tunnel. Multilink PPP functionality is also supported.

The L2TP services are supported on the following routers only:

- M7i routers with AS PICs
- M10i routers with AS and MultiServices 100 PICs
- M120 routers with AS, MultiServices 100, and MultiServices 400 PICs

For more information, see "L2TP Services Configuration Overview" on page 311.

Voice Services Overview

Adaptive services interfaces include a voice services feature that allows you to specify interface type **lsq-*fpc/pic/port*** to accommodate voice over IP (VoIP) traffic. This interface uses compressed RTP (CRTP), which is defined in RFC 2508, *Compressing IP/UDP/RTP Headers for Low-Speed Serial Links*.

CRTP enables VoIP traffic to use low-speed links more effectively, by compressing the 40-byte IP/UDP/RTP header down to 2 to 4 bytes in most cases.

Copyright © 2013, Juniper Networks, Inc.

Voice services on the AS and MultiServices PICs support single-link PPP-encapsulated IPv4 traffic over the following physical interface types: ATM2, DS3, E1, E3, OC3, OC12, STM1, and T1, including the channelized versions of these interfaces.

Voice services do not require a separate service rules configuration.

Voice services also support LFI on Juniper Networks M Series Multiservice Edge routers, except the M320 router. For more information about configuring voice services, see "Configuring Services Interfaces for Voice Services" on page 422.

For link services IQ interfaces (**lsq**) only, you can configure CRTP with multiclass MLPPP (MCML). MCML greatly simplifies packet ordering issues that occur when multiple links are used. Without MCML, all voice traffic belonging to a single flow is hashed to a single link in order to avoid packet ordering issues. With MCML, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about MCML support on link services IQ interfaces, see "Configuring Link Services and CoS on Services PICs" on page 377.

HTTP URL Tracking and Policy Control for Client Requests

The URL manipulation capability in the service plane allows an administrator to enter all the URLs associated with an action. An administrator can now enter URLs and actions in the service plane and apply them to the traffic associated with an interface and a subscriber by using the existing functions.



NOTE: The URL manipulation capability is supported only when the Junos OS Extension-Provider packages are installed and configured on the device.

In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.

The HTTP URL manipulation allows routers to:

- Monitor HTTP transactions in the service plane and identify matching incoming HTTP requests with preconfigured URLs. When a matching HTTP request is found, the associated action and preconfigured URL are applied to the transaction.
- Maintain the statistics of HTTP requests that match the preconfigured URLs. Such statistics must be maintained per policy on a per-rule or per-term basis across all service sets.
- Display the current statistics for all visible HTTP-URL matches.
- Scale 50,000 HTTP transactions per second on a single MS-DPC network processing unit.

Guidelines for Configuring HTTP URL Monitoring for Client Requests

HTTP URL manipulation can be configured with service rules containing a sequence of terms. The URL rules are evaluated based on the longest prefix or suffix matches. The rules and terms are configured if:

Copyright © 2013, Juniper Networks, Inc.

- Multiple hostnames are specified in a match condition with an OR operator (that is, “any” is a potential match).
- Multiple request-URLs are specified in a match condition with an OR operator (that is, “any” is a potential match).
- The “discard” action causes HTTP requests matched to URLs to be dropped (stateful).
- The “accept” action causes HTTP requests matched to URLs to be allowed (stateful).
- The “log-request” action causes HTTP requests matched to URLs to be logged (fast log, once per transaction).
- The “count” action causes HTTP requests matched to URLs to be counted against the specified rule or term. The total number is the total of all service sets in which the rule or terms exists.
- The “discard” action can be combined only with the “log-request” action.
- The “accept” action can be combined with the “log-request” and “count” actions.
- If any given HTTP request matches more than one rule or term, the action applied is undeterministic and may be any of the matched rules or terms.
- The number of rules, terms, url-lists, and individual clauses in the rules, terms, and url-lists are limited by service-plane memory. No hard limits are imposed.
- If a hostname is not specified, then “*” (any) is assumed. Similarly, if a request-URL is not specified, then “*” (any) is assumed. Matching of hostnames and request-URLs to the HTTP request follows the same process described for url-lists. However, a match happens only when both the hostname and the request-URL match an entry in the same term of a rule.
- Multiple “url” and “url-list” clauses may be entered for the same “from {}” clause.

Configuring HTTP URL Tracking and Policy Control

To configure an HTTP URL Tracking and Policy Control, include the **url-rule *url-rule-name*** statement at the **[edit services hcm]** hierarchy level:

```
services {
  hcm {
    url-rule url-rule-name {
      term term-num {
        from {
          url-list url-list-name;
          url url_identifier {
            host hostname;
            request-url page-name;
          }
        }
        then {
          discard;
          accept;
          count;
          log-request;
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
    url-rule-set url-rule-set-name {
      url-rule rule1;
      url-rule rule2;
    }
  }
}
```

Each HCM rule consists of a set of terms, similar to a firewall filter. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

Class of Service Overview

The CoS configuration available for the AS PIC enables you to configure Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets transiting the AS PIC. You can configure the CoS service alongside the stateful firewall and NAT services, using a similar rule structure. The component structures are described in detail in the Junos OS Class of Service Configuration Guide.

Standards for Differentiated Services are described in the following documents:

- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2475, *An Architecture for Differentiated Services*



NOTE: **CoS BA classification is not supported on services interfaces.**

For more information about configuring CoS services, see Class-of-Service Properties.

Examples: Services Interfaces Configuration

This section includes the following examples:

- Example: Service Interfaces Configuration on page 63
- Example: VPN Routing and Forwarding (VRF) and Service Configuration on page 65
- Example: Dynamic Source NAT as a Next-Hop Service on page 67
- Example: NAT Between VRFs Configuration on page 68
- Example: BOOTP and Broadcast Addresses on page 71

Copyright © 2013, Juniper Networks, Inc.

Example: Service Interfaces Configuration

The following configuration includes all the items necessary to configure services on an interface. For examples showing individual service configurations, see the chapters that describe each service in detail.

```
[edit]
interfaces {
    fe-0/1/0 {
        unit 0 {
            family inet {
                service {
                    input {
                        service-set Firewall-Set;
                    }
                    output {
                        service-set Firewall-Set;
                    }
                }
                address 10.1.3.2/24;
            }
        }
    }
    fe-0/1/1 {
        unit 0 {
            family inet {
                filter {
                    input Sample;
                }
                address 172.16.1.2/24;
            }
        }
    }
    sp-1/0/0 {
        unit 0 {
            family inet {
                address 172.16.1.3/24 {
                }
            }
        }
    }
}
forwarding-options {
    sampling {
        input {
            family inet {
                rate 1;
            }
        }
        output {
            cflowd 10.1.3.1 {
                port 2055;
                version 5;
            }
            flow-inactive-timeout 15;
```

Copyright © 2013, Juniper Networks, Inc.

```
        flow-active-timeout 60;
        interface sp-1/0/0 {
          engine-id 1;
          engine-type 136;
          source-address 10.1.3.2;
        }
      }
    }
}
firewall {
  filter Sample {
    term Sample {
      then {
        count Sample;
        sample;
        accept;
      }
    }
  }
}
services {
  stateful-firewall {
    rule Rule1 {
      match-direction input;
      term 1 {
        from {
          application-sets Applications;
        }
        then {
          accept;
        }
      }
      term accept {
        then {
          accept;
        }
      }
    }
    rule Rule2 {
      match-direction output;
      term Local {
        from {
          source-address {
            10.1.3.2/32;
          }
        }
        then {
          accept;
        }
      }
    }
  }
  ids {
    rule Attacks {
      match-direction output;
      term Match {
```


Copyright © 2013, Juniper Networks, Inc.

```
            from {
              application-sets Applications;
            }
            then {
              logging {
                syslog;
              }
            }
          }
        }
      }
      nat {
        pool public {
          address-range low 172.16.2.1 high 172.16.2.32;
          port automatic;
        }
        rule Private-Public {
          match-direction input;
          term Translate {
            then {
              translated {
                source-pool public;
                translation-type source dynamic;
              }
            }
          }
        }
      }
      service-set Firewall-Set {
        stateful-firewall-rules Rule1;
        stateful-firewall-rules Rule2;
        nat-rules Private-Public;
        ids-rules Attacks;
        interface-service {
          service-interface sp-1/0/0;
        }
      }
    }
    applications {
      application ICMP {
        application-protocol icmp;
      }
      application FTP {
        application-protocol ftp;
        destination-port ftp;
      }
      application-set Applications {
        application ICMP;
        application FTP;
      }
    }
```

Example: VPN Routing and Forwarding (VRF) and Service Configuration

The following example combines VPN routing and forwarding (VRF) and services configuration:

Copyright © 2013, Juniper Networks, Inc.

```
[edit policy-options]
policy-statement test-policy {
   term t1 {
      then reject;
   }
}
[edit routing-instances]
test {
   interface ge-0/2/0.0;
   interface sp-1/3/0.20;
   instance-type vrf;
   route-distinguisher 10.58.255.1:37;
   vrf-import test-policy;
   vrf-export test-policy;
   routing-options {
      static {
         route 0.0.0.0/0 next-table inet.0;
      }
   }
}
[edit interfaces]
ge-0/2/0 {
   unit 0 {
      family inet {
         service {
            input service-set nat-me;
            output service-set nat-me;
         }
      }
   }
}
sp-1/3/0 {
   unit 0 {
      family inet;
   }
   unit 20 {
      family inet;
      service-domain inside;
   }
   unit 21 {
      family inet;
      service-domain outside;
   }
   [edit services]
   stateful-firewall {
      rule allow-any-input {
         match-direction input;
         term t1 {
            then accept;
         }
      }
   }
   nat {
      pool hide-pool {
         address 10.58.16.100;
         port automatic;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    rule hide-all-input {
      match-direction input;
      term t1 {
        then {
          translated {
            source-pool hide-pool;
            translation-type source dynamic;
          }
        }
      }
    }
  }
  service-set nat-me {
    stateful-firewall-rules allow-any-input;
    nat-rules hide-all-input;
    interface-service {
      service-interface sp-1/3/0.20;
    }
  }
}
```

Example: Dynamic Source NAT as a Next-Hop Service

The following example shows dynamic-source NAT applied as a next-hop service:

```
[edit interfaces]
ge-0/2/0 {
  unit 0 {
    family mpls;
  }
}
sp-1/3/0 {
  unit 0 {
    family inet;
  }
  unit 20 {
    family inet;
  }
  unit 32 {
    family inet;
  }
}
[edit routing-instances]
protected-domain {
  interface ge-0/2/0.0;
  interface sp-1/3/0.20;
  instance-type vrf;
  route-distinguisher 10.58.255.17:37;
  vrf-import protected-domain-policy;
  vrf-export protected-domain-policy;
  routing-options {
    static {
      route 0.0.0.0/0 next-hop sp-1/3/0.20;
    }
  }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
[edit policy-options]
policy-statement protected-domain-policy {
   term t1 {
      then reject;
   }
}
[edit services]
stateful-firewall {
   rule allow-all {
      match-direction input;
      term t1 {
         then {
            accept;
         }
      }
   }
}
nat {
   pool my-pool {
      address 10.58.16.100;
      port automatic;
   }
   rule hide-all {
      match-direction input;
      term t1 {
         then {
            translated {
               source-pool my-pool;
               translation-type source dynamic;
            }
         }
      }
   }
}
service-set null-sfw-with-nat {
   stateful-firewall-rules allow-all;
   nat-rules hide-all;
   next-hop-service {
      inside-service-interface sp-1/3/0.20;
      outside-service-interface sp-1/3/0.32;
   }
}
```

Example: NAT Between VRFs Configuration

The following example configuration enables NAT between VRFs with overlapping private addresses, using distinct public addresses for the source and destination NAT in this scenario:

- A host in **vrf-a** traverses **10.58.16.201** to reach **10.58.0.2** in **vrf-b**.
- A host in **vrf-b** traverses **10.58.16.101** to reach **10.58.0.2** in **vrf-a**.

```
[edit interfaces]
ge-0/2/0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family inet {
        address 10.58.0.1/24;
        service {
          input service-set vrf-a-svc-set;
          output service-set vrf-a-svc-set;
        }
      }
    }
}
ge-0/3/0 {
    unit 0 {
      family inet {
        address 10.58.0.1/24;
        service {
          input service-set vrf-b-svc-set;
          output service-set vrf-b-svc-set;
        }
      }
    }
}
sp-1/3/0 {
    unit 0 {
      family inet;
    }
    unit 10 {
      family inet;
      service-domain inside;
    }
    unit 20 {
      family inet;
      service-domain inside;
    }
}
[edit policy-options]
policy-statement test-policy {
    term t1 {
      then reject;
    }
}
[edit routing-instances]
vrf-a {
    interface ge-0/2/0.0;
    interface sp-1/3/0.10;
    instance-type vrf;
    route-distinguisher 10.1.1.1:1;
    vrf-import test-policy;
    vrf-export test-policy;
    routing-options {
      static {
        route 0.0.0.0/0 next-table inet.0;
      }
    }
}
vrf-b {
    interface ge-0/3/0.0;
```

Copyright © 2013, Juniper Networks, Inc.

```
    interface sp-1/3/0.20;
    instance-type vrf;
    route-distinguisher 10.2.2.2:2;
    vrf-import test-policy;
    vrf-export test-policy;
    routing-options {
      static {
        route 0.0.0.0/0 next-table inet.0;
      }
    }
}
[edit services]
stateful-firewall {
    rule allow-all {
      match-direction input-output;
      term t1 {
        then {
          accept;
        }
      }
    }
}
nat {
    pool vrf-a-src-pool {
      address 10.58.16.100;
      port automatic;
    }
    pool vrf-a-dst-pool {
      address 10.58.0.2;
    }
    rule vrf-a-input {
      match-direction input;
      term t1 {
        then {
          translated {
            source-pool vrf-a-src-pool;
            translation-type napt-44;
          }
        }
      }
    }
    rule vrf-a-output {
      match-direction output;
      term t1 {
        from {
          destination-address 10.58.16.101;
        }
        then {
          translated {
            destination-pool vrf-a-dst-pool;
            translation-type destination static;
          }
        }
      }
    }
    pool vrf-b-src-pool {
```


Copyright © 2013, Juniper Networks, Inc.

```
    address 10.58.16.200;
    port automatic;
  }
  pool vrf-b-dst-pool {
    address 10.58.0.2;
  }
  rule vrf-b-input {
    match-direction input;
    term t1 {
      then {
        translated {
          source-pool vrf-b-src-pool;
          translation-type source dynamic;
        }
      }
    }
  }
  rule vrf-b-output {
    match-direction output;
    term t1 {
      from {
        destination-address 10.58.16.201;
      }
      then {
        translated {
          destination-pool vrf-b-dst-pool;
          translation-type destination static;
        }
      }
    }
  }
}
service-set vrf-a-svc-set {
  stateful-firewall-rules allow-all;
  nat-rules vrf-a-input;
  nat-rules vrf-a-output;
  interface-service {
    service-interface sp-1/3/0.10;
  }
}
service-set vrf-b-svc-set {
  stateful-firewall-rules allow-all;
  nat-rules vrf-b-input;
  nat-rules vrf-b-output;
  interface-service {
    service-interface sp-1/3/0.20;
  }
}
```

Example: BOOTP and Broadcast Addresses

The following example supports Bootstrap Protocol (BOOTP) and broadcast addresses:

```
[edit applications]
application bootp {
  application-protocol bootp;
```

Copyright © 2013, Juniper Networks, Inc.

```
    protocol udp;
    destination-port 67;
}
[edit services]
stateful-firewall bootp-support {
    rule bootp-allow {
        direction input;
        term bootp-allow {
            from {
                destination-address {
                    any-unicast;
                    255.255.255.255;
                }
                application bootp;
            }
            then {
                accept;
            }
        }
    }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 4

Applications Configuration Guidelines

You can define application protocols for the stateful firewall and Network Address Translation (NAT) services to use in match condition rules. An application protocol, or application layer gateway (ALG), defines application parameters using information from network Layer 3 and above. Examples of such applications are FTP and H.323.

To configure applications that are used with services, include the following statements at the **[edit applications]** hierarchy level:

```
[edit applications]
application application-name {
    application-protocol protocol-name;
    destination-port port-number;
    icmp-code value;
    icmp-type value;
    inactivity-timeout value;
    learn-sip-register;
    protocol type;
    rpc-program-number number;
    sip-call-hold-timeout seconds;
    snmp-command command;
    source-port port-number;
    ttl-threshold value;
    uuid hex-value;
}
application-set application-set-name {
    application application-name;
}
```

This chapter includes the following sections:

- Configuring Application Protocol Properties on page 74
- Configuring Application Sets on page 85
- ALG Descriptions on page 85
- Verifying the Output of ALG Sessions on page 105
- Junos Default Groups on page 111
- Examples: Configuring Application Protocols on page 118

Copyright © 2013, Juniper Networks, Inc.

Configuring Application Protocol Properties

To configure application properties, include the **application** statement at the **[edit applications]** hierarchy level:

```
[edit applications]
application application-name {
  application-protocol protocol-name;
  destination-port port-number;
  icmp-code value;
  icmp-type value;
  inactivity-timeout value;
  protocol type;
  rpc-program-number number;
  snmp-command command;
  source-port port-number;
  ttl-threshold value;
  uuid hex-value;
}
```

You can group application objects by configuring the **application-set** statement; for more information, see "Configuring Application Sets" on page 85.

This section includes the following tasks for configuring applications:

- Configuring an Application Protocol on page 74
- Configuring the Network Protocol on page 76
- Configuring the ICMP Code and Type on page 78
- Configuring Source and Destination Ports on page 79
- Configuring the Inactivity Timeout Period on page 82
- Configuring SIP on page 82
- Configuring an SNMP Command for Packet Matching on page 84
- Configuring an RPC Program Number on page 84
- Configuring the TTL Threshold on page 84
- Configuring a Universal Unique Identifier on page 84

Configuring an Application Protocol

The **application-protocol** statement allows you to specify which of the supported application protocols (ALGs) to configure and include in an application set for service processing. To configure application protocols, include the **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
application-protocol protocol-name;
```

Table 5 on page 75 shows the list of supported protocols. For more information about specific protocols, see "ALG Descriptions" on page 85.

Copyright © 2013, Juniper Networks, Inc.

Table 5: Application Protocols Supported by Services Interfaces

Protocol Name	CLI Value	Comments
Bootstrap protocol (BOOTP)	**bootp**	Supports BOOTP and dynamic host configuration protocol (DHCP).
Distributed Computing Environment (DCE) remote procedure call (RPC)	**dce-rpc**	Requires the **protocol** statement to have the value **udp** or **tcp**. Requires a **uuid** value. You cannot specify **destination-port** or **source-port** values.
DCE RPC portmap	**dce-rpc-portmap**	Requires the **protocol** statement to have the value **udp** or **tcp**. Requires a **destination-port** value.
Domain Name System (DNS)	**dns**	Requires the **protocol** statement to have the value **udp**. This application protocol closes the DNS flow as soon as the DNS response is received.
Exec	**exec**	Requires the **protocol** statement to have the value **tcp** or to be unspecified. Requires a **destination-port** value.
FTP	**ftp**	Requires the **protocol** statement to have the value **tcp** or to be unspecified. Requires a **destination-port** value.
H.323	**h323**	–
Internet Control Message Protocol (ICMP)	**icmp**	Requires the **protocol** statement to have the value **icmp** or to be unspecified.
Internet Inter-ORB Protocol	**iiop**	–
IP	**ip**	–
Login	**login**	–
NetBIOS	**netbios**	Requires the **protocol** statement to have the value **udp** or to be unspecified. Requires a **destination-port** value.
NetShow	**netshow**	Requires the **protocol** statement to have the value **tcp** or to be unspecified. Requires a **destination-port** value.
Point-to-Point Tunneling Protocol	**pptp**	–
RealAudio	**realaudio**	–
Real-Time Streaming Protocol (RTSP)	**rtsp**	Requires the **protocol** statement to have the value **tcp** or to be unspecified. Requires a **destination-port** value.
RPC User Datagram Protocol (UDP) or TCP	**rpc**	Requires the **protocol** statement to have the value **udp** or **tcp**. Requires a **rpc-program-number** value. You cannot specify **destination-port** or **source-port** values.
RPC port mapping	**rpc-portmap**	Requires the **protocol** statement to have the value **udp** or **tcp**. Requires a **destination-port** value.

Copyright © 2013, Juniper Networks, Inc.

Table 5: Application Protocols Supported by Services Interfaces *(continued)*

Protocol Name	CLI Value	Comments
Shell	**shell**	Requires the **protocol** statement to have the value **tcp** or to be unspecified. Requires a **destination-port** value.
Session Initiation Protocol	**sip**	–
SNMP	**snmp**	Requires the **protocol** statement to have the value **udp** or to be unspecified. Requires a **destination-port** value.
SQLNet	**sqlnet**	Requires the **protocol** statement to have the value **tcp** or to be unspecified. Requires a **destination-port** or **source-port** value.
Talk Program	**talk**	
Trace route	**traceroute**	Requires the **protocol** statement to have the value **udp** or to be unspecified. Requires a **destination-port** value.
Trivial FTP (TFTP)	**tftp**	Requires the **protocol** statement to have the value **udp** or to be unspecified. Requires a **destination-port** value.
WinFrame	**winframe**	–



NOTE: You can configure application-level gateways (ALGs) for ICMP and trace route under stateful firewall, NAT, or CoS rules when twice NAT is configured in the same service set. These ALGs cannot be applied to flows created by the Packet Gateway Controller Protocol (PGCP). Twice NAT does not support any other ALGs. NAT applies only the IP address and TCP or UDP headers, but not the payload.

For more information about configuring twice NAT, see Network Address Translation.

Configuring the Network Protocol

The **protocol** statement allows you to specify which of the supported network protocols to match in an application definition. To configure network protocols, include the **protocol** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
protocol type;
```

You specify the protocol type as a numeric value; for the more commonly used protocols, text names are also supported in the command-line interface (CLI). Table 6 on page 77 shows the list of the supported protocols.

Copyright © 2013, Juniper Networks, Inc.

Table 6: Network Protocols Supported by Services Interfaces

Network Protocol Type	CLI Value	Comments
IP Security (IPsec) authentication header (AH)	**ah**	–
External Gateway Protocol (EGP)	**egp**	–
IPsec Encapsulating Security Payload (ESP)	**esp**	–
Generic routing encapsulation (GR)	**gre**	–
ICMP	**icmp**	Requires an **application-protocol** value of **icmp**.
Internet Group Management Protocol (IGMP)	**igmp**	–
IP in IP	**ipip**	–
OSPF	**ospf**	–
Protocol Independent Multicast (PIM)	**pim**	–
Resource Reservation Protocol (RSVP)	**rsvp**	–
TCP	**tcp**	Requires a **destination-port** or **source-port** value unless you specify **application-protocol rcp** or **dce-rcp**.
UDP	**udp**	Requires a **destination-port** or **source-port** value unless you specify **application-protocol rcp** or **dce-rcp**.
Virtual Router Redundancy Protocol (VRRP)	**vrrp**	–

For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.



NOTE: IP version 6 (IPv6) is not supported as a network protocol in application definitions.

By default, the twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages. You can include the **protocol tcp** and **protocol udp** statements with the application statement for twice NAT configurations. For more information about configuring twice NAT, see Network Address Translation.

Copyright © 2013, Juniper Networks, Inc.

Configuring the ICMP Code and Type

The ICMP code and type provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ICMP settings, include the **icmp-code** and **icmp-type** statements at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
icmp-code value;
icmp-type value;
```

You can include only one ICMP code and type value. The **application-protocol** statement must have the value **icmp**. Table 7 on page 78 shows the list of supported ICMP values.

Table 7: ICMP Codes and Types Supported by Services Interfaces

CLI Statement	Description
icmp-code	This value or keyword provides more specific information than **icmp-type**. Because the value's meaning depends upon the associated **icmp-type** value, you must specify **icmp-type** along with **icmp-code**. For more information, see the Routing Policy Configuration Guide. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated: parameter-problem: **ip-header-bad** (0), **required-option-missing** (1) redirect: **redirect-for-host** (1), **redirect-for-network** (0), **redirect-for-tos-and-host** (3), **redirect-for-tos-and-net** (2) time-exceeded: **ttl-eq-zero-during-reassembly** (1), **ttl-eq-zero-during-transit** (0) unreachable: **communication-prohibited-by-filtering** (13), **destination-host-prohibited** (10), **destination-host-unknown** (7), **destination-network-prohibited** (9), **destination-network-unknown** (6), **fragmentation-needed** (4), **host-precedence-violation** (14), **host-unreachable** (1), **host-unreachable-for-TOS** (12), **network-unreachable** (0), **network-unreachable-for-TOS** (11), **port-unreachable** (3), **precedence-cutoff-in-effect** (15), **protocol-unreachable** (2), **source-host-isolated** (8), **source-route-failed** (5)
icmp-type	Normally, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port. For more information, see the Routing Policy Configuration Guide. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): **echo-reply** (0), **echo-request** (8), **info-reply** (16), **info-request** (15), **mask-request** (17), **mask-reply** (18), **parameter-problem** (12), **redirect** (5), **router-advertisement** (9), **router-solicit** (10), **source-quench** (4), **time-exceeded** (11), **timestamp** (13), **timestamp-reply** (14), or **unreachable** (3).

Copyright © 2013, Juniper Networks, Inc.

NOTE: If you configure an interface with an input firewall filter that includes a reject action and with a service set that includes stateful firewall rules, the router executes the input firewall filter before the stateful firewall rules are run on the packet. As a result, when the Packet Forwarding Engine sends an ICMP error message out through the interface, the stateful firewall rules might drop the packet because it was not seen in the input direction.

Possible workarounds are to include a forwarding-table filter to perform the reject action, because this type of filter is executed after the stateful firewall in the input direction, or to include an output service filter to prevent the locally generated ICMP packets from going to the stateful firewall service.

Configuring Source and Destination Ports

The TCP or UDP source and destination port provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ports, include the **destination-port** and **source-port** statements at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
destination-port value;
source-port value;
```

You must define one source or destination port. Normally, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port; for constraints, see Table 5 on page 75.

You can specify either a numeric value or one of the text synonyms listed in Table 8 on page 79.

Table 8: Port Names Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67

Copyright © 2013, Juniper Networks, Inc.

Table 8: Port Names Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
domain	53
eklogin	2105
ekshell	2106
exec	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760
kshell	544
ldap	389
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137

Copyright © 2013, Juniper Networks, Inc.

Table 8: Port Names Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmptrap	162
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs-ds	65
talk	517
telnet	23
tftp	69

Copyright © 2013, Juniper Networks, Inc.

Table 8: Port Names Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
timed	525
who	513
xdmcp	177
zephyr-clt	2103
zephyr-hm	2104

For more information about matching criteria, see the Routing Policy Configuration Guide.

Configuring the Inactivity Timeout Period

You can specify a timeout period for application inactivity. If the software has not detected any activity during the duration, the flow becomes invalid when the timer expires. To configure a timeout period, include the **inactivity-timeout** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
inactivity-timeout seconds;
```

The default value is 30 seconds. The value you configure for an application overrides any global value configured at the **[edit interfaces *interface-name* service-options]** hierarchy level; for more information, see "Configuring Default Timeout Settings for Services Interfaces" on page 516.

Configuring SIP

The Session Initiation Protocol (SIP) is a generalized protocol for communication between endpoints involved in Internet services such as telephony, fax, video conferencing, instant messaging, and file exchange. The supported standard is described in RFC 3261, SIP: Session Initiation Protocol, which includes stateful firewall and Network Address Translation (NAT) support for SIP dialogs and UDP IPv4 transport of SIP messages.

To implement SIP on adaptive services interfaces, you configure the **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level with the value **sip**. For more information about this statement, see "Configuring an Application Protocol" on page 74. In addition, there are two other statements you can configure to modify how SIP is implemented:

- You can enable the router to accept any incoming SIP calls for the endpoint devices that are behind the NAT firewall. When a device behind the firewall registers with the proxy that is outside the firewall, the AS or Multiservices PIC maintains the registration state. When the learn-sip-register statement is enabled, the router can use this information to accept inbound calls. If this statement is not configured, no inbound calls are accepted; only the devices behind the firewall can call devices outside the firewall.

Copyright © 2013, Juniper Networks, Inc.

To configure SIP registration, include the **learn-sip-register** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
learn-sip-register;
```

You can also manually inspect the SIP register by issuing the **show services stateful-firewall sip-register** command; for more information, see the *Junos OS System Basics and Services Command Reference*.

- You can specify a timeout period for the duration of SIP calls that are placed on hold. When a call is put on hold, there is no activity and flows might time out after the configured **inactivity-timeout** period expires, resulting in call state teardown. To avoid this, when a call is put on hold, the flow timer is reset to the **sip-call-hold-timeout** cycle to preserve the call state and flows for longer than the **inactivity-timeout** period.

  To configure a timeout period, include the **sip-call-hold-timeout** statement at the **[edit applications application *application-name*]** hierarchy level:

  ```
  [edit applications application application-name]
  sip-call-hold-timeout seconds;
  ```

  The default value is 7200 seconds and the range is from 0 through 36,000 seconds (10 hours).

Limitations

The following limitations apply to configuration of the SIP ALG:

- Only the methods described in RFC 3261 are supported.
- Only SIP version 3 is supported.
- TCP is not supported as a transport mechanism for signaling messages.
- IPv6 signaling data is not supported.
- Authentication is not supported.
- Encrypted messages are not supported.
- SIP fragmentation is not supported.
- The maximum UDP packet size containing a SIP message is assumed to be 4 KB. SIP messages larger than this are not supported.
- The maximum number of media channels in a SIP message is assumed to be six.
- Fully qualified domain names (FQDNs) are not supported in critical fields.
- QoS is not supported.
- High availability is not supported, except for warm standby.
- A timeout setting of never is not supported on SIP or NAT.

Copyright © 2013, Juniper Networks, Inc.

- Multicast (forking proxy) is not supported.
- When clients use STUN/TURN to detect the firewall or NAT devices between the caller and responder or proxy, the client attempts to best-guess the NAT device behavior and act accordingly to place the call. In such cases, you should not configure the ALG.

Configuring an SNMP Command for Packet Matching

You can specify an SNMP command setting for packet matching. To configure SNMP, include the **snmp-command** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
snmp-command value;
```

The supported values are **get**, **get-next**, **set**, and **trap**. You can configure only one value for matching. The **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level must have the value **snmp**. For information about specifying the application protocol, see "Configuring an Application Protocol" on page 74.

Configuring an RPC Program Number

You can specify an RPC program number for packet matching. To configure an RPC program number, include the **rpc-program-number** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
rpc-program-number number;
```

The range of values used for DCE or RPC is from 100,000 through 400,000. The **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level must have the value **rpc**. For information about specifying the application protocol, see "Configuring an Application Protocol" on page 74.

Configuring the TTL Threshold

You can specify a trace route time-to-live (TTL) threshold value, which controls the acceptable level of network penetration for trace routing. To configure a TTL value, include the **ttl-threshold** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
ttl-threshold value;
```

The **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level must have the value **traceroute**. For information about specifying the application protocol, see "Configuring an Application Protocol" on page 74.

Configuring a Universal Unique Identifier

You can specify a Universal Unique Identifier (UUID) for DCE RPC objects. To configure a UUID value, include the **uuid** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
```

 Copyright © 2013, Juniper Networks, Inc.

uuid *hex-value*;

The **uuid** value is in hexadecimal notation. The **application-protocol** statement at the **[edit applications application *application-name*** hierarchy level must have the value **dce-rpc**. For information about specifying the application protocol, see "Configuring an Application Protocol" on page 74. For more information on UUID numbers, see **http://www.opengroup.org/onlinepubs/9629399/apdxa.htm**.

Configuring Application Sets

You can group the applications you have defined into a named object by including the **application-set** statement at the **[edit applications]** hierarchy level with an **application** statement for each application:

```
[edit applications]
  application-set application-set-name {
  application application;
}
```

For an example of a typical application set, see "Examples: Configuring Application Protocols" on page 118.

ALG Descriptions

This topic describes the Application Layer Gateways (ALGs) supported by the Junos OS. ALG support includes managing pinholes and parent-child relationships for the supported all ALGs. This topic includes the following section:

- Supported ALGs on page 85
- ALG Support Details on page 86
- Juniper Networks Defaults on page 95

Supported ALGs

Table 9 on page 85 lists ALGs supported by the Junos OS.

Table 9: ALGs Supported by the Junos OS

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
Basic TCP ALG	Yes	Yes	Yes	Yes
Basic UPD ALG	Yes	Yes	Yes	Yes
BOOTP	Yes	No	No	No
DCE RPC Services	Yes	No	No	No
DNS	Yes	Yes	No	No
FTP	Yes	No	No	Yes

Copyright © 2013, Juniper Networks, Inc.

Table 9: ALGs Supported by the Junos OS *(continued)*

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
H323	Yes	No	No	No
ICMP	Yes	Yes	Yes	Yes
IIOP	Yes	No	No	No
IP	Yes	No	No	No
NETBIOS	Yes	No	No	No
NETSHOW	Yes	No	No	No
PPTP	Yes	No	No	Yes
REALAUDIO	Yes	No	No	No
Sun RPC and RPC Port Map Services	Yes	No	No	No
RTSP	Yes	No	No	Yes
SIP	Yes	No	No	No
SNMP	Yes	No	No	No
SQLNET	Yes	No	No	No
TFTP	Yes	No	No	Yes
Traceroute	Yes	Yes	No	Yes
Unix Remote Shell Service	Yes	No	No	No
WINFrame	Yes	No	No	No

ALG Support Details

This section includes details about the ALGs. It includes the following:

- Basic TCP ALG on page 87
- Basic UDP ALG on page 88
- BOOTP on page 88
- DCE RPC Services on page 88
- DNS on page 88
- FTP on page 89

Copyright © 2013, Juniper Networks, Inc.

- H323 on page 89
- ICMP on page 89
- IIOP on page 90
- IP on page 90
- NetBIOS on page 90
- NetShow on page 90
- ONC RPC Services on page 90
- PPTP on page 91
- RealAudio on page 91
- Sun RPC and RPC Portmap Services on page 92
- RTSP on page 93
- SIP on page 93
- SNMP on page 94
- SQLNet on page 94
- TFTP on page 94
- Traceroute on page 94
- UNIX Remote-Shell Services on page 95
- Winframe on page 95

Basic TCP ALG

This ALG performs basic sanity checking on TCP packets. If it finds errors, it generates the following anomaly events and system log messages:

- TCP source or destination port zero
- TCP header length check failed
- TCP sequence number zero and no flags are set
- TCP sequence number zero and FIN/PSH/RST flags are set
- TCP FIN/RST or SYN(URG|FIN|RST) flags are set

The TCP ALG performs the following steps:

1. When the router receives a SYN packet, the ALG creates TCP forward and reverse flows and groups them in a *conversation*. It tracks the TCP three-way handshake.
2. The SYN-defense mechanism tracks the TCP connection establishment state. It expects the TCP session to be established within a small time interval (currently 4 seconds). If the TCP three-way handshake is not established in that period, the session is terminated.
3. A keepalive mechanism detects TCP sessions with nonresponsive endpoints.
4. ICMP errors are allowed only if there is a flow that matches the selector information specified in the ICMP data.

Copyright © 2013, Juniper Networks, Inc.

Basic UDP ALG

This ALG performs basic sanity checking on UDP headers. If it finds errors,. it generates the following anomaly events and system log messages:

- UDP source or destination port 0
- UDP header length check failed

The UDP ALG performs the following steps:

1. When it receives the first packet, the ALG creates bidirectional flows to accept forward and reverse UDP session traffic.
2. If the session is idle for more than the maximum allowed idle time (the default is 30 seconds), the flows are deleted.
3. ICMP errors are allowed only if there is a flow that matches the selector information specified in the ICMP data.

BOOTP

The Bootstrap Protocol (BOOTP) client retrieves its networking information from a server across the network. It sends out a general broadcast message to request the information, which is returned by the BOOTP server. For the protocol specification, see **ftp://ftp.isi.edu/in-notes/rfc951.txt**.

Stateful firewall support requires that you configure the BOOTP ALG on UDP server port 67 and client port 68. If the client sends a broadcast message, you should configure the broadcast address in the **from** statement of the service rule. Network Address Translation (NAT) is not performed on the BOOTP traffic, even if the NAT rule matches the traffic. If the BOOTP relay feature is activated on the router, the remote BOOTP server is assumed to assign addresses for clients masked by NAT translation.

DCE RPC Services

Distributed Computing Environment (DCE) Remote Procedure Call (RPC) services are mainly used by Microsoft applications. The ALG uses well-known TCP port 135 for port mapping services, and uses the universal unique identifier (UUID) instead of the program number to identify protocols. The main application-based DCE RPC is the Microsoft Exchange Protocol.

Support for stateful firewall and NAT services requires that you configure the DCE RPC portmap ALG on TCP port 135. The DCE RPC ALG uses the TCP protocol with application-specific UUIDs.

DNS

The Domain Name Service (DNS) ALG handles data associated with locating and translating domain names into IP addresses. The ALG typically runs on port 53. The ALG monitors DNS query and reply packets and supports only UDP traffic. The ALG does not support payload translations. The DNS ALG will only close the session when a reply is received or an idle timeout is reached.

Copyright © 2013, Juniper Networks, Inc.

FTP

FTP is the File Transfer Protocol, specified in RFC 959. In addition to the main control connection, data connections are also made for any data transfer between the client and the server; and the host, port, and direction are negotiated through the control channel.

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

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

Support for stateful firewall and NAT services requires that you configure the FTP ALG on TCP port 21 to enable the FTP control protocol. The ALG performs the following tasks:

- Automatically allocates data ports and firewall permissions for dynamic data connection
- Creates flows for the dynamically negotiated data connection
- Monitors the control connection in both active and passive modes
- Rewrites the control packets with the appropriate NAT address and port information

H323

H323 is a suite of ITU protocols for audio and video conferencing and collaboration applications. H323 consists of H.225 call signaling protocols and H.245 control protocol for media communication. During H.225 negotiation, the endpoints create a call by exchanging call signaling messages on the control channel and negotiate a new control channel for H.245. A new control connection is created for H.245 messages. Messages are exchanged on the H.245 control channel to open media channels.

Stateful firewall monitors the H.225 control channel to open the H.245 control channel. After the H.245 channel is created, stateful firewall also monitors this channel for media channel information and allows the media traffic throught the firewall.

H323 ALG supports static destination, static and dynamic source NAT by rewriting the appropriate addresses and ports in the H.225 and H.245 messages.

ICMP

The Internet Control Message Protocol (ICMP) is defined in RFC 792. The Junos OS stateful firewall service allows ICMP messages to be filtered by specific type or specific type code value. ICMP error packets that lack a specifically configured type and code are matched against any existing flow in the opposite direction to check for the legitimacy

Copyright © 2013, Juniper Networks, Inc.

of the error packet. ICMP error packets that pass the filter matching are subject to NAT translation.

The ICMP ALG always tracks ping traffic statefully using the ICMP sequence number. Each echo reply is forwarded only if there is an echo request with the corresponding sequence number. For any ping flow, only 20 echo requests can be forwarded without receiving an echo reply. When you configure dynamic NAT, the PING packet identifier is translated to allow additional hosts in the NAT pool to use the same identifier.

Support for stateful firewall and NAT services requires that you configure the ICMP ALG if the protocol is needed. You can configure the ICMP type and code for additional filtering.

IIOP

The Oracle Application Server NameServer Internet Inter-ORB Protocol (IIOP). This ALG is used in Common Object Request Broker Architecture (CORBA) based on distributed computing. Even though CORBA and IIOP are Object Management Group (OMG) standards, there is no fixed port assigned for IIOP. Each vendor implementing CORBA chooses a port. Java Virtual machine uses port 1975 by default, while ORBIX uses port 3075 as a default.

Stateful firewall and NAT require ALG IIOP be configured for TCP port 1975 for Java VM IIOP, and 3075 for CORBA applications ORBIX, a CORBA framework from Iona Technologies.

IP

NetBIOS

A NetBIOS ALG translates NetBIOS IP addresses and port numbers when NAT is used.

NetBIOS supports the TCP and UDP transport protocols. Support for stateful firewall and NAT services requires that you configure the NetBIOS ALG on UDP port 138 and TCP port 139.

NetShow

The Microsoft protocol ms-streaming is used by NetShow, the Microsoft media server. This protocol supports several transport protocols: TCP, UDP, and HTTP. The client starts a TCP connection on port 1755 and sends the PORT command to the server. The server then starts UDP on that port to the client. Support for stateful firewall and NAT services requires that you configure the NetShow ALG on UDP port 1755.

ONC RPC Services

Open Networks Computing (ONC) RPC services function similarly to DCE RCP services. However, the ONC RPC ALG uses TCP/UDP port 111 for port mapping services, and uses the program number to identify protocols rather than the UUID.

Support for stateful firewall and NAT services requires that you configure the ONC RPC portmap ALG on TCP port 111. The ONC RPC ALG uses the TCP protocol with application-specific program numbers.

Copyright © 2013, Juniper Networks, Inc.

PPTP

The Point-to-Point Tunneling Protocol (PPTP) ALG is a TCP-based ALG. PPTP allows the Point-to-Point Protocol (PPP) to be tunneled through an IP network. PPTP defines a client-server architecture, a PPTP Network Server, and a PPTP Access Concentrator. The PPTP ALG requires a control connection and a data tunnel. The control connection uses TCP to establish and disconnect PPP sessions, and runs on port 1723. The data tunnel carries PPP traffic in generic routing encapsulated (GRE) packets that are carried over IP.

RealAudio

Real Networks PNA protocol RealVideo is not a separate service. It is part of the RealPlayer and most likely uses another channel for video. The RealPlayer versions G2, 7, and 8 use PNA and RTSP. For this version to work, the ALG must allow both PNA(7070) and RTSP(554). For the media, the server selects from a range of UDP ports(6970 through 7170), or TCP port 7071, or HTTP. The client can be configured to use a particular port. The RealPlayer versions 4.0 and 5.0 use control channel 7070 media UDP ports 6970 through 7170, or TCP port 7071, or HTTP. RealAudio player version 3.0 uses control channel 7070 media, UDP ports 6770-7170, or TCP port 7071.

Real products use the ports and ranges of ports shown in Table 10 on page 91.

Table 10: RealAudio Product Port Usage

Real Product	Port Usage
4.0 and 5.0 Servers/Players	Control channel (bidirectional) on TCP port 7070. Data channel from server to player on TCP port 7070 or UDP port 6970-7170.
4.0 and 5.0 Servers/Encoders	Control channel (bidirectional) on TCP port 7070. Data channel from encoder or server on TCP port 7070.
G2 Servers/Players	Control channel (bidirectional) on TCP port 80, 554, 7070, or 8080. Data channel from server to player on TCP port 80, 554, 7070, 8080 or UDP port 6970-32,000.
G2 Server/3.1, and 5.x Encoders	Control channel (bidirectional) on TCP port 7070. Data channel from encoder to server on TCP port 7070.
G2 Server/G2 Producer	Control channel (bidirectional) on TCP port 4040. Data channel from encoder to server on TCP port 4040 and UDP port 6970-32,000.
2 Server/G2 Producer (TCP ONLY)	Control channel (bidirectional) on TCP port 4040 Data channel from encoder to server on TCP port 4040. Note: TCP-ONLY option available in version 6.1 or above.

NOTE: RealAudio was the original protocol by RealPlayers. Newer versions of RealPlayer use RTSP. Stateful firewall and NAT require ALG RealAudio to be programmed on TCP port 7070.

Copyright © 2013, Juniper Networks, Inc.

Sun RPC and RPC Portmap Services

The Remote Procedure Call (RPC) ALG uses well-known ports TCP 111 and UDP 111 for port mapping, which dynamically assigns and opens ports for RPC services. The RPC Portmap ALG keeps track of port requests and dynamically opens the firewall for these requested ports. The RPC ALG can further restrict the RPC protocol by specifying allowed program numbers.

The ALG includes the RPC services listed in Table 11 on page 92.

Table 11: Supported RPC Services

Name	Description	Comments
rpc-mountd	Network File Server (NFS) mount daemon; for details, see the UNIX man page for **rpc.mountd(8)**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
rpc-nfsprog	Used as part of NFS. For details, see RFC 1094. See also RFC1813 for NFS v3.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
rpc-nisplus	Network Information Service Plus (NIS+), designed to replace NIS; it is a default naming service for Sun Solaris and is not related to the old NIS. No protocol information is available.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
rpc-nlockmgr	Network lock manager.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-nlockmgr** service can be allowed or blocked based on RPC program 100021.
rpc-pcnfsd	Kernel statistics server. For details, see the UNIX man pages for **rstatd** and **rpc.rstatd**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-rstat** service can be allowed or blocked based on RPC program 150001.
rpc-rwall	Used to write a message to users; for details, see the UNIX man page for **rpc.rwalld**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-rwall** service can be allowed or blocked based on RPC program 150008.
rpc-ypbind	NIS binding process. For details, see the UNIX man page for **ypbind**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-ypbind** service can be allowed or blocked based on RPC program 100007.
rpc-yppasswd	NIS password server. For details, see the UNIX man page for **yppasswd**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-yppasswd** service can be allowed or blocked based on RPC program 100009.
rpc-ypserv	NIS server. For details, see the UNIX man page for **ypserv**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-ypserv** service can be allowed or blocked based on RPC program 100004.

Copyright © 2013, Juniper Networks, Inc.

Table 11: Supported RPC Services *(continued)*

Name	Description	Comments
rpc-ypupdated	Network updating tool.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-ypupdated** service can be allowed or blocked based on RPC program 100028.
rpc-ypxfrd	NIS map transfer server. For details, see the UNIX man page for **rpc.ypxfrd**.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, **rpc-ypxfrd** service can be allowed or blocked based on RPC program 100069.

Support for stateful firewall and NAT services that use port mapping requires that you configure the RPC portmap ALG on TCP/UDP destination port 111 and the RPC ALG for both TCP and UDP. You can specify one or more **rpc-program-number** values to further restrict allowed RPC protocols.

RTSP

The Real-Time Streaming Protocol (RTSP) controls the delivery of data with real-time properties such as audio and video. The streams controlled by RTSP can use RTP, but it is not required. Media can be transmitted on the same RTSP control stream. This is an HTTP-like text-based protocol, but client and server maintain session information. A session is established using the SETUP message and terminated using the TEARDOWN message. The transport (the media protocol, address, and port numbers) is negotiated in the setup and the setup-response.

Support for stateful firewall and NAT services requires that you configure the RTSP ALG for TCP port 554.

The ALG monitors the control connection, opens flows dynamically for media (RTP/RTSP) streams, and performs NAT address and port rewrites.

SIP

The Session Initiation Protocol (SIP) is an application layer protocol that can establish, maintain, and terminate media sessions. It is a widely used voice over IP (VoIP) signaling protocol. The SIP ALG monitors SIP traffic and dynamically creates and manages pinholes on the signaling and media paths. The ALG only allows packets with the correct permissions. The SIP ALG also performs the following functions:

- Manages parent-child session relationships.
- Enforces security policies.
- Manages pinholes for VoIP traffic.

Starting with Junos OS Release 11.4, the SIP ALG supports Network Address Translation (NAT) and stateful firewall configuration on JSF. The SIP ALG supports the following features:

Copyright © 2013, Juniper Networks, Inc.

- Stateful firewall
- Static source NAT
- Dynamic address only source NAT
- Network Address Port Translation (NAPT)



NOTE: The SIP ALG does not support destination NAT, class of service (CoS), multicast, or IP version 6 (IPv6).

At present, the SIP ALG does not support the following features:

- Encryption and authentication of SIP messages
- Transport of SIP messages over TCP

SNMP

SNMP is a communication protocol for managing TCP/IP networks, including both individual network devices and aggregated devices. The protocol is defined by RFC 1157. SNMP runs on top of UDP.

The Junos OS stateful firewall service implements the SNMP ALG to inspect the SNMP type. SNMP does not enforce stateful flow. Each SNMP type needs to be specifically enabled. Full SNMP support of stateful firewall services requires that you configure the SNMP ALG on UDP port 161. This enables the SNMP ***get*** and **get-next** commands, as well as their response traffic in the reverse direction: UDP port 161 enables the SNMP **get-response** command. If SNMP traps are permitted, you can configure them on UDP port 162, enabling the SNMP **trap** command.

SQLNet

The SQLNet protocol is used by Oracle SQL servers to execute SQL commands from clients, including load balancing and application-specific services.

Support of stateful firewall and NAT services requires that you configure the SQLNet ALG for TCP port 1521.

The ALG monitors the control packets, opens flows dynamically for data traffic, and performs NAT address and port rewrites.

TFTP

The Trivial File Transfer Protocol (TFTP) is specified in RFC 1350. The initial TFTP requests are sent to UDP destination port 69. Additional flows can be created to **get** or **put** individual files. Support of stateful firewall and NAT services requires that you configure the TFTP ALG for UDP destination port 69.

Traceroute

Traceroute is a tool for displaying the route that packets take to a network host. It uses the IP time-to-live (TTL) field to trigger ICMP time-exceeded messages from routers or

Copyright © 2013, Juniper Networks, Inc.

gateways. It sends UDP datagrams to destination ports that are believed to be not in use; destination ports are numbered using the formula: + *n*hops – 1. The default base port is 33434. To support traceroute through the firewall, two types of traffic must be passed through:

1. UDP probe packets (UDP destination port > 33000, IP TTL < 30)
2. ICMP response packets (ICMP type time-exceeded)

When NAT is applied, the IP address and port within the ICMP error packet also must be changed.

Support of stateful firewall and NAT services requires you to configure the Traceroute ALG for UDP destination port 33434 to 33450. In addition, you can configure the TTL threshold to prevent UDP flood attacks with large TTL values.

UNIX Remote-Shell Services

Three protocols form the basis for UNIX remote-shell services:

- Exec—Remote command execution; enables a user on the client system to execute a command on the remote system. The first command from client (**rcmd**) to server (**rshd**) uses well-known TCP port 512. A second TCP connection can be opened at the request of **rcmd**. The client port number for the second connection is sent to the server as an ASCII string.
- Login—Better known as **rlogin**; uses well-known TCP port 513. For details, see RFC 1282. No special firewall processing is required.
- Shell—Remote command execution; enables a user on the client system to execute a command on the remote system. The first command from client (**rcmd**) to server (**rshd**) uses well-known TCP port 514. A second TCP connection can be opened at the request of **rcmd**. The client port number for the second connection is sent to the server as an ASCII string.

Support of stateful firewall services requires that you configure the Exec ALG on TCP port 512, the Login ALG on TCP port 513, and the Shell ALG on TCP port 514. NAT remote-shell services require that any dynamic source port assigned be within the port range 512 to 1023. If you configure a NAT pool, this port range is reserved exclusively for remote shell applications.

Winframe

WinFrame application server software provides access to virtually any Windows application, across any type of network connection to any type of client.

This protocol is mainly used by Citrix Windows applications.

Stateful firewall and NAT require the ALG Winframe to be configured on TCP destination port 1494 and UDP port 1604.

Juniper Networks Defaults

```
applications {
```

Copyright © 2013, Juniper Networks, Inc.

```
#
# File Transfer Protocol
#
application junos-ftp {
   application-protocol ftp;
   protocol tcp;
   destination-port 21;
}
#
# Trivial File Transfer Protocol
#
application junos-tftp {
   application-protocol tftp;
   protocol udp;
   destination-port 69;
}
#
# RPC portmapper on TCP
#
application junos-rpc-portmap-tcp {
   application-protocol rpc-portmap;
   protocol tcp;
   destination-port 111;
}
#
# RPC portmapper on UDP
#
application junos-rpc-portmap-udp {
   application-protocol rpc-portmap;
   protocol udp;
   destination-port 111;
}
#
# SNMP get
#
application junos-snmp-get {
   application-protocol snmp;
   protocol udp;
   destination-port 161;
   snmp-command get;
}
#
# SNMP get next
#
application junos-snmp-get-next {
   application-protocol snmp;
   protocol udp;
   destination-port 161;
   snmp-command get-next;
}
#
# SNMP response
#
application junos-snmp-response {
   application-protocol snmp;
   protocol udp;
```


Copyright © 2013, Juniper Networks, Inc.

```
    source-port 161;
    snmp-command get-response;
}
#
# SNMP trap
#
application junos-snmp-trap {
    application-protocol snmp;
    protocol udp;
    destination-port 162;
    snmp-command trap;
}
#
# remote exec
#
application junos-rexec {
    application-protocol exec;
    protocol tcp;
    destination-port 512;
}
#
# remote login
#
application junos-rlogin {
    application-protocol shell;
    protocol tcp;
    destination-port 513;
}
#
# remote shell
#
application junos-rsh {
    application-protocol shell;
    protocol tcp;
    destination-port 514;
}
#
# Real Time Streaming Protocol
#
application junos-rtsp {
    application-protocol rtsp;
    protocol tcp;
    destination-port 554;
}
#
# Citrix windows application server protocol
# windows applications remotely on windows/non-windows clients
#
# citrix needs udp 1604 to be open
#
application junos-citrix-winframe {
    application-protocol winframe;
    protocol tcp;
    destination-port 1494;
}
application junos-citrix-winframe-udp {
```

Copyright © 2013, Juniper Networks, Inc.

```
    protocol udp;
    destination-port 1604;
}
#
# Oracle SQL servers use this protocol to execute sql commands
# from clients, load balance, use application-specific servers, etc
#
application junos-sqlnet {
    application-protocol sqlnet;
    protocol tcp;
    destination-port 1521;
}
#
# H.323 Protocol for audio/video conferencing
#
application junos-h323 {
    application-protocol h323;
    protocol tcp;
    destination-port 1720;
}
#
# Internet Inter-ORB Protocol - used for CORBA applications
# The ORB protocol in Java virtual machines uses port 1975 as default
#
application junos-iiop-java {
    application-protocol iiop;
    protocol tcp;
    destination-port 1975;
}
#
# Internet Inter-ORB Protocol - used for CORBA applications
# ORBIX is a CORBA framework from Iona Technologies that uses port
# 3075 as default
#
application junos-iiop-orbix {
    application-protocol iiop;
    protocol tcp;
    destination-port 3075;
}
#
# Real players use this protocol for real time streaming
# This was the original protocol for real players.
# RTSP is more widely used by real players
# but they still support realaudio.
#
application junos-realaudio {
    application-protocol realaudio;
    protocol tcp;
    destination-port 7070;
}
#
# traceroute application.
#
application junos-traceroute {
    application-protocol traceroute;
    protocol udp;
```


Copyright © 2013, Juniper Networks, Inc.

```
    destination-port 33435-33450;
    ttl-threshold 30;
}
#
# The full range of known RPC programs using UDP
# The program numbers can be more specific to certain applications.
#
application junos-rpc-services-udp {
    application-protocol rpc;
    protocol udp;
    rpc-program-number 100000-400000;
}
#
# The full range of known RPC programs using TCP
# The program numbers can be more specific to certain applications.
#
application junos-rpc-services-tcp {
    application-protocol rpc;
    protocol tcp;
    rpc-program-number 100000-400000;
}
#
# All ICMP traffic
# This can be made to be more restrictive by specifying ICMP type
# and code.
#
application junos-icmp-all {
    application-protocol icmp;
}
#
# Protocol used by Windows media server and windows media player
#
application junos-netshow {
    application-protocol netshow;
    protocol tcp;
    destination-port 1755;
}
#
# NetBIOS - networking protocol used on
# Windows networks name service port, both UDP and TCP
#
application junos-netbios-name-udp {
    application-protocol netbios;
    protocol udp;
    destination-port 137;
}
application junos-netbios-name-tcp {
    protocol tcp;
    destination-port 137;
}
#
# NetBIOS - networking protocol used on
# Windows networks datagram service port
#
application junos-netbios-datagram {
    application-protocol netbios;
```

Copyright © 2013, Juniper Networks, Inc.

```
    protocol udp;
    destination-port 138;
  }
  #
  # NetBIOS - networking protocol used on
  # Windows networks session service port
  #
  application junos-netbios-session {
    protocol tcp;
    destination-port 139;
  }
  #
  # DCE-RPC portmapper on TCP
  #
  application junos-dce-rpc-portmap {
    application-protocol dce-rpc-portmap;
    protocol tcp;
    destination-port 135;
  }
  #
  # DCE-RPC application on TCP sample UUID
  # This application requires user to specify the UUID value
#
  # application junos-dcerpc {
    # application-protocol dce-rpc;
    # protocol tcp;
    #
    # # UUID also needs to be defined as shown below
    # UUID 11223344 22334455 33445566 44556677;
    #
    # }
    #
    # ms-exchange needs these 3 UUIDs
    #
    application junos-dcerpc-endpoint-mapper-service {
      application-protocol dce-rpc;
      protocol tcp;
      uuid e1af8308-5d1f-11c9-91a4-08002b14a0fa;
    }
    application junos-dcerpc-msexchange-directory-rfr {
      application-protocol dce-rpc;
      protocol tcp;
      uuid 1544f5e0-613c-11d1-93df-00c04fd7bd09;
    }
    application junos-dcerpc-msexchange-information-store {
      application-protocol dce-rpc;
      protocol tcp;
      uuid a4f1db00-ca47-1067-b31f-00dd010662da;
    }
    application junos-ssh {
      protocol tcp;
      destination-port 22;
    }
    application junos-telnet {
      protocol tcp;
      destination-port 23;
```


Copyright © 2013, Juniper Networks, Inc.

```
}
application junos-smtp {
   protocol tcp;
   destination-port 25;
}
application junos-dns-udp {
   protocol udp;
   destination-port 53;
}
application junos-dns-tcp {
   protocol tcp;
   destination-port 53;
}
application junos-tacacs {
   protocol tcp;
   destination-port 49;
}
# TACACS Database Service
application junos-tacacs-ds {
   protocol tcp;
   destination-port 65;
}
application junos-dhcp-client {
   protocol udp;
   destination-port 68;
}
application junos-dhcp-server {
   protocol udp;
   destination-port 67;
}
application junos-bootpc {
   protocol udp;
   destination-port 68;
}
application junos-bootps {
   protocol udp;
   destination-port 67;
}
application junos-finger {
   protocol tcp;
   destination-port 79;
}
application junos-http {
   protocol tcp;
   destination-port 80;
}
application junos-https {
   protocol tcp;
   destination-port 443;
}
application junos-pop3 {
   protocol tcp;
   destination-port 110;
}
application junos-ident {
   protocol tcp;
```

Copyright © 2013, Juniper Networks, Inc.

```
    destination-port 113;
}
application junos-nntp {
    protocol tcp;
    destination-port 119;
}
application junos-ntp {
    protocol udp;
    destination-port 123;
}
application junos-imap {
    protocol tcp;
    destination-port 143;
}
application junos-imaps {
    protocol tcp;
    destination-port 993;
}
application junos-bgp {
    protocol tcp;
    destination-port 179;
}
application junos-ldap {
    protocol tcp;
    destination-port 389;
}
application junos-snpp {
    protocol tcp;
    destination-port 444;
}
application junos-biff {
    protocol udp;
    destination-port 512;
}
# UNIX who
application junos-who {
    protocol udp;
    destination-port 513;
}
application junos-syslog {
    protocol udp;
    destination-port 514;
}
# line printer daemon, printer, spooler
application junos-printer {
    protocol tcp;
    destination-port 515;
}
# UNIX talk
application junos-talk-tcp {
    protocol tcp;
    destination-port 517;
}
application junos-talk-udp {
    protocol udp;
    destination-port 517;
```


Copyright © 2013, Juniper Networks, Inc.

```
}
application junos-ntalk {
    protocol udp;
    destination-port 518;
}
application junos-rip {
    protocol udp;
    destination-port 520;
}
# INA sanctioned RADIUS port numbers
application junos-radius {
    protocol udp;
    destination-port 1812;
}
application junos-radacct {
    protocol udp;
    destination-port 1813;
}
application junos-nfsd-tcp {
    protocol tcp;
    destination-port 2049;
}
application junos-nfsd-udp {
    protocol udp;
    destination-port 2049;
}
application junos-cvspserver {
    protocol tcp;
    destination-port 2401;
}
#
# Label Distribution Protocol
#
application junos-ldp-tcp {
    protocol tcp;
    destination-port 646;
}
application junos-ldp-udp {
    protocol udp;
    destination-port 646;
}
#
# JUNOScript and JUNOScope management
#
application junos-xnm-ssl {
    protocol tcp;
    destination-port 3220;
}
application junos-xnm-clear-text {
    protocol tcp;
    destination-port 3221;
}
#
# IPsec tunnel
#
application junos-ipsec-esp {
```

Copyright © 2013, Juniper Networks, Inc.

```
    protocol esp;
}
application junos-ike {
    protocol udp;
    destination-port 500;
}
#
# 'junos-algs-outbound' defines a set of all applications
# requiring an ALG. Useful for defining rule to the the public
# internet allowing private network users to use all JUNOS OS
# supported ALGs initiated from the private network.
#
# NOTE: the contents of this set might grow in future JUNOS OS versions.
#
application-set junos-algs-outbound {
    application junos-ftp;
    application junos-tftp;
    application junos-rpc-portmap-tcp;
    application junos-rpc-portmap-udp;
    application junos-snmp-get;
    application junos-snmp-get-next;
    application junos-snmp-response;
    application junos-snmp-trap;
    application junos-rexec;
    application junos-rlogin;
    application junos-rsh;
    application junos-rtsp;
    application junos-citrix-winframe;
    application junos-citrix-winframe-udp;
    application junos-sqlnet;
    application junos-h323;
    application junos-iiop-java;
    application junos-iiop-orbix;
    application junos-realaudio;
    application junos-traceroute;
    application junos-rpc-services-udp;
    application junos-rpc-services-tcp;
    application junos-icmp-all;
    application junos-netshow;
    application junos-netbios-name-udp;
    application junos-netbios-datagram;
    application junos-dcerpc-endpoint-mapper-service;
    application junos-dcerpc-msexchange-directory-rfr;
    application junos-dcerpc-msexchange-information-store;
}
#
# 'junos-management-inbound' represents the group of applications
# that might need access the router from public network for
# for management purposes.
#
# Set is intended for a UI to display management choices.
#
# NOTE: It is not recommended the user to use the entire set
# directly in a firewall rule and open up firewall to all
# of these applications. Also, the user should always
# specify the source and destination prefixes when using
```


Copyright © 2013, Juniper Networks, Inc.

```
    # each application.
    #
    # NOTE: the contents of this set may grow in future JUNOS versions.
    #
    application-set junos-management-inbound {
      application junos-snmp-get;
      application junos-snmp-get-next;
      application junos-snmp-response;
      application junos-snmp-trap;
      application junos-ssh;
      application junos-telnet;
      application junos-http;
      application junos-https;
      application junos-xnm-ssl;
      application junos-xnm-clear-text;
    }
    #
    # 'junos-routing-inbound' represents routing protocols that might
    # need to access the router from public network.
    #
    # Set is intended for a UI to display routing involvement choices.
    #
    # NOTE: It is not recommended the user to use the entire set
    # directly in a firewall rule and open up firewall to all
    # of these applications. Also, the user should always
    # specify the source and destination prefixes when using
    # each application.
    #
    # NOTE: the contents of this set might grow in future JUNOS OS versions.
    #
    application-set junos-routing-inbound {
      application junos-bgp;
      application junos-rip;
      application junos-ldp-tcp;
      application junos-ldp-udp;
    }
  }
```

Verifying the Output of ALG Sessions

This section contains examples of successful output from ALG sessions and information on system log configuration. You can compare the results of your sessions to check whether the configurations are functioning correctly.

- FTP Example on page 106
- RTSP ALG Example on page 108
- System Log Messages on page 110

Copyright © 2013, Juniper Networks, Inc.

FTP Example

This example analyzes the output during an active FTP session. It consists of four different flows; two are control flows and two are data flows. The example consists of the following parts:

- Sample Output on page 106
- FTP System Log Messages on page 107
- Analysis on page 107
- Troubleshooting Questions on page 108

Sample Output

The following is a complete sample output from the **show services stateful-firewall conversations application-protocol ftp** operational mode command:

```
user@host>show services stateful-firewall conversations application-protocol ftp
Interface: ms-1/3/0, Service set: CLBJI1-AAF001
Conversation: ALG protocol: ftp
  Number of initiators: 2, Number of responders: 2
Flow                                              State    Dir       Frm count
TCP           1.1.79.2:14083 ->        2.2.2.2:21    Watch    I              13
    NAT source        1.1.79.2:14083   ->   194.250.1.237:50118
TCP           1.1.79.2:14104 ->        2.2.2.2:20    Forward  I               3
    NAT source        1.1.79.2:14104   ->   194.250.1.237:50119
TCP            2.2.2.2:21    ->  194.250.1.237:50118 Watch    O              12
    NAT dest     194.250.1.237:50118   ->        1.1.79.2:14083
TCP            2.2.2.2:20    ->  194.250.1.237:50119 Forward  O               5
    NAT dest     194.250.1.237:50119   ->        1.1.79.2:14104
```

For each flow, the first line shows flow information, including protocol (TCP), source address, source port, destination address, destination port, flow state, direction, and frame count.

- The state of a flow can be **Watch**, **Forward**, or **Drop**:
  - A **Watch** flow state indicates that the control flow is monitored by the ALG for information in the payload. NAT processing is performed on the header and payload as needed.
  - A **Forward** flow forwards the packets without monitoring the payload. NAT is performed on the header as needed.
  - A **Drop** flow drops any packet that matches the 5 tuple.
- The frame count (**Frm count**) shows the number of packets that were processed on that flow.

The second line shows the NAT information.

- **source** indicates source NAT.
- **dest** indicates destination NAT.

Copyright © 2013, Juniper Networks, Inc.

- The first address and port in the NAT line are the original address and port being translated for that flow.
- The second address and port in the NAT line are the translated address and port for that flow.

FTP System Log Messages

System log messages are generated during an FTP session. For more information about system logs, see "System Log Messages" on page 110.

The following system log messages are generated during creation of the FTP control flow:

- Rule Accept system log:

  Oct 27 11:42:54 (FPC Slot 1, PIC Slot 1) {ss_ftp}[FWNAT]: ASP_SFW_RULE_ACCEPT: proto 6 (TCP) application: ftp, fe-3/3/3.0:1.1.1.2:4450 -> 2.2.2.2:21, Match SFW accept rule-set:, rule: ftp, term: 1

- Create Accept Flow system log:

  Oct 27 11:42:54 (FPC Slot 1, PIC Slot 1) {ss_ftp}[FWNAT]: ASP_SFW_CREATE_ACCEPT_FLOW: proto 6 (TCP) application: ftp, fe-3/3/3.0:1.1.1.2:4450 -> 2.2.2.2:21, creating forward or watch flow

- System log for data flow creation:

  Oct 27 11:43:30 (FPC Slot 1, PIC Slot 1) {ss_ftp}[FWNAT]: ASP_SFW_FTP_ACTIVE_ACCEPT: proto 6 (TCP) application: ftp, so-2/1/2.0:2.2.2.2:20 -> 1.1.1.2:50726, Creating FTP active mode forward flow

Analysis

Control Flows

The control flows are established after the three-way handshake is complete.

- Control flow from FTP client to FTP server. TCP destination port is 21.

```
TCP           1.1.79.2:14083 ->        2.2.2.2:21    Watch    I
13
    NAT source        1.1.79.2:14083   ->   194.250.1.237:50118
```

- Control flow from FTP server to FTP client. TCP source port is 21.

```
TCP            2.2.2.2:21    ->  194.250.1.237:50118 Watch    O
12
    NAT dest     194.250.1.237:50118   ->        1.1.79.2:14083
```

Data Flows

A data port of 20 is negotiated for data transfer during the course of the FTP control protocol. These two flows are data flows between the FTP client and the FTP server:

```
TCP           1.1.79.2:14104 ->        2.2.2.2:20    Forward  I              3
    NAT source        1.1.79.2:14104   ->   194.250.1.237:50119
```

Copyright © 2013, Juniper Networks, Inc.

```
TCP            2.2.2.2:20     ->  194.250.1.237:50119 Forward  0              5
    NAT dest     194.250.1.237:50119    ->        1.1.79.2:14104
```

Troubleshooting Questions

1. How do I know if the FTP ALG is active?
   - The ALG protocol field in the conversation should display **ftp**.
   - There should be a valid frame count (**Frm count**) in the control flows.
   - A valid frame count in the data flows indicates that data transfer has taken place.
2. What do I need to check if the FTP connection is established but data transfer does not take place?
   - Most probably, the control connection is up, but the data connection is down.
   - Check the conversations output to determine whether both the control and data flows are present.
3. How do I interpret each flow? What does each flow mean?
   - FTP control flow initiator flow—Flow with destination port 21
   - FTP control flow responder flow—Flow with source port ;21
   - FTP data flow initiator flow—Flow with destination port 20
   - FTP data flow responder flow—Flow with source port 20

RTSP ALG Example

The following is an example of an RTSP conversation. The application uses the RTSP protocol for control connection. Once the connection is set up, the media is sent using UDP protocol (RTP).

This example consists of the following:

- Sample Output on page 108
- Analysis on page 109
- Troubleshooting Questions on page 109

Sample Output

Here is the output from the **show services stateful-firewall conversations** operational mode command:

```
user@host# show services stateful-firewall conversations
Interface: ms-3/2/0, Service set: svc_set
Conversation: ALG protocol: rtsp
  Number of initiators: 5, Number of responders: 5
Flow                                                 State    Dir       Frm count
TCP            1.1.1.3:58795 ->        2.2.2.2:554   Watch    I              7
UDP            1.1.1.3:1028  ->        2.2.2.2:1028  Forward  I              0
UDP            1.1.1.3:1029  ->        2.2.2.2:1029  Forward  I              0
UDP            1.1.1.3:1030  ->        2.2.2.2:1030  Forward  I              0
UDP            1.1.1.3:1031  ->        2.2.2.2:1031  Forward  I              0
```


Copyright © 2013, Juniper Networks, Inc.

```
TCP            2.2.2.2:554   ->        1.1.1.3:58795 Watch    0              5
UDP            2.2.2.2:1028  ->        1.1.1.3:1028  Forward  0              6
UDP            2.2.2.2:1029  ->        1.1.1.3:1029  Forward  0              0
UDP            2.2.2.2:1030  ->        1.1.1.3:1030  Forward  0              3
UDP            2.2.2.2:1031  ->        1.1.1.3:1031  Forward  0              0
```

Analysis

An RTSP conversation should consist of TCP flows corresponding to the RTSP control connection. There should be two flows, one in each direction, from client to server and from server to client:

```
TCP            1.1.1.3:58795 ->        2.2.2.2:554   Watch    I              7
TCP            2.2.2.2:554   ->        1.1.1.3:58795 Watch    O              5
```

- The RTSP control connection for the initiator flow is sent from destination port 554.
- The RTSP control connection for the responder flow is sent from source port 554.

The UDP flows correspond to RTP media sent over the RTSP connection.

Troubleshooting Questions

1. Media does not work when the RTSP ALG is configured. What do I do?
   - Check RTSP conversations to see whether both TCP and UDP flows exist.
   - The ALG protocol should be displayed as **rtsp**.

   

   NOTE: The state of the flow is displayed as Watch, because the ALG processing is taking place and the client is essentially “watching” or processing payload corresponding to the application. For FTP and RTSP ALG flows, the control connections are always Watch flows.

2. How do I check for ALG errors?
   - You can check for errors by issuing the following command. Each ALG has a separate field for ALG packet errors.

     ```
     user@host# show services stateful-firewall statistics extensive
     Interface: ms-3/2/0
       Service set: svc_set
         New flows:
           Accepts: 1347, Discards: 0, Rejects: 0
         Existing flows:
           Accepts: 144187, Discards: 0, Rejects: 0
         Drops:
           IP option: 0, TCP SYN defense: 0
           NAT ports exhausted: 0
         Errors:
           IP: 0, TCP: 276
           UDP: 0, ICMP: 0
           Non-IP packets: 0, ALG: 0
         IP errors:
           IP packet length inconsistencies: 0
           Minimum IP header length check failures: 0
           Reassembled packet exceeds maximum IP length: 0
     ```

Copyright © 2013, Juniper Networks, Inc.

```
  Illegal source address: 0
  Illegal destination address: 0
  TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
  Land attack: 0
  Non-IPv4 packets: 0, Bad checksum: 0
  Illegal IP fragment length: 0
  IP fragment overlap: 0
  IP fragment reassembly timeout: 0
  Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combinations: 0
  SYN attack (multiple SYN messages seen for the same flow): 276
  First packet not a SYN message: 0
  TCP port scan (TCP handshake, RST seen from server for SYN): 0
  Bad SYN cookie response: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port number is zero: 0
  UDP port scan (ICMP error seen for UDP flow): 0
ICMP errors:
  IP data length less than minimum ICMP header length (8 bytes): 0
  ICMP error length inconsistencies: 0
  Duplicate ping sequence number: 0
  Mismatched ping sequence number: 0
ALG errors:
  BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
  DNS: 0, Exec: 0, FTP: 0
  ICMP: 0
  Login: 0, NetBIOS: 0, NetShow: 0
  RPC: 0, RPC portmap: 0
  RTSP: 0, Shell: 0
  SNMP: 0, SQLNet: 0, TFTP: 0
  Traceroute: 0
```

System Log Messages

Enabling system log generation and checking the system log are also helpful for ALG flow analysis. This section contains the following:

- System Log Configuration on page 110
- System Log Output on page 111

System Log Configuration

You can configure the enabling of system log messages at a number of different levels in the Junos OS CLI. As shown in the following sample configurations, the choice of level depends on how specific you want the event logging to be and what options you want to include. For details on the configuration options, see the Junos OS System Basics Configuration Guide (system level) or the Junos Services Interfaces Configuration Release 12.3 (all other levels).

1. At the topmost global level:

```
user@host# show system syslog
file messages {
  any any;
```


Copyright © 2013, Juniper Networks, Inc.

```
}
```

2. At the service set level:

```
user@host# show services service-set svc_set
syslog {
  host local {
    services any;
  }
}
stateful-firewall-rules allow_rtsp;
interface-service {
  service-interface ms-3/2/0;
}
```

3. At the service rule level:

```
user@host# show services stateful-firewall rule allow_rtsp
match-direction input-output;
term 0 {
  from {
    applications junos-rtsp;
  }
  then {
    accept;
    syslog;
  }
}
```

System Log Output

System log messages are generated during flow creation, as shown in the following examples:

The following system log message indicates that the ASP matched an accept rule:

```
Oct 25 16:11:37 (FPC Slot 3, PIC Slot 2) {svc_set}[FWNAT]: ASP_SFW_RULE_ACCEPT:
  proto 6 (TCP) application: rtsp, ge-2/0/1.0:1.1.1.2:35595 -> 2.2.2.2:554, Match SFW accept
  rule-set: , rule: allow_rtsp, term: 0
```

For a complete listing of system log messages, see the *Junos OS System Log Messages Reference*.

Junos Default Groups

The Junos OS provides a default, hidden configuration group called **junos-defaults** that is automatically applied to the configuration of your router. The **junos-defaults** group contains preconfigured statements that contain predefined values for common applications. Some of the statements must be referenced to take effect, such as applications like FTP or Telnet. Other statements are applied automatically, such as terminal settings. All of the preconfigured statements begin with the reserved name **junos-**.

Copyright © 2013, Juniper Networks, Inc.



NOTE: You can override the Junos default configuration values, but you cannot delete or edit them. If you delete a configuration, the defaults return when a new configuration is added.

You cannot use the **apply-groups** statement with the Junos defaults group.

To view the full set of available preset statements from the Junos default group, issue the **show groups junos-defaults** configuration mode command. The following example displays a partial list of Junos default groups that use application protocols (ALGs).

```
user@host# show groups junos-defaults
... output for other groups defined at the  [edit groups junos-defaults] hierarchy level ...
applications {
   # File Transfer Protocol
   application junos-ftp {
      application-protocol ftp;
      protocol tcp;
      destination-port 21;
   }
   # Trivial File Transfer Protocol
   application junos-tftp {
      application-protocol tftp;
      protocol udp;
      destination-port 69;
   }
   # RPC port mapper on TCP
   application junos-rpc-portmap-tcp {
      application-protocol rpc-portmap;
      protocol tcp;
      destination-port 111;
   }
   # RPC port mapper on UDP
   application junos-rpc-portmap-udp {
      application-protocol rpc-portmap;
      protocol udp;
      destination-port 111;
   }
   # IP Protocol
   application junos-ip {
      application-protocol ip;
   }
   # remote exec
   application junos-rexec {
      application-protocol exec;
      protocol tcp;
      destination-port 512;
   }
   # remote login
   application junos-rlogin {
      application-protocol login;
      protocol tcp;
      destination-port 513;
   }
```


Copyright © 2013, Juniper Networks, Inc.

```
# remote shell
application junos-rsh {
   application-protocol shell;
   protocol tcp;
   destination-port 514;
}
# Real-Time Streaming Protocol
application junos-rtsp {
   application-protocol rtsp;
   protocol tcp;
   destination-port 554;
}
# Oracle SQL servers use this protocol to execute SQL commands
# from clients, load balance, use application-specific servers, and so on.
application junos-sqlnet {
   application-protocol sqlnet;
   protocol tcp;
   destination-port 1521;
}
# H.323 Protocol for audio/video conferencing
protocol tcp;
   destination-port 1720;
}
# Internet Inter-ORB Protocol is used for CORBA applications.
# The ORB protocol in Java virtual machine uses port 1975 as a default.
protocol tcp;
   destination-port 1975;
}
# Internet Inter-ORB Protocol is used for CORBA applications.
# ORBIX is a CORBA framework from Iona Technologies that uses
# port 3075 as a default.
protocol tcp;
   destination-port 3075;
}
# This was the original RealPlayer protocol.
# RTSP is more widely used by RealPlayer,
protocol tcp;
   destination-port 7070;
}
# Traceroute application
application junos-traceroute {
   application-protocol traceroute;
   protocol udp;
   destination-port 33435-33450;
   ttl-threshold 30;
}
# Traceroute application that stops at device supporting firewall
# (packets with ttl > 1 will be discarded).
application junos-traceroute-ttl-1 {
   application-protocol traceroute;
   protocol udp;
   destination-port 33435-33450;
   ttl-threshold 1;
}
# The full range of known RPC programs using UDP.
# Specific program numbers are assigned to certain applications.
```

Copyright © 2013, Juniper Networks, Inc.

```
application junos-rpc-services-udp {
   application-protocol rpc;
   protocol udp;
   rpc-program-number 100001-400000;
}
# The full range of known RPC programs using TCP.
# Specific program numbers are assigned to certain applications.
application junos-rpc-services-tcp {
   application-protocol rpc;
   protocol tcp;
   rpc-program-number 100001-400000;
}
# All ICMP traffic
# This can be made more restrictive by specifying ICMP type and code.
application junos-icmp-all {
   application-protocol icmp;
}
# ICMP ping; the echo reply is allowed upon return.
application junos-icmp-ping {
   application-protocol icmp;
   icmp-type echo-request;
}
# Protocol used by Windows Media Server and Windows Media Player
application junos-netshow {
   application-protocol netshow;
   protocol tcp;
   destination-port 1755;
}
# NetBIOS, the networking protocol used on Windows networks;
# includes name service port, both UDP and TCP.
application junos-netbios-name-udp {
   application-protocol netbios;
   protocol udp;
   destination-port 137;
}
application junos-netbios-name-tcp {
   protocol tcp;
   destination-port 137;
}
# NetBIOS, the networking protocol used on Windows networks;
# includes datagram service port.
application junos-netbios-datagram {
   application-protocol netbios;
   protocol udp;
   destination-port 138;
}
# NetBIOS, the networking protocol used on Windows networks;
# includes session service port.
application junos-netbios-session {
   protocol tcp;
   destination-port 139;
}
# DCE-RPC port mapper on TCP
application junos-dce-rpc-portmap {
   application-protocol dce-rpc-portmap;
   protocol tcp;
```


Copyright © 2013, Juniper Networks, Inc.

```
    destination-port 135;
}
# MS Exchange requires these three UUID values.
application junos-dcerpc-endpoint-mapper-service {
    application-protocol dce-rpc;
    protocol tcp;
    uuid e1af8308-5d1f-11c9-91a4-08002b14a0fa;
}
application junos-ssh {
    protocol tcp;
    destination-port 22;
}
application junos-telnet {
    protocol tcp;
    destination-port 23;
}
application junos-smtp {
    protocol tcp;
    destination-port 25;
}
application junos-dns-udp {
    protocol udp;
    destination-port 53;
}
application junos-dns-tcp {
    protocol tcp;
    destination-port 53;
}
application junos-tacacs {
    protocol tcp;
    destination-port 49;
}
# TACACS Database Service
application junos-tacacs-ds {
    protocol tcp;
    destination-port 65;
}
application junos-dhcp-client {
    protocol udp;
    destination-port 68;
}
application junos-dhcp-server {
    protocol udp;
    destination-port 67;
}
application junos-bootpc {
    protocol udp;
    destination-port 68;
}
application junos-bootps {
    protocol udp;
    destination-port 67;
}
application junos-http {
    protocol tcp;
    destination-port 80;
```

Copyright © 2013, Juniper Networks, Inc.

```
}
application junos-https {
   protocol tcp;
   destination-port 443;
}
# " junos-algs-outbound" defines a set of all applications
# requiring an ALG. Useful for defining a rule for an untrusted
# network to allow trusted network users to use all the
# Junos-supported ALGs initiated from the trusted network.
application-set junos-algs-outbound {
   application junos-ftp;
   application junos-tftp;
   application junos-rpc-portmap-tcp;
   application junos-rpc-portmap-udp;
   application junos-snmp-get;
   application junos-snmp-get-next;
   application junos-snmp-response;
   application junos-snmp-trap;
   application junos-rexec;
   application junos-rlogin;
   application junos-rsh;
   application junos-rtsp;
   application junos-sqlnet;
   application junos-traceroute;
   application junos-rpc-services-udp;
   application junos-rpc-services-tcp;
   application junos-icmp-all;
   application junos-netshow;
   application junos-netbios-name-udp;
   application junos-netbios-datagram;
   application junos-dce-rpc-portmap;
   application junos-dcerpc-msexchange-directory-rfr;
   application junos-dcerpc-msexchange-information-store;
   application junos-dcerpc-msexchange-directory-nsp;
}
# " junos-management-inbound" represents the group of applications
# that might need access to the trusted network from the untrusted
# network for management purposes.
# The set is intended for a UI to display management choices.
# NOTE: It is not recommended that you use the entire set directly in
# a firewall rule and open up firewall to all of these
# applications. Also, you should always specify the source
# and destination prefixes when using each application.
application-set junos-management-inbound {
   application junos-snmp-get;
   application junos-snmp-get-next;
   application junos-snmp-response;
   application junos-snmp-trap;
   application junos-ssh;
   application junos-telnet;
   application junos-http;
   application junos-https;
   application junos-xnm-ssl;
   application junos-xnm-clear-text;
   application junos-icmp-ping;
   application junos-traceroute-ttl-1;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  }
  }
```

To reference statements available from the **junos-defaults** group, include the selected **junos-*default-name*** statement at the applicable hierarchy level. To configure application protocols, see "Configuring Application Protocol Properties" on page 74; for details about a specific protocol, see "ALG Descriptions" on page 85.

Examples: Referencing the Preset Statement from the Junos Default Group

The following example is a preset statement from the Junos default groups that is available for FTP in a stateful firewall:

```
[edit]
groups {
  junos-defaults {
    applications {
      application junos-ftp { # Use FTP default configuration
      application-protocol ftp;
      protocol tcp;
      destination-port 21;
    }
  }
}
```

To reference a preset Junos default statement from the Junos default groups, include the **junos-*default-name*** statement at the applicable hierarchy level. For example, to reference the Junos default statement for FTP in a stateful firewall, include the **junos-ftp** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* from applications]** hierarchy level.

```
[edit]
services {
  stateful-firewall {
    rule my-rule {
      term my-term {
        from {
          applications junos-ftp; #Reference predefined statement, junos-ftp,
        }
      }
    }
  }
}
```

The following example shows configuration of the default Junos IP ALG:

```
[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          applications junos-ip;
```

Copyright © 2013, Juniper Networks, Inc.

```
            }
            then {
               accept;
               syslog;
            }
         }
      }
   }
}
```

If you configure the IP ALG in the stateful firewall rule, it is matched by any IP traffic, but if there is any other more specific application that matches the same traffic, the IP ALG will not be matched. For example, in the following configuration, both the ICMP ALG and the IP ALG are configured, but traffic is matched for ICMP packets, because it is the more specific match.

```
[edit]
services {
   stateful-firewall {
      rule r1 {
         match-direction input;
         term t1 {
            from {
               applications [ junos-ip junos-icmp-all ];
            }
            then {
               accept;
               syslog;
            }
         }
      }
   }
}
```

Examples: Configuring Application Protocols

The following example shows an application protocol definition describing a special FTP application running on port 78:

```
[edit applications]
application my-ftp-app {
   application-protocol ftp;
   protocol tcp;
   destination-port 78;
   timeout 100; # inactivity timeout for FTP service
}
```

The following example shows a special ICMP protocol (**application-protocol icmp**) of type 8 (ICMP echo):

```
[edit applications]
application icmp-app {
   application-protocol icmp;
   protocol icmp;
   icmp-type icmp-echo;
```


Copyright © 2013, Juniper Networks, Inc.

```
}
```

The following example shows a possible application set:

```
[edit applications]
application-set basic {
    http;
    ftp;
    telnet;
    nfs;
    icmp;
}
```

The software includes a predefined set of well-known application protocols. The set includes applications for which the TCP and UDP destination ports are already recognized by stateless firewall filters.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 5

Summary of Applications Configuration Statements

The following sections explain each of the applications configuration statements. The statements are organized alphabetically.

application

Syntax

```
application application-name {
  application-protocol protocol-name;
  destination-port port-number;
  icmp-code value;
  icmp-type value;
  inactivity-timeout value;
  protocol type;
  rpc-program-number number;
  snmp-command command;
  source-port port-number;
  ttl-threshold number;
  uuid hex-value;
}
```

Hierarchy Level

[edit applications],
[edit applications application-set *application-set-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure properties of an application and whether to include it in an application set.

Options

application-name—Identifier of the application.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring Application Protocol Properties" on page 74.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

application-protocol

Syntax application-protocol *protocol-name*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.
login options introduced in Junos OS Release 7.4.
ip option introduced in Junos OS Release 8.2.

Description Identify the application protocol name. Application protocols are also called application layer gateways (ALGs).

Options *protocol-name*—Name of the protocol. The following protocols are supported:

bootp—Bootstrap protocol

dce-rpc—DCE RPC

dce-rpc-portmap—DCE RPC portmap

dns—Domain Name Service

exec—Remote Execution Protocol

ftp—File Transfer Protocol

h323—H.323

icmp—ICMP

iiop—Internet Inter-ORB Protocol

ip—IP

login—Login

netbios—NetBIOS

netshow—NetShow

pptp—Point-to-Point Tunneling Protocol

realaudio—RealAudio

rpc—RPC

rpc-portmap—RPC portmap

rtsp—Real Time Streaming Protocol

shell—Shell

sip—Session Initiation Protocol

snmp—SNMP

sqlnet—SQLNet

talk–Talk Program

Copyright © 2013, Juniper Networks, Inc.

tftp—Trivial File Transfer Protocol

traceroute—Traceroute

winframe—WinFrame

Usage Guidelines See "Configuring an Application Protocol" on page 74.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

application-set

Syntax

```
application-set application-set-name {
    application application-name;
}
```

Hierarchy Level [edit applications]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure one or more applications to include in an application set.

Options ***application-set-name***—Identifier of an application set.

Usage Guidelines See "Configuring Application Sets" on page 85.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

applications

Syntax applications { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the applications used in services.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Application Properties

Copyright © 2013, Juniper Networks, Inc.

destination-port

Syntax destination-port *port-value*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number.

Options ***port-value***—Identifier for the port. For a complete list, see "Configuring Source and Destination Ports" on page 79.

Usage Guidelines See "Configuring Source and Destination Ports" on page 79.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

icmp-code

Syntax icmp-code *value*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Internet Control Message Protocol (ICMP) code value.

Options ***value***—The ICMP code value. For a complete list, see "Configuring the ICMP Code and Type" on page 78.

Usage Guidelines See "Configuring the ICMP Code and Type" on page 78.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

icmp-type

Syntax icmp-type *value*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description ICMP packet type value.

Options *value*—The ICMP type value, such as **echo** or **echo-reply**. For a complete list, see "Configuring the ICMP Code and Type" on page 78.

Usage Guidelines See "Configuring the ICMP Code and Type" on page 78.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

inactivity-timeout

Syntax inactivity-timeout *seconds*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Inactivity timeout period, in seconds.

Options *seconds*—Length of time the application is inactive before it times out.
Default: 30 seconds

Usage Guidelines See "Configuring the Inactivity Timeout Period" on page 82.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

learn-sip-register

Syntax learn-sip-register;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Activate SIP register to accept potential incoming SIP calls.

Usage Guidelines See "Configuring SIP" on page 74.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

protocol

Syntax protocol *type*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Networking protocol type or number.

Options *type*—Networking protocol type. The following text values are supported:

- **ah**
- **egp**
- **esp**
- **gre**
- **icmp**
- **igmp**
- **ipip**
- **ospf**
- **pim**
- **rsvp**
- **tcp**
- **udp**
- **vrrp**



NOTE: IP version 6 (IPv6) is not supported as a network protocol in application definitions.

Usage Guidelines See "Configuring the Network Protocol" on page 76.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rpc-program-number

Syntax rpc-program-number *number*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Remote procedure call (RPC) or Distributed Computing Environment (DCE) value.

Options *number*—RPC or DCE program value.
Range: 100,000 through 400,000

Usage Guidelines See "Configuring an RPC Program Number" on page 84.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

sip-call-hold-timeout

Syntax sip-call-hold-timeout *seconds*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Timeout period for SIP calls placed on hold, in seconds.

Options *seconds*—Length of time the application holds a SIP call open before it times out.
Default: 7200 seconds
Range: 0 through 36,000 seconds (10 hours)

Usage Guidelines See "Configuring SIP" on page 74.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

snmp-command

Syntax snmp-command *command*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description SNMP command format.

Options ***command***—Supported commands are SNMP **get**, **get-next**, **set**, and **trap**.

Usage Guidelines See "Configuring an SNMP Command for Packet Matching" on page 84.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

source-port

Syntax source-port *port-number*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Source port identifier.

Options ***port-value***—Identifier for the port. For a complete list, see "Configuring Source and Destination Ports" on page 79.

Usage Guidelines See "Configuring Source and Destination Ports" on page 79.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ttl-threshold

Syntax ttl-threshold *number*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the traceroute time-to-live (TTL) threshold value. This value sets the acceptable level of network penetration for trace routing.

Options ***number***—TTL threshold value.

Usage Guidelines See "Configuring the TTL Threshold" on page 84.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

uuid

Syntax uuid *hex-value*;

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the Universal Unique Identifier (UUID) for DCE RPC objects.

Options ***hex-value***—Hexadecimal value.

Usage Guidelines See "Configuring a Universal Unique Identifier" on page 84.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 6

Stateful Firewall Services Configuration Guidelines

To configure stateful firewall services, include the **stateful-firewall** statement at the **[edit services]** hierarchy level:

```
[edit services]
stateful-firewall {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        application-sets set-name;
        applications [ application-names ];
        destination-address (address | any-unicast) <except>;
        destination-address-range low minimum-value high maximum-value <except>;
        destination-prefix-list list-name <except>;
        source-address (address | any-unicast) <except>;
        source-address-range low minimum-value high maximum-value <except>;
        source-prefix-list list-name <except>;
      }
      then {
        (accept | discard | reject);
        allow-ip-options [ values ];
        syslog;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
}
```

This chapter contains the following sections:

- Configuring Stateful Firewall Rules on page 132
- Configuring Stateful Firewall Rule Sets on page 136
- Examples: Configuring Stateful Firewall Rules on page 136

Copyright © 2013, Juniper Networks, Inc.

Configuring Stateful Firewall Rules

To configure a stateful firewall rule, include the **rule *rule-name*** statement at the **[edit services stateful-firewall]** hierarchy level:

```
[edit services stateful-firewall]
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address address <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address address <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      (accept | discard | reject);
      allow-ip-options [ values ];
      syslog;
    }
  }
}
```

Each stateful firewall rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded. The **from** statement is optional in stateful firewall rules.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software. The **then** statement is mandatory in stateful firewall rules.

The following sections explain how to configure the components of stateful firewall rules:

- Configuring Match Direction for Stateful Firewall Rules on page 132
- Configuring Match Conditions in Stateful Firewall Rules on page 133
- Configuring Actions in Stateful Firewall Rules on page 134

Configuring Match Direction for Stateful Firewall Rules

Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services stateful-firewall rule *rule-name*]** hierarchy level:

```
[edit services stateful-firewall rule rule-name]
match-direction (input | output | input-output);
```


Copyright © 2013, Juniper Networks, Inc.

If you configure **match-direction input-output**, sessions initiated from both directions might match this rule.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output. For more information on inside and outside interfaces, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

On the PIC, a flow lookup is performed. If no flow is found, rule processing is performed. Rules in this service set are considered in sequence until a match is found. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered. Most packets result in the creation of bidirectional flows.

Configuring Match Conditions in Stateful Firewall Rules

To configure stateful firewall match conditions, include the **from** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services stateful-firewall rule rule-name term term-name]
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (address | any-unicast) <except>;
  destination-address-range low minimum-value high maximum-value <except>;
  destination-prefix-list list-name <except>;
  source-address (address | any-unicast) <except>;
  source-address-range low minimum-value high maximum-value <except>;
  source-prefix-list list-name <except>;
}
```

The source address and destination address can be either IPv4 or IPv6. You can use either the source address or the destination address as a match condition, in the same way that you would configure a firewall filter; for more information, see the Routing Policy Configuration Guide. You can use the wildcard value **any-unicast**, which denotes matching all unicast addresses.

Alternatively, you can specify a list of source or destination prefixes by configuring the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or the **source-prefix-list** statement in the stateful firewall rule. For an example, see "Examples: Configuring Stateful Firewall Rules" on page 136.

If you omit the **from** term, the stateful firewall accepts all traffic and the default protocol handlers take effect:

Copyright © 2013, Juniper Networks, Inc.

- User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP) create a bidirectional flow with a predicted reverse flow.
- IP creates a unidirectional flow.

You can also include application protocol definitions you have configured at the **[edit applications]** hierarchy level; for more information, see "Configuring Application Protocol Properties" on page 74.

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services stateful-firewall rule** *rule-name* **term** *term-name* **from]** hierarchy level.
- To apply one or more sets of application protocol definitions you have defined, include the **application-sets** statement at the **[edit services stateful-firewall rule** *rule-name* **term** *term-name* **from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit applications]** hierarchy level; you cannot specify these properties as match conditions.

Configuring Actions in Stateful Firewall Rules

To configure stateful firewall actions, include the **then** statement at the **[edit services stateful-firewall rule** *rule-name* **term** *term-name***]** hierarchy level:

```
[edit services stateful-firewall rule rule-name term term-name]
then {
    (accept | discard | reject);
    allow-ip-options [ values ];
    syslog;
}
```

You must include one of the following three possible actions:

- **accept**—The packet is accepted and sent on to its destination.
- **discard**—The packet is not accepted and is not processed further.
- **reject**—The packet is not accepted and a rejection message is returned; UDP sends an ICMP unreachable code and TCP sends RST. Rejected packets can be logged or sampled.

You can optionally configure the firewall to record information in the system logging facility by including the **syslog** statement at the **[edit services stateful-firewall rule** *rule-name* **term** *term-name* **then]** hierarchy level. This statement overrides any **syslog** setting included in the service set or interface default configuration.

Copyright © 2013, Juniper Networks, Inc.

Configuring IP Option Handling

You can optionally configure the firewall to inspect IP header information by including the **allow-ip-options** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* then]** hierarchy level. When you configure this statement, all packets that match the criteria specified in the **from** statement are subjected to additional matching criteria. A packet is accepted only when all of its IP option types are configured as values in the **allow-ip-options** statement. If you do not configure **allow-ip-options**, only packets without IP header options are accepted.

The additional IP header option inspection applies only to the **accept** and **reject** stateful firewall actions. This configuration has no effect on the **discard** action. When the IP header inspection fails, reject frames are not sent; in this case, the **reject** action has the same effect as **discard**.

If an IP option packet is accepted by the stateful firewall, Network Address Translation (NAT) and intrusion detection service (IDS) are applied in the same way as to packets without IP option headers. The IP option configuration appears only in the stateful firewall rules; NAT applies to packets with or without IP options.

When a packet is dropped because it fails the IP option inspection, this exception event generates both IDS event and system log messages. The event type depends on the first IP option field rejected.

Table 12 on page 135 lists the possible values for the **allow-ip-options** statement. You can include a range or set of numeric values, or one or more of the predefined IP option settings. You can enter either the option name or its numeric equivalent. For more information, refer to http://www.iana.org/assignments/ip-parameters.

Table 12: IP Option Values

IP Option Name	Numeric Value	Comment
any	**0**	Any IP option
ip-security	**130**	–
ip-stream	**136**	–
loose-source-route	**131**	–
route-record	**7**	–
router-alert	**148**	–
strict-source-route	**137**	–
timestamp	**68**	–

Copyright © 2013, Juniper Networks, Inc.

Configuring Stateful Firewall Rule Sets

The **rule-set** statement defines a collection of stateful firewall rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services stateful-firewall]** hierarchy level with a **rule** statement for each rule:

```
[edit services stateful-firewall]
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Examples: Configuring Stateful Firewall Rules

The following example show a stateful firewall configuration containing two rules, one for input matching on a specified application set and the other for output matching on a specified source address:

```
[edit services]
stateful-firewall {
  rule Rule1 {
    match-direction input;
    term 1 {
      from {
        application-sets Applications;
      }
      then {
        accept;
      }
    }
    term accept {
      then {
        accept;
      }
    }
  }
  rule Rule2 {
    match-direction output;
    term Local {
      from {
        source-address {
          10.1.3.2/32;
        }
      }
      then {
        accept;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

The following example has a single rule with two terms. The first term rejects all traffic in **my-application-group** that originates from the specified source address, and provides a detailed system log record of the rejected packets. The second term accepts Hypertext Transfer Protocol (HTTP) traffic from anyone to the specified destination address.

```
[edit services stateful-firewall]
rule my-firewall-rule {
  match-direction input-output;
  term term1 {
    from {
      source-address 10.1.3.2/32;
      application-sets my-application-group;
    }
    then {
      reject;
      syslog;
    }
  }
  term term2 {
    from {
      destination-address 10.2.3.2/32;
      applications http;
    }
    then {
      accept;
    }
  }
}
```

The following example shows use of source and destination prefix lists. This requires two separate configuration items.

You configure the prefix list at the **[edit policy-options]** hierarchy level:

```
[edit]
policy-options {
  prefix-list p1 {
    1.1.1.1/32;
    2.2.2.0/24;
  }
  prefix-list p2 {
    3.3.3.3/32;
    4.4.4.0/24;
  }
}
```

You reference the configured prefix list in the stateful firewall rule:

```
[edit]
services {
  stateful-firewall {
    rule r1 {
```

Copyright © 2013, Juniper Networks, Inc.

```
        match-direction input;
        term t1 {
          from {
            source-prefix-list {
              p1;
            }
            destination-prefix-list {
              p2;
            }
          }
          then {
            accept;
          }
        }
      }
    }
  }
```

This is equivalent to the following configuration:

```
[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          source-address {
            1.1.1.1/32;
            2.2.2.0/24;
          }
          destination-address {
            3.3.3.3/32;
            4.4.4.0/24;
          }
        }
        then {
          accept;
        }
      }
    }
  }
}
```

You can use the **except** qualifier with the prefix lists, as in the following example. In this case, the **except** qualifier applies to all prefixes included in prefix list **p2**.

```
[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          source-prefix-list {
            p1;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
            destination-prefix-list {
              p2 except;
            }
          }
          then {
            accept;
          }
        }
      }
    }
  }
```

For additional examples that combine stateful firewall configuration with other services and with virtual private network (VPN) routing and forwarding (VRF) tables, see the configuration examples.

Related Documentation

- Example: BOOTP and Broadcast Addresses on page 71
- Example: NAT Between VRFs Configuration on page 68
- Example: Dynamic Source NAT as a Next-Hop Service on page 67
- Example: VPN Routing and Forwarding (VRF) and Service Configuration on page 65
- Example: Service Interfaces Configuration on page 63
- Example: Configuring Layer 3 Services and the Services SDK on Two PICs

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 7

Summary of Stateful Firewall Configuration Statements

The following sections explain each of the stateful firewall services statements. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

allow-ip-options

Syntax allow-ip-options [*values*];

Hierarchy Level [edit services (stateful-firewall) stateful-firewall rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure how the stateful firewall handles IP header information. This statement is optional.

Options ***value***—Can be a set or range of numeric values, or one or more of the following predefined option types. You can enter either the option name or its numeric equivalent.

Option Name	Numeric Value
any	0
ip-security	130
ip-stream	8
loose-source-route	3
route-record	7
router-alert	148
strict-source-route	9
timestamp	4

Usage Guidelines See "Configuring Actions in Stateful Firewall Rules" on page 134.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

application-sets

Syntax applications-sets *set-name*;

Hierarchy Level [edit services (stateful-firewall) stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define one or more target application sets.

Options ***set-name***—Name of the target application set.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

applications

Syntax applications [*application-names*];

Hierarchy Level [edit services (stateful-firewall) stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define one or more applications to which the stateful firewall services apply.

Options ***application-name***—Name of the target application.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination-address

Syntax destination-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services (stateful-firewall) stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.
any-unicast and **except** options introduced in Junos OS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the destination address for rule matching.

Options ***address***—Destination IPv4 or IPv6 address or prefix value.

any-unicast—Match all unicast packets.

except—(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

destination-address-range

Syntax destination-address-range low *minimum-value* high *maximum-value* <except>;

Hierarchy Level [edit services (stateful-firewall) stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 7.6.
minimum-value and ***maximum-value*** options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the destination address range for rule matching.

Options ***minimum-value***—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exclude the specified address range from rule matching.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination-prefix-list

Syntax destination-prefix-list *list-name* <except>;

Hierarchy Level [edit services (stateful-firewall) stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the destination prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Routing Policy Configuration Guide

Copyright © 2013, Juniper Networks, Inc.

from

Syntax

```
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (address | any-unicast) <except>;
  destination-address-range low minimum-value high maximum-value <except>;
  destination-prefix-list list-name <except>;
  source-address (address | any-unicast) <except>;
  source-address-range low minimum-value high maximum-value <except>;
  source-prefix-list list-name <except>;
}
```

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify input conditions for a stateful firewall term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Stateful Firewall Rules" on page 132.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

match-direction

Syntax match-direction (input | output | input-output);

Hierarchy Level [edit services stateful-firewall rule *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the direction in which the rule match is applied.

Options **input**—Apply the rule match on the input side of the interface.

output—Apply the rule match on the output side of the interface.

input-output—Apply the rule match bidirectionally.

Usage Guidelines See "Configuring Stateful Firewall Rules" on page 132.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rule

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (address | any-unicast) <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address (address | any-unicast) <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      (accept | discard | reject);
      syslog;
    }
  }
}
```

Hierarchy Level [edit services stateful-firewall],
[edit services stateful-firewall rule-set *rule-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule the router uses when applying this service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Stateful Firewall Rules" on page 132.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rule-set

Syntax
```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services stateful-firewall]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

Usage Guidelines See "Configuring Stateful Firewall Rule Sets" on page 136.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

services (stateful-firewall)

Syntax services stateful-firewall { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 10.4.

Description Define the service rules to be applied to traffic.

Options **stateful-firewall**—Identifies the stateful firewall set of rules statements.

Usage Guidelines See Stateful Firewall for Juniper Service Framework (JSF).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

source-address

Syntax source-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.
any-unicast and **except** options introduced in Junos OS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Source address for rule matching.

Options ***address***—Source IPv4 or IPv6 address or prefix value.

any-unicast—Any unicast packet.

except—(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

source-address-range

Syntax source-address-range low *minimum-value* high *maximum-value* <except>;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 7.6.
minimum-value and ***maximum-value*** options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Source address range for rule matching.

Options ***minimum-value***—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

source-prefix-list

Syntax source-prefix-list *list-name* <except>;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the source prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Usage Guidelines See "Configuring Match Conditions in Stateful Firewall Rules" on page 133.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Routing Policy Configuration Guide

syslog

Syntax syslog;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the **/var/log** directory. This setting overrides any **syslog** statement setting included in the service set or interface default configuration.

Usage Guidelines See "Configuring Actions in Stateful Firewall Rules" on page 134.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

term

Syntax

```
term term-name {
  from {
    application-sets set-name;
    applications [ application-names ];
    destination-address (address | any-unicast) <except>;
    destination-address-range low minimum-value high maximum-value <except>;
    destination-prefix-list list-name <except>;
    source-address (address | any-unicast) <except>;
    source-address-range low minimum-value high maximum-value <except>;
    source-prefix-list list-name <except>;
  }
  then {
    (accept | discard | reject);
    syslog;
  }
}
```

Hierarchy Level [edit services stateful-firewall rule *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the stateful firewall term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Stateful Firewall Rules" on page 132.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

then

Syntax

```
then {
  (accept | discard | reject);
  syslog;
}
```

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the stateful firewall term actions. You can configure the router to accept, discard, or reject the targeted traffic. The other actions are optional.

Options

accept—Accept the traffic and send it on to its destination.

discard—Do not accept traffic or process it further.

reject—Do not accept the traffic and return a rejection message. Rejected traffic can be logged or sampled.

The remaining statement is explained separately.

Usage Guidelines See "Configuring Actions in Stateful Firewall Rules" on page 134.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Routing Policy Configuration Guide

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 8

Stateful Firewall on the Embedded Junos OS Platform Configuration Guidelines

Till now, all services run only on the Juniper microkernel software platform. However, some services will now be deployed on the Embedded Junos software platform. This allows such services to be coupled with third-party applications. Starting with Junos OS Release 9.5, the stateful firewall service has been implemented using the embedded Junos Application Framework (eJAF). The stateful firewall plug-in described in the following sections supports many of the features of the existing stateful firewall service that runs on the Juniper microkernel.

This chapter contains the following sections:

- Loading the Stateful Firewall Plug-In on page 153
- Configuring Memory for the Stateful Firewall Plug-In on page 155
- Configuring rsh, rlogin, rexec for Stateful Firewall on page 155

Loading the Stateful Firewall Plug-In

As of Junos OS Release 9.5, a stateful firewall plug-in is provided as part of the jbundle package. To load this plug-in on the PIC, include the **package jservices-sfw** statement at the **[edit chassis fpc *slot-number* pic *slot-number* adaptive-services service-package extension-provider]** hierarchy level. For example:

```
user@host# show chassis
fpc 0 {
    pic 2 {
        adaptive-services {
            service-package {
                extension-provider {
                    control-cores 1;
                    data-cores 7;
                    object-cache-size 512;
                    package jservices-sfw; #Loads stateful firewall plug-in.
                    policy-db-size 64;
                }
            }
        }
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

You can load both the **jservices-sfw** package and a Junos SDK application package on the same PIC.

The following example demonstrates the stateful firewall plug-in coexisting with a provider's plug-in:

```
[edit]
services {
   service-set sset {
      stateful-firewall-rules rule1;
      interface-service {
         service-interface ms-0/0/0;
      }
      extension-service customer-plugin;
      service-order {
         forward-flow [ stateful-firewall customer-plugin ];
      }
   }
   stateful-firewall {
      rule rule1 {
         match-direction input-output;
         term term1 {
            from {
               applications junos-ftp;
            }
            then {
               accept;
            }
         }
      }
      rule rule2 {
         match-direction input;
         term term1 {
            from {
               source-address {
                  192.1.1.2/32;
               }
               then {
                  reject;
                  syslog;
               }
            }
         }
      }
   }
}
```

The following stateful firewall operational commands support the **ms-** interface:

- **show services stateful-firewall flows**—Display stateful firewall flow table entries.
- **show services stateful-firewall statistics**—Display stateful firewall statistics. For this command, only rule and ALG statistics are given. In the extensive option, other statistics appear but do not populate correctly; those values are all zeroes.
- **clear services stateful-firewall flows**—Remove established flows from the flow table.

Copyright © 2013, Juniper Networks, Inc.

The commands are described in the Junos OS Operational Mode Commands.

Related Documentation

- Configuring Memory for the Stateful Firewall Plug-In on page 155
- extension-provider on page 160

Configuring Memory for the Stateful Firewall Plug-In

When configuring the stateful firewall internal plug-in, some questions remain regarding the upper limit to specify for the **policy-db-size**, **object-cache-size**, and **forwarding-db-size** statements when the application needs to use a large number of rules, causing the total memory required to approach the size of the object cache configured. The following limits, which are specific to the stateful firewall configuration, await additional review:

- Maximum number of terms (with one rule per term) per service set: 1200
- Maximum number of service sets per Multiservices PIC: 4000 (Juniper Networks M Series Multiservice Edge Routers and T Series Core Routers), 6000 (Juniper Networks MX Series 3D Universal Edge Routers and M120 Multiservice Edge Routers)
- Maximum object cache size: 1280 MB (Multiservices 400 PICs and DPCs), 512 MB (Multiservices 100 PICs)
- Maximum policy database size: Still to be determined.

If the policy database is set too small, an error message is logged in the router message file even though the commit may appear to be successful. It is necessary to check the logs to make sure that no message file error is found to be sure that the stateful firewall commit was indeed successful. The remedial action is to increase the size of the policy database.

Related Documentation

- extension-provider on page 160

Configuring rsh, rlogin, rexec for Stateful Firewall

Some implementations of the rsh, rlogin, rexec mechanism require the remote host to authenticate the request by opening a separate TCP session to port 113 on the client host. By default, the stateful firewall does not allow this authentication flow to go through.

To open the authentication flow, include the **applications junos-ident** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* from]** hierarchy level:

```
[edit]
services {
  stateful-firewall {
    rule rule1 {
      term term1 {
        from {
          (source-address | destination-address);
          applications junos-ident;
        }
        then {
```

```
            accept;
          }
        }
      }
    }
  }
```

To allow Kerberos-enabled rsh, rlogin, rexec through the stateful firewall, configure the following additional applications and include them in the stateful firewall terms:

```
[edit]
applications {
  application test-kerberos-kshell {
    Protocol tcp;
    destination-port kshell;
  }
  application test kerberos-klogin {
    protocol tcp;
    destination-port klogin;
  }
}

services {
  stateful-firewall {
    rule rule1 {
      term term1 {
        from {
          applications [kerberos-klogin kerberos-kshell];
        }
        then {
          accept;
        }
      }
    }
  }
}
```

Related Documentation

- Configuring Memory for the Stateful Firewall Plug-In on page 155

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 9

Summary of Stateful Firewall on the Embedded Junos OS Platform Configuration Statements

The following sections explain stateful firewall statements used in SDK applications. The statements are organized alphabetically.

control-cores

Syntax control-cores *control-number*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure control cores. Any cores not configured as either control or data cores are treated as user cores. When the number of control cores is changed, the PIC reboots.

Options ***control-number***—Number of control cores. At least one core must be a control core.

Range: 1 through 8

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Control and Data Cores
- data-cores on page 158

Copyright © 2013, Juniper Networks, Inc.

data-cores

Syntax data-cores *data-number*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure data cores. Any cores not configured as either data or control cores are treated as user cores. When the number of data cores is changed, the PIC reboots.

Options ***data-number***—Number of data cores. Although it is not mandatory to dedicate any cores as data cores, it is advisable, depending on the nature of the application, to dedicate a minimum of five as data cores to achieve good performance.

Range: 0 through 7

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Control and Data Cores
- control-cores on page 157

data-flow-affinity

Syntax
```
data-flow-affinity {
  hash-key (layer-3 | layer-4);
}
```

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable flow affinity distribution for packets over data CPUs on the PIC. Once enabled, the default behavior distributing data packets changes from a round-robin distribution to a flow affinity distribution based on a hash distribution. Adding or deleting this statement causes the PIC to reboot.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Packet Distribution Settings

Copyright © 2013, Juniper Networks, Inc.

destination (Chassis)

Syntax destination *destination*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider syslog *facility*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure where log messages go. By default, all messages go to the **/var/log** directory on the Routing Engine. Enhancements to the existing infrastructure make debugging on the Multiservices PIC easier by giving the user the option of redirecting log messages. When the **syslog destination** statement is configured to redirect the log messages, you can use the **set system syslog** command, a command available in the native Junos OS CLI, to override the syslog settings made on the Multiservices PIC.

Options ***destination***—Choose one of the following options:

- **routing-engine**—Forward log messages to the Routing Engine.
- **pic-console**—Forward log messages to the console of the PIC.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring System Log Messages
- extension-provider on page 160

Copyright © 2013, Juniper Networks, Inc.

extension-provider

Syntax

```
extension-provider {
  control-cores control-number;
  data-cores data-number;
  data-flow-affinity {
    hash-key (layer-3 | layer-4);
  }
  forwarding-db-size size;
  object-cache-size size;
  package package-name;
  policy-db-size size;
  syslog {
    facility {
      severity;
      destination destination;
    }
  }
  wired-process-mem-size mem-size;
}
```

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure an application on a PIC. When the **extension-provider** statement is first configured, the PIC reboots.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Control and Data Cores
- Configuring Packet Distribution Settings
- Configuring Memory Settings
- Configuring Packages on the PIC
- Configuring System Log Messages

Copyright © 2013, Juniper Networks, Inc.

forwarding-db-size

Syntax forwarding-db-size *size*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the size of the forwarding database (FDB). When this setting is changed, the PIC reboots.



NOTE: You need to enable the forwarding-options sampling statement for the FDB to be created.

Options *size*—Size of the FDB, in megabytes (MB). The size of the FDB and the size of the policy database together must be smaller than the size of the object cache.

Range: 0 through 12879 MB

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Memory Settings
- policy-db-size on page 164
- wired-process-mem-size on page 166
- object-cache-size on page 163

Copyright © 2013, Juniper Networks, Inc.

hash-key (Chassis)

Syntax hash-key (layer-3 | layer-4);

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider data-flow-affinity]

Release Information Statement introduced in Junos OS Release 10.2.

Description Set the hashing distribution of flow affinity. This is an optional setting. Once the **data-flow-affinity** statement is enabled, you may need to choose the hashing distribution. Modifying this statement causes the PIC to reboot.

Default If you do not configure the **hash-key** statement, the hashing distribution is 5-tuple hashing, or **layer-4**.

Options **layer-3**—3-tuple hashing (source IP address, destination IP address, and IP protocol).

layer-4—5-tuple hashing (3-tuple plus source and destination TCP or UDP ports).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Packet Distribution Settings
- extension-provider on page 160

 Copyright © 2013, Juniper Networks, Inc.

object-cache-size

Syntax object-cache-size *value*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Description Configure the size of the object cache. When this setting is changed, the PIC reboots.

Options *value*—Amount of object cache, in MB. Only values in increments of 128 MB are allowed.

Range: For Multiservices 100 PIC, range is 128 MB through 512 MB. If the **wired-process-mem-size** statement at the same hierarchy level has a value of 512 MB, the maximum value for this statement is 128 MB.

Range: For Multiservices 400 PIC, range is 128 MB through 1280 MB. If the **wired-process-mem-size** statement at the same hierarchy level has a value of 512 MB, the maximum value for this statement is 512 MB.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Memory Settings
- forwarding-db-size on page 161
- policy-db-size on page 164
- wired-process-mem-size on page 166

package (Loading on PIC)

Syntax package *package-name*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Release Information Statement introduced in Junos OS Release 9.1.

Description Identify a package to be loaded on the PIC. When a package is added or removed, the PIC reboots.

Options *package-name*—Name of the package to be loaded on the PIC. There can be up to eight packages loaded on a PIC; however, only one data package is allowed per PIC. An error message is displayed if more than eight packages are specified.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Packages on the PIC

Copyright © 2013, Juniper Networks, Inc.

policy-db-size

Syntax policy-db-size *size*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Description Configure the size of the policy database. When this setting is changed, the PIC reboots.



NOTE: At least one data core must be configured to configure the size of the policy database.

Options ***size***—Size of the policy database, in megabytes (MB). The size of the forwarding database and the size of the policy database together must be smaller than the size of the object cache.

Range: 0 through 1279 MB

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Memory Settings
- forwarding-db-size on page 161
- object-cache-size on page 163
- wired-process-mem-size on page 166

Copyright © 2013, Juniper Networks, Inc.

syslog (Chassis)

Syntax

```
syslog {
  facility {
    severity;
    destination destination;
  }
}
```

Hierarchy Level

[edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Release Information

Statement introduced in Junos OS Release 9.2.
Options **daemon** and **kernel** (for ***facility***) introduced in Junos OS Release 9.5.

Description

Enable PIC system logging to record or view system log messages on a specific PIC. The system log information is passed to the kernel for logging in the **/var/log** directory.

Options

facility—Group of messages that are either generated by the same software process or concern a similar condition or activity. Possible values include the following: **daemon**, **external**, **kernel**, and **pfe**.

severity—Classification of effect on functioning. Possible values are the following options:

- **any**—Include all severity levels.
- **none**—Disable logging of the associated facility to a destination.
- **emergency**—System panic or other condition that causes the routing platform to stop functioning.
- **alert**—Conditions that require immediate correction, such as a corrupted system database.
- **critical**—Critical conditions, such as hard errors.
- **error**—Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels.
- **warning**—Conditions that warrant monitoring.
- **notice**—Conditions that are not errors but might warrant special handling.
- **info**—Events or nonerror conditions of interest.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring System Log Messages

Copyright © 2013, Juniper Networks, Inc.

wired-process-mem-size

Syntax wired-process-mem-size *mem-size*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]

Description Configure the size of the reserved wired process memory. You can also configure object cache. If this setting is changed, the PIC reboots.

Options ***megabytes***—Size of the reserved wired process memory, in MB. The only size you can set for this statement is 512 MB.

Default: 512 MB

Range: 0 through 512 MB

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Memory Settings
- forwarding-db-size on page 161
- object-cache-size on page 163
- policy-db-size on page 164
- wired-max-processes

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 10

Load Balancing Configuration Guidelines

As of now, most router services are provisioned using service sets in Junos OS. Each service set directs traffic to a specific preconfigured services PIC only. This leads to inefficient use of networking resources within a system. Load balancing resolves this situation by allowing distribution of ingress and egress traffic across multiple services PICs. Load balancing works by hashing each packet and then redirecting the packet to the appropriate services PIC. Load balancing can be accomplished only on MX Series 3D Universal Edge routers because services PICs require symmetric hashing to ensure that ingress and egress traffic are directed properly.

- Configuring Load Balancing on AMS Infrastructure on page 167
- Example: Configuring Static Source Translation on AMS Infrastructure on page 169

Configuring Load Balancing on AMS Infrastructure

Configuring load balancing requires an aggregated Multiservices (AMS) system. AMS involves grouping several Multiservices PICs together. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs. Starting with Junos OS 11.4, high availability (HA) is supported on AMS infrastructure on all MX Series 3D Universal Edge routers. AMS has several benefits:

- Support for configuring behavior if a Multiservices PIC that is part of the AMS configuration fails
- Support for specifying hash keys for each service set in either direction
- Support for adding routes to individual PICs within the AMS system

Configuring AMS Infrastructure

AMS supports load balancing across multiple service sets. All ingress or egress traffic for a service set can be load balanced across different services PICs. To enable load balancing, you have to configure an aggregate interface with existing services interfaces.

To configure failure behavior in AMS, include the **member-failure-options** statement:

```
[edit interfaces ams1]
load-balancing-options {
   member-failure-options {
```

Copyright © 2013, Juniper Networks, Inc.

```
        drop-member-traffic {
          rejoin-timeout rejoin-timeout;
        }
        redistribute-all-traffic {
          enable-rejoin;
        }
      }
    }
```

If a PIC fails, the traffic to the failed PIC can be configured to be redistributed by using the **redistribute-all-traffic** statement at the **[edit interfaces *interface-name* load-balancing-options member-failure-options]** hierarchy level. If the **drop-member-traffic** statement is used, all traffic to the failed PIC is dropped. Both options are mutually exclusive.



NOTE: If **member-failure-options** is not explicitly configured, the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Only **mams-** interfaces (services interfaces that are part of AMS) can be aggregated. After an AMS interface has been configured, the constituent **mams-** interfaces cannot be individually configured. A **mams-** interface cannot be used as an **rms** interface. AMS supports only IPv4; inet6 family is not supported. It is not possible to configure addresses on an AMS interface. Network Address Translation (NAT) is the only application that runs on AMS infrastructure at this time.



NOTE: Unit 0 on an AMS interface cannot be configured.

To support multiple applications and different types of translation, AMS infrastructure supports configuring hashing for each service set. The hash keys can be configured separately for ingress and egress. The default configuration uses source IP, destination IP, and the protocol for hashing; incoming-interface for ingress and outgoing-interface for egress are also available.

Configuring High Availability

In an AMS system configured with high availability, a designated Multiservices PIC acts as a backup for other active PICs that are part of the AMS system. Presently, only N:1 backup for high availability is supported; only one PIC is available as backup for all other active PICs. High availability for load balancing is configured by adding the **high-availability-options** statement at the **[edit interfaces *interface-name* load-balancing-options]** hierarchy level.

To configure high availability, include the **high-availability-options** statement:

```
[edit interfaces ams1]
load-balancing-options {
  high-availability-options {
    many-to-one {
      preferred-backup preferred-backup;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
```

Load Balancing Network Address Translation Flows

Starting with Junos OS Release 11.4, Network Address Translation (NAT) has been programmed as a plug-in and is a function of load balancing and high availability. The plug-in runs on AMS infrastructure. All flows for translation are automatically distributed to different services PICs that are part of the AMS infrastructure. In case of failure of an active Multiservices PIC, the configured backup Multiservices PIC will take over the NAT pool resources of the failed PIC. The hashing method selected depends on the type of NAT. Using NAT on AMS infrastructure has a few limitations:

- NAT flows to failed PICs cannot be restored.
- There is no support for IPv6 flows.
- Twice NAT is not supported for load balancing.

See "Example: Configuring Static Source Translation on AMS Infrastructure" on page 169 for more details on configuring NAT flows for load balancing.

Example: Configuring Static Source Translation on AMS Infrastructure

This example shows a static source translation configured on an AMS interface. The flows will be load balanced across member interfaces with this example.

Configure the AMS interface **ams0** with load balancing options.

```
[edit interfaces ams0]
load-balancing-options {
member-interface mams-5/0/0;
member-interface mams-5/1/0;
}
unit 1 {
    family inet;
}
unit 2 {
    family inet;
}
```

Configure hashing for the service set for both ingress and egress traffic.

```
[edit services service-set ss1]
interface-service {
    service-interface ams0.1;
    load-balancing-options {
        hash-keys {
            ingress-key destination-ip;
            egress-key source-ip;
        }
    }
}
```

Copyright © 2013, Juniper Networks, Inc.



NOTE: Hashing is determined based on whether the service set is applied on the ingress or egress interface.

Configure two NAT pools because you have configured two member interfaces for the AMS interface.

```
[edit services]
nat {
    pool p1 {
        address-range low 20.1.1.80 high 20.1.1.80;
    }
pool p2 {
        address 20.1.1.81/32;
    }
}
```

Configure the NAT rule and translation.

```
[edit services]
nat {
    rule r1 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    20.1.1.2/32;
                }
            }
            then {
                translated {
                    source-pool p1;
                    translation-type {
                        basic-nat44;
                    }
                }
           }
        term t1 {
              from {
                  source-address {
                      40.1.1.2/32;
                  }
             }
             then {
                 translated {
                     source-pool p2;
                     translation-type {
                         basic-nat44;
                     }
                 }
            }
      }
 }
```


Copyright © 2013, Juniper Networks, Inc.



NOTE: A similar configuration can be applied for translation types **dynamic-nat44** and **napt-44**. Twice NAT cannot run on AMS infrastructure at this time.

Related Documentation

- Configuring Load Balancing on AMS Infrastructure on page 167

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 11

Summary of Load Balancing Configuration Statements

The following sections explain each of the load balancing and aggregated Multiservices (AMS) statements. The statements are organized alphabetically.

drop-member-traffic (Aggregated Multiservices)

Syntax

```
drop-member-traffic {
  rejoin-timeout rejoin-timeout;
}
```

Hierarchy Level [edit interfaces *interface-name* load-balancing-options member-failure-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify whether the broadband gateway should drop traffic to a Multiservices PIC when it fails.

For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration is valid only when two or more Multiservices PICs have failed.

The remaining statement is explained separately.

Default If this statement is not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- member-failure-options (Aggregated Multiservices) on page 179

Copyright © 2013, Juniper Networks, Inc.

enable-rejoin (aggregated Multiservices)

Syntax enable-rejoin;

Hierarchy Level [edit interfaces *interface-name* load-balancing-options member-failure-options redistribute-all-traffic]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enable the failed member to rejoin the aggregated Multiservices (AMS) interface after the member comes back online.

For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration allows the failed members to rejoin the pool of active members automatically.

Default If you do not configure this option, then the failed members do not automatically rejoin the **ams** interface even after coming back online.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- redistribute-all-traffic (Aggregated Multiservices) on page 182

family (aggregated Multiservices)

Syntax family *family*;

Hierarchy Level [edit interfaces *interface-name* unit *interface-unit-number*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure protocol family information for the logical interface.

Options ***family***—Protocol family. Currently, only one option, **inet** (IP version 4 suite), is supported.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- unit (Aggregated Multiservices) on page 183

Copyright © 2013, Juniper Networks, Inc.

high-availability-options (aggregated Multiservices)

Syntax

```
high-availability-options {
  many-to-one {
    preferred-backup preferred-backup;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* load-balancing-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the high availability options for the aggregated Multiservices (AMS) interface. For service applications, if only the load-balancing feature is being used, then this configuration is optional.

For many-to-one (N:1) high availability support for service applications like Network Address Translation (NAT), the preferred backup Multiservices PIC, in hot standby mode, backs up one or more (N) active Multiservices PICs.



NOTE: In both cases, if one of the active Multiservices PICs goes down, then the backup replaces it as the active Multiservices PIC. When the failed PIC comes back up, it becomes the new backup. This is called floating backup.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- load-balancing-options on page 177

Copyright © 2013, Juniper Networks, Inc.

interfaces (Aggregated Multiservices)

Syntax

```
interfaces interface-name {
  load-balancing-options {
    high-availability-options {
      many-to-one {
        preferred-backup preferred-backup;
      }
    }
    member-failure-options {
      drop-member-traffic {
        rejoin-timeout rejoin-timeout;
      }
      redistribute-all-traffic {
        enable-rejoin;
      }
    }
    member-interface interface-name;
  }
  unit interface-unit-number {
    family family;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the aggregated Multiservices (AMS) interface. The AMS interface provides the infrastructure for load balancing and high availability (HA).

NOTE: The interfaces must be valid aggregated Multiservices interfaces (ams)—for example, ams0 or ams1, and so on. The ams infrastructure is supported only in chassis with Trio-based modules and Multiservices Dense Port Concentrators (MS-DPCs).

The remaining statements are explained separately.

Options **interface-name**—Name of the aggregated Multiservices interface (**ams**)—for example, **ams0** or **ams1**, and so on.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Load Balancing on AMS Infrastructure on page 167

Copyright © 2013, Juniper Networks, Inc.

load-balancing-options (Aggregated Multiservices)

Syntax

```
load-balancing-options {
  high-availability-options {
    many-to-one {
      preferred-backup preferred-backup;
    }
  }
  member-failure-options {
    drop-member-traffic {
      rejoin-timeout rejoin-timeout;
    }
    redistribute-all-traffic {
      enable-rejoin;
    }
  }
  member-interface interface-name;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the high availability (HA) options for the aggregated Multiservices (AMS) interface.

Many-to-one (N:1) high availability mode for service applications like Network Address Translation (NAT) is supported. In this case, one Multiservices PIC is the backup (in hot standby mode) for one or more (N) active Multiservices PICs. If one of the active Multiservices PICs goes down, then the backup replaces it as the active Multiservices PIC. When the failed PIC comes back online, it becomes the new backup. This is called floating backup mode.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- interfaces on page 176

Copyright © 2013, Juniper Networks, Inc.

many-to-one (Aggregated Multiservices)

Syntax

```
many-to-one {
  preferred-backup preferred-backup;
}
```

Hierarchy Level [edit interfaces *interface-name* load-balancing-options high-availability-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the initial preferred backup for the aggregated Multiservices (AMS) interface.



NOTE: The preferred backup must be one of the member interfaces (mams–) that have already been configured at the [edit interfaces *interface-name* load-balancing-options] hierarchy level. Even in the case of mobile control plane redundancy, which is one-to-one (1:1), the initial preferred backup is configured at this hierarchy level.

The remaining statements are explained separately.

Options

preferred-backup *preferred-backup*—Name of the preferred backup member interface. The member interface format is **mams-a/b/0**, where **a** is the Flexible PIC Concentrator (FPC) slot number and **b** is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- high-availability-options (aggregated Multiservices) on page 175

Copyright © 2013, Juniper Networks, Inc.

member-failure-options (Aggregated Multiservices)

Syntax

```
member-failure-options {
  drop-member-traffic {
    rejoin-timeout rejoin-timeout;
  }
  redistribute-all-traffic {
    enable-rejoin;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* load-balancing-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the possible behavior for the aggregated Multiservices (AMS) interface in case of failure of more than one active member.

NOTE: The **drop-member-traffic** configuration and the **redistribute-all-traffic** configuration are mutually exclusive.

Table 13 on page 179 displays the behavior of the member interface after the failure of the first Multiservices PIC. Table 14 on page 180 displays the behavior of the member interface after the failure of two Multiservices PICs.

NOTE: The AMS infrastructure has been designed to handle one failure automatically. However, in the unlikely event that more than one Multiservices PIC fails, the AMS infrastructure provides configuration options to minimize the impact on existing traffic flows.

Table 13: Behavior of Member Interface After One Multiservices PIC Fails

High Availability Mode	Member Interface Behavior
Many-to-one (N:1) high availability support for service applications	Automatically handled by the AMS infrastructure

Copyright © 2013, Juniper Networks, Inc.

Table 14: Behavior of Member Interface After Two Multiservices PICs Fail

High Availability Mode	Configuration	rejoin-timeout	Behavior when member rejoins before rejoin-timeout expires	Behavior when member rejoins after rejoin-timeout expires
Many-to-one (N:1) high availability support for service applications	**drop-member-traffic**	Configured	The existing traffic for the second failed member will *not* be redistributed to the other members. The first member to rejoin becomes an active member. The second member to rejoin becomes the backup. This behavior is handled automatically by the AMS infrastructure.	The existing traffic for the second failed member will *not* be redistributed to the other members. The first member will rejoin the AMS automatically. However, the other members who are rejoining will be moved to the discard state.
Many-to-one (N:1) high availability support for service applications	**redistribute-all-traffic**	Not applicable	Before rejoin, the traffic is redistributed to existing active members. After a failed member rejoins, the traffic is load-balanced afresh. This may impact existing traffic flows.	

The remaining statements are explained separately.

Default If **member-failure-options** are not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- load-balancing-options (Aggregated Multiservices) on page 177

Copyright © 2013, Juniper Networks, Inc.

member-interface (Aggregated Multiservices)

Syntax member-interface *interface-name*;

Hierarchy Level [edit interfaces *interface-name* load-balancing-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the member interfaces for the aggregated Multiservices (AMS) interface. You can configure multiple interfaces by specifying each interface in a separate statement.

For high availability service applications like Network Address Translation (NAT) that support many-to-one (N:1) redundancy, you can specify two or more interfaces.



NOTE: **The member interfaces that you specify must be members of aggregated Multiservices interfaces (mams-).**

The remaining statements are explained separately.

Options ***interface-name***—Name of the member interface. The member interface format is **mams-a/b/0**, where **a** is the Flexible PIC Concentrator (FPC) slot number and **b** is the PIC slot number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- load-balancing-options (Aggregated Multiservices) on page 177

Copyright © 2013, Juniper Networks, Inc.

redistribute-all-traffic (Aggregated Multiservices)

Syntax

```
redistribute-all-traffic {
  enable-rejoin;
}
```

Hierarchy Level [edit interfaces *interface-name* load-balancing-options member-failure-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enable the option to redistribute traffic of a failed active member to the other active members.

For many-to-one (N:1) high availability support for Network Address Translation (NAT), the traffic for the failed member is automatically redistributed to the other active members.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- member-failure-options (Aggregated Multiservices) on page 179

rejoin-timeout (Aggregated Multiservices)

Syntax rejoin-timeout *rejoin-timeout*;

Hierarchy Level [edit interfaces *interface-name* load-balancing-options member-failure-options drop-member-traffic]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the time by when a failed member should rejoin the aggregated Multiservices (AMS) interface automatically. If the failed member does not rejoin by the configured time, then the member is moved to the “inactive” state and the traffic meant for this member is dropped.

Default If you do not configure a value, the default value of 120 seconds is used.

Options ***rejoin-timeout***—Time, in seconds, by which a failed member must rejoin.

Default: 120 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- drop-member-traffic (Aggregated Multiservices) on page 173

Copyright © 2013, Juniper Networks, Inc.

unit (Aggregated Multiservices)

Syntax

```
unit interface-unit-number {
  family family;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

The remaining statements are explained separately.

Options ***interface-unit-number***—Number of the logical unit.



NOTE: Unit 0 is reserved and cannot be configured under the aggregated Multiservices interface (ams).

Range: 1 through 16,384

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- interfaces on page 176

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 12

Intrusion Detection Service Configuration Guidelines

The Adaptive Services (AS) or Multiservices PIC supports a limited set of intrusion detection services (IDS) to perform attack detection. You can use IDS to perform the following tasks:

- Detect various types of denial-of-service (DoS) and directed denial-of-service (DDoS) attacks.
- Detect attempts at network scanning and probing.
- Detect anomalies in traffic patterns, such as sudden bursts or a decline in bandwidth.
- Prevent some types of attacks.
- Redirect attack traffic to a collector for analysis.
- Specify thresholds for limiting the number of flows, the packet rate, and the session rate.

IDS enables you to focus attack detection and remedial actions on specific hosts or networks that you specify in the IDS terms. Signature detection is not supported.

To configure IDS, include the **ids** statement at the **[edit services]** hierarchy level:

```
[edit services]
ids {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      rule {
        application-sets set-name;
        applications [ application-names ];
        destination-address (address | any-unicast) <except>;
        destination-address-range low minimum-value high maximum-value <except>;
        destination-prefix-list list-name <except>;
        source-address (address | any-unicast) <except>;
        source-address-range low minimum-value high maximum-value <except>;
        source-prefix-list list-name <except>;
      }
      then {
        aggregation {
```

Copyright © 2013, Juniper Networks, Inc.

```
            destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
            source-prefix prefix-value | source-prefix-ipv6 prefix-value;
          }
          (force-entry | ignore-entry);
          logging {
            syslog;
            threshold rate;
          }
          session-limit {
            by-destination {
              hold-time seconds;
              maximum number;
              packets number;
              rate number;
            }
            by-pair {
              hold-time seconds;
              maximum number;
              packets number;
              rate number;
            }
            by-source {
              hold-time seconds;
              maximum number;
              packets number;
              rate number;
            }
          }
          syn-cookie {
            mss value;
            threshold rate;
          }
        }
      }
    }
    rule-set rule-set-name {
      [ rule rule-names ];
    }
  }
```



NOTE: The Junos OS uses stateful firewall settings as a basis for performing IDS. You must commit a stateful firewall configuration in the same service set for IDS to function properly.

This chapter contains the following sections:

- Configuring IDS Rules on page 187
- Configuring IDS Rule Sets on page 193
- Examples: Configuring IDS Rules on page 193

Copyright © 2013, Juniper Networks, Inc.

Configuring IDS Rules

IDS rules identify traffic for which you want the router software to count events. Because IDS is based on stateful firewall properties, you must configure at least one stateful firewall rule and include it in the service set with the IDS rules; for more information, see "Configuring Stateful Firewall Rules" on page 132.

To configure an IDS rule, include the **rule *rule-name*** statement at the **[edit services ids]** hierarchy level:

```
[edit services ids]
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (address | any-unicast) <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      destination-prefix-list list-name <except>;
      source-address (address | any-unicast) <except>;
      source-address-range low minimum-value high maximum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      aggregation {
        destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
        source-prefix prefix-value | source-prefix-ipv6 prefix-value;
      }
      (force-entry | ignore-entry);
      logging {
        syslog;
        threshold rate;
      }
      session-limit {
        by-destination {
          hold-time seconds;
          maximum number;
          packets number;
          rate number;
        }
        by-pair {
          hold-time seconds;
          maximum number;
          packets number;
          rate number;
        }
        by-source {
          hold-time seconds;
          maximum number;
          packets number;
          rate number;
        }
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
        syn-cookie {
          mss value;
          threshold rate;
        }
      }
    }
  }
```

Each IDS rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections describe IDS rule content in more detail:

- Configuring Match Direction for IDS Rules on page 188
- Configuring Match Conditions in IDS Rules on page 189
- Configuring Actions in IDS Rules on page 190

Configuring Match Direction for IDS Rules

Each rule must include a **match-direction** statement that specifies whether the match is applied on the input or output side of the interface. To configure where the match is applied, include the **match-direction (input | input-output | output)** statement at the **[edit services ids rule *rule-name*]** hierarchy level:

```
[edit services ids rule rule-name]
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, bidirectional rule creation is allowed.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output. For more information on inside and outside interfaces, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

On the AS or Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that match the packet direction are considered.

Copyright © 2013, Juniper Networks, Inc.

Configuring Match Conditions in IDS Rules

To configure IDS match conditions, include the **from** statement at the **[edit services ids rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ids rule rule-name term term-name]
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (address | any-unicast) <except>;
  destination-address-range low minimum-value high maximum-value <except>;
  destination-prefix-list list-name <except>;
  source-address (address | any-unicast) <except>;
  source-address-range low minimum-value high maximum-value <except>;
  source-prefix-list list-name <except>;
}
```

If you omit the **from** statement, the software accepts all events and places them in the IDS cache for processing.

The source address and destination address can be either IPv4 or IPv6. You can use the destination address, a range of destination addresses, a source address, or a range of source addresses as a match condition, in the same way that you would configure a firewall filter; for more information, see the Routing Policy Configuration Guide.

Alternatively, you can specify a list of source or destination prefixes by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or **source-prefix-list** statement in the IDS rule. For an example, see "Examples: Configuring Stateful Firewall Rules" on page 136.

You can also include application protocol definitions that you have configured at the **[edit applications]** hierarchy level; for more information, see "Configuring Application Protocol Properties" on page 74.

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services ids rule *rule-name* term *term-name* from]** hierarchy level.
- To apply one or more sets of application protocol definitions that you have defined, include the **application-sets** statement at the **[edit services ids rule *rule-name* term *term-name* from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit applications]** hierarchy level; you cannot specify these properties as match conditions.

If a match occurs on an application, the application protocol is displayed separately in the **show services ids** command output. For more information, see the Junos OS Operational Mode Commands.

Copyright © 2013, Juniper Networks, Inc.

Configuring Actions in IDS Rules

To configure IDS actions, include the **then** statement at the **[edit services ids rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ids rule rule-name term term-name]
then {
  aggregation {
    destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
    source-prefix prefix-value | source-prefix-ipv6 prefix-value;
  }
  (force-entry | ignore-entry);
  logging {
    syslog;
    threshold rate;
  }
  session-limit {
    by-destination {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-pair {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-source {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
  }
  syn-cookie {
    mss value;
    threshold rate;
  }
}
```

You can configure the following possible actions:

- **aggregation**—The router aggregates traffic labeled with the specified source or destination prefixes before passing the events to IDS processing. This is helpful if you want to examine all the traffic connected with a particular source or destination host. To collect traffic with some other marker, such as a particular application or port, configure that value in the match conditions.

  To configure aggregation prefixes, include the **aggregation** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level and specify values for **source-prefix**, **destination-prefix source-prefix-ipv6**, or **destination-prefix-ipv6**:

  ```
  [edit services ids rule rule-name term term-name then]
  ```


Copyright © 2013, Juniper Networks, Inc.

```
aggregation {
destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
source-prefix prefix-value | source-prefix-ipv6 prefix-value;
}
```

The value of **source-prefix** and **destination-prefix** must be an integer between 1 and 32. The value of **source-prefix-ipv6** and **destination-prefix-ipv6** must be an integer between 1 and 128.

- **(force-entry | ignore-entry)**—**force-entry** provides a permanent spot in IDS caches for subsequent events after one event is registered. By default, the IDS software does not record information about "good" packets that do not exhibit suspicious behavior. You can use the **force-entry** statement to record all traffic from a suspect host, even traffic that would not otherwise be counted.

  ignore-entry ensures that all IDS events are ignored. You can use this statement to disregard all traffic from a host you trust, including any temporary anomalies that IDS would otherwise count as events.

  To configure an entry behavior different from the default, include the **force-entry** or **ignore-entry** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level:

  ```
  [edit services ids rule rule-name term term-name then]
  (force-entry | ignore-entry);
  ```

- **logging**—The event is logged in the system log file.

  To configure logging, include the **logging** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level:

  ```
  [edit services ids rule rule-name term term-name then]
  logging {
    syslog;
    threshold rate;
  }
  ```

  You can optionally include a threshold rate to trigger the generation of system log messages. The threshold rate is specified in events per second. IDS logs are generated once every 60 seconds for each anomaly that is reported. The logs are generated as long as the events continue.

- **session-limit**—The router limits open sessions when the specified threshold is reached.

  To configure a threshold, include the **session-limit** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level:

  ```
  [edit services ids rule rule-name term term-name then]
  session-limit {
    by-destination {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-pair {
      hold-time seconds;
      maximum number;
  ```

Copyright © 2013, Juniper Networks, Inc.

```
    packets number;
    rate number;
  }
  by-source {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
}
```

You configure the thresholds for flow limitation based on traffic direction:

- To limit the number of outgoing sessions from one internal host or subnet, configure the **by-source** statement.
- To limit the number of sessions between a pair of IP addresses, subnets, or applications, configure the **by-pair** statement.
- To limit the number of incoming sessions to one external public IP address or subnet, configure the **by-destination** statement.

For each direction, you can configure the following threshold values:

- **hold-time** *seconds*—When the **rate** or **packets** measurement reaches the threshold value, stop all new flows for the specified number of seconds. Once **hold-time** is in effect, the traffic is blocked for the specified time even if the rate subsides below the specified limit. By default, **hold-time** has a value of 0; the range is 0 through 60 seconds.
- **maximum** *number*—Maximum number of open sessions per IP address or subnet per application. The range is 1 through 32,767.
- **packets** *number*—Maximum number of packets per second (pps) per IP address or subnet per application. The range is 4 through 2,147,483,647.
- **rate** *number*—Maximum number of sessions per second per IP address or subnet per application. The range is 4 through 32,767.

  If you include more than one source address in the match conditions configured at the **[edit services ids rule** *rule-name* **term** *term-name* **from]** hierarchy level, limits are applied for each source address independently. For example, the following configuration allows 20 connections from each source address (**10.1.1.1** and **10.1.1.2**), not 20 connections total. The same logic applies to the **applications** and **destination-address** match conditions.

```
[edit services ids rule rule-name term term-name]
  from {
  source-address 10.1.1.1;
  source-address 10.1.1.2;
  }
  then {
  session-limit by-source {
    maximum 20;
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

NOTE: IDS limits are applied to packets that are accepted by stateful firewall rules. They are not applied to packets discarded or rejected by stateful firewall rules. For example, if the stateful firewall accepts 75 percent of the incoming traffic and the remaining 25 percent is rejected or discarded, the IDS limit applies only to 75 percent of the traffic.

- **syn-cookie**—The router activates SYN-cookie defensive mechanisms.

  To configure SYN-cookie values, include the **syn-cookie** statement at the **[edit services ids rule *rule-name* term *term-name* then]** hierarchy level:

  ```
  [edit services ids rule rule-name term term-name then]
  syn-cookie {
    mss value;
    threshold rate;
  }
  ```

  If you enable SYN-cookie defenses, you must include both a threshold rate to trigger SYN-cookie activity and a Transmission Control Protocol (TCP) maximum segment size (MSS) value for TCP delayed binding. The threshold rate is specified in SYN attacks per second. By default, the TCP MSS value is 1500; the range is from 128 through 8192.

Configuring IDS Rule Sets

The **rule-set** statement defines a collection of IDS rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services ids]** hierarchy level with a **rule** statement for each rule:

```
[edit services ids]
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Examples: Configuring IDS Rules

The following configuration adds a permanent entry to the IDS anomaly table when it encounters a flow with the destination address 10.410.6.2:

```
[edit services ids]
rule simple_ids {
  term 1 {
    from {
      destination-address 10.410.6.2/32;
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
        then {
            force-entry;
            logging {
                threshold 1;
                syslog;
            }
        }
    }
    term default {
        then {
            aggregation {
                source-prefix 24;
            }
        }
    }
    match-direction input;
}
```

The IDS configuration works in conjunction with the stateful firewall mechanism and relies heavily on the anomalies reported by the stateful firewall. The following configuration example shows this relationship:

```
[edit services ids]
rule simple_ids {
    term 1 {
        from {
            source-address 10.30.20.2/32;
            destination-address {
                10.30.10.2/32;
                10.30.1.2/32 except;
            }
            applications appl-ftp;
        }
        then {
            force-entry;
            logging {
                threshold 5;
                syslog;
            }
            syn-cookie {
                threshold 10;
            }
        }
    }
    match-direction input;
}
```

The following example shows configuration of flow limits:

```
[edit services ids]
rule ids-all {
    match-direction input;
    term t1 {
        from {
            application-sets alg-set;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
        then {
            aggregation {
                destination-prefix 30; /* IDS action aggregation */
            }
            logging {
                threshold 10;
            }
            session-limit {
                by-destination {
                    hold-time 0;
                    maximum 10;
                    packets 200;
                    rate 100;
                }
                by-pair {
                    hold-time 0;
                    maximum 10;
                    packets 200;
                    rate 100;
                }
                by-source {
                    hold-time 5;
                    maximum 10;
                    packets 200;
                    rate 100;
                }
            }
        }
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 13

Summary of Intrusion Detection Service Configuration Statements

The following sections explain each of the intrusion detection service (IDS) statements. The statements are organized alphabetically.

aggregation

Syntax

```
aggregation {
    destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
    source-prefix prefix-value | source-prefix-ipv6 prefix-value;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the type of data to be aggregated.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring IDS Rules" on page 187.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

application-sets (Services IDS)

Syntax application-sets *set-name*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define one or more target application sets.

Options ***set-name***—Name of the target application set.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

applications (Services IDS)

Syntax applications [*application-names*];

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define one or more applications to which IDS applies.

Options ***application-name***—Name of the target application.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

by-destination

Syntax

```
by-destination {
  hold-time seconds;
  maximum number;
  packets number;
  rate number;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then session-limit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply limit to sessions based on numbers generated from the configured destination (IP or subnet) or application.

Options

hold-time ***seconds***—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the **maximum**, **packets**, or **rate** statements.

maximum ***number***—Maximum number of open sessions per application or IP address.

packets ***number***—Maximum peak packets per second per application or IP address.

rate ***number***—Maximum number of sessions per second per application or IP address.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

by-pair

Syntax

```
by-pair {
  hold-time seconds;
  maximum number;
  packets number;
  rate number;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then session-limit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply limit to paired stateful firewall and NAT flows (forward and reverse).

Options

hold-time ***seconds***—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the **maximum**, **packets**, or **rate** statements.

maximum ***number***—Maximum number of open sessions per application or IP address.

packets ***number***—Maximum peak packets per second per application or IP address.

rate ***number***—Maximum number of sessions per second per application or IP address.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

by-source

Syntax

```
by-source {
  hold-time seconds;
  maximum number;
  packets number;
  rate number;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then session-limit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply limit to sessions based on numbers generated from the configured source (IP or subnet) or application.

Options

hold-time ***seconds***—Length of time for which to stop all new flows once the rate of events exceeds the threshold set by one or more of the **maximum**, **packets**, or **rate** statements.

maximum ***number***—Maximum number of open sessions per application or IP address.

packets ***number***—Maximum peak packets per second per application or IP address.

rate ***number***—Maximum number of sessions per second per application or IP address.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination-address (Services IDS)

Syntax destination-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the destination address for rule matching.

Options ***address***—Destination IPv4 or IPv6 address or prefix value.

any-unicast—Any unicast packet.

except—(Optional) Exempt the specified address, prefix, or unicast packets from rule matching.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

destination-address-range (Services IDS)

Syntax destination-address-range low *minimum-value* high *maximum-value* <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 7.6.
minimum-value and ***maximum-value*** options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the destination address range for rule matching.

Options ***minimum-value***—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exempt the specified address range from rule matching.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

destination-prefix (Services IDS)

Syntax destination-prefix *prefix-value*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then aggregation]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the prefix value for destination IPv4 address aggregation.

Options ***prefix-value***—Integer value.
Range: 1 through 32

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

destination-prefix-ipv6

Syntax destination-prefix-ipv6 *prefix*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then aggregation]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the prefix value for destination IPv6 address aggregation.

Options ***prefix-value***—Integer value.
Range: 1 through 128

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination-prefix-list (Services IDS)

Syntax destination-prefix-list *list-name* <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the destination prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Routing Policy Configuration Guide

force-entry

Syntax (force-entry | ignore-entry);

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify handling of entries in the IDS events cache:

- **force-entry**—Ensure that the entry has a permanent place in the IDS cache after one event is registered.
- **ignore-entry**—Ensure that all IDS events are ignored.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

from (Services IDS)

Syntax

```
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (address | any-unicast) <except>;
  destination-address-range low minimum-value high maximum-value <except>;
  source-address (address | any-unicast) <except>;
  source-address-range low minimum-value high maximum-value <except>;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify input conditions for the IDS term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

ignore-entry

See force-entry

logging (Services IDS)

Syntax

```
logging {
  syslog;
  threshold rate;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set logging values for this IDS term.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

match-direction (Services IDS)

Syntax match-direction (input | output | input-output);

Hierarchy Level [edit services ids rule *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the direction in which the rule match is applied.

Options
input—Apply the rule match on input.

output—Apply the rule match on output.

input-output—Apply the rule match bidirectionally.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

mss

Syntax mss *value*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then syn-cookie]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the maximum segment size (MSS) value used in Transmission Control Protocol (TCP) delayed binding.

Options
value—MSS value.
Default: 1500
Range: 128 through 8192

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rule (Services IDS)

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (address | any-unicast) <except>;
      destination-address-range low minimum-value high maximum-value <except>;
      source-address (address | any-unicast) <except>;
      source-address-range low minimum-value high maximum-value <except>;
    }
    then {
      aggregation {
        destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
        source-prefix prefix-value | source-prefix-ipv6 prefix-value;
      }
      (force-entry | ignore-entry);
      logging {
        syslog;
        threshold rate;
      }
      session-limit {
        by-destination {
          hold-time seconds;
          maximum number;
          packets number;
          rate number;
        }
        by-pair {
          hold-time seconds;
          maximum number;
          packets number;
          rate number;
        }
        by-source {
          hold-time seconds;
          maximum number;
          packets number;
          rate number;
        }
      }
      syn-cookie {
        mss value;
        threshold rate;
      }
    }
  }
}
```

Hierarchy Level

[edit services ids],
[edit services ids rule-set *rule-set-name*]

Copyright © 2013, Juniper Networks, Inc.

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule the router uses when applying this service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

Usage Guidelines See "Configuring IDS Rules" on page 187.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

rule-set (Services IDS)

Syntax
```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services ids]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

Usage Guidelines See "Configuring IDS Rule Sets" on page 193.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

services (IDS)

Syntax services ids { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service rules to be applied to traffic.

Options **ids**—Identifies the IDS set of rules statements.

Usage Guidelines See "Configuring IDS Rules" on page 187.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

session-limit

Syntax

```
session-limit {
  by-destination {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
  by-pair {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
  by-source {
    hold-time seconds;
    maximum number;
    packets number;
    rate number;
  }
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable flow limitation by configuring thresholds on source, destination, or stateful firewall and network address translation (NAT) paired traffic flows.

Options The remaining statements are described separately.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

source-address (Services IDS)

Syntax source-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the source address for rule matching.

Options ***address***—Source IPv4 or IPv6 address or prefix value.

any-unicast—Any unicast packet.

except—(Optional) Exempt the specified address, prefix, or unicast packets from rule matching.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

source-address-range (Services IDS)

Syntax source-address-range low *minimum-value* high *maximum-value* <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 7.6.
minimum-value and ***maximum-value*** options enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the source address range for rule matching.

Options ***minimum-value***—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exempt the specified address range from rule matching.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

source-prefix (Services IDS)

Syntax source-prefix *prefix-value*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then aggregation]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the prefix value for source IPv4 address aggregation.

Options ***prefix-value***—Integer value.
Range: 1 through 32

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

source-prefix-ipv6

Syntax source-prefix-ipv6 *prefix-value*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then aggregation]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the prefix value for source IPv6 address aggregation.

Options ***prefix-value***—Integer value.
Range: 1 through 128

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

source-prefix-list (Services IDS)

Syntax source-prefix-list *list-name* <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the source prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 189.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Routing Policy Configuration Guide

syn-cookie

Syntax
```
syn-cookie {
  mss value;
  threshold rate;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable SYN-cookie defenses against SYN attacks. By default, SYN-cookie techniques are not applied.

Options The remaining statements are described separately.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

syslog (Services IDS)

Syntax syslog;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then logging]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable system logging. The system log information from the Adaptive Services or Multiservices Physical Interface Card (PIC) is passed to the kernel for logging in the **/var/log** directory.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

term (Services IDS)

Syntax

```
term term-name {
  from {
    application-sets set-name;
    applications [ application-names ];
    destination-address (address | any-unicast) <except>;
    destination-address-range low minimum-value high maximum-value <except>;
    source-address (address | any-unicast) <except>;
    source-address-range low minimum-value high maximum-value <except>;
  }
  then {
    aggregation {
      destination-prefix prefix-value | destination-prefix-ipv6 prefix-value;
      source-prefix prefix-value | source-prefix-ipv6 prefix-value;
    }
    (force-entry | ignore-entry);
    logging {
      syslog;
      threshold rate;
    }
    session-limit {
      by-destination {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
      }
      by-pair {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
      }
      by-source {
        hold-time seconds;
        maximum number;
        packets number;
        rate number;
      }
    }
    syn-cookie {
      mss value;
      threshold rate;
    }
  }
}
```

Hierarchy Level [edit services ids rule *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IDS term properties.

Copyright © 2013, Juniper Networks, Inc.

Options

term-name—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring IDS Rules" on page 187.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

then (Services IDS)

Syntax

```
then {
  aggregation {
    destination-prefix prefix-number | destination-prefix-ipv6 prefix-value;
    source-prefix prefix-number | source-prefix-ipv6 prefix-value;
  }
  (force-entry | ignore-entry);
  logging {
    syslog;
    threshold rate;
  }
  session-limit {
    by-destination {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-pair {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
    by-source {
      hold-time seconds;
      maximum number;
      packets number;
      rate number;
    }
  }
  syn-cookie {
    mss value;
    threshold rate;
  }
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IDS term actions.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring IDS Rules" on page 187.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

threshold (Services)

Syntax threshold *rate*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then logging],
[edit services ids rule *rule-name* term *term-name* then syn-cookie]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the threshold for logging or applying SYN-cookie defenses.

Options ***rate***—Logging threshold number of events per second.

rate—SYN-cookie defense number of SYN attacks per second.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 190.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 14

IPsec Services Configuration Guidelines

To configure IP Security (IPsec) services, include the following statements at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
clear-ike-sas-on-pic-restart;
clear-ipsec-sas-on-pic-restart;
ike {
  proposal proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256);
    authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2 | group5 | group14);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
  }
  policy policy-name {
    description description;
    local-certificate identifier;
    local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
    version (1 | 2);
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
    remote-id {
      any-remote-id;
      ipv4_addr [ values ];
      ipv6_addr [ values ];
      key_id [ values ];
    }
  }
}
ipsec {
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    description description;
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
  }
  policy policy-name {
    description description;
    perfect-forward-secrecy {
```

Copyright © 2013, Juniper Networks, Inc.

```
        keys (group1 | group2);
      }
      proposals [ proposal-names ];
    }
  }
  rule rule-name {
    match-direction (input | output);
    term term-name {
      from {
        destination-address address;
        ipsec-inside-interface interface-name;
        source-address address;
      }
      then {
        anti-replay-window-size bits;
        backup-remote-gateway address;
        clear-dont-fragment-bit;
        dynamic {
          ike-policy policy-name;
          ipsec-policy policy-name;
        }
        initiate-dead-peer-detection;
        manual {
          direction (inbound | outbound | bidirectional) {
            authentication {
              algorithm (hmac-md5-96 | hmac-sha1-96);
              key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi spi-value;
            encryption {
              algorithm algorithm;
              key (ascii-text key | hexadecimal key);
            }
            protocol (ah | bundle | esp);
            spi spi-value;
          }
        }
        no-anti-replay;
        remote-gateway address;
        syslog;
        tunnel-mtu bytes;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
  no-ipsec-tunnel-in-traceroute;
  traceoptions {
    file {
      files number;
      size bytes;
    }
    flag flag;
    level level;
  }
```


Copyright © 2013, Juniper Networks, Inc.

This chapter includes the following sections:

- Minimum Security Association Configurations on page 221
- Configuring Security Associations on page 222
- Configuring IKE Proposals on page 228
- Configuring IKE Policies on page 232
- Configuring IPsec Proposals on page 237
- Configuring IPsec Policies on page 240
- IPsec Policy for Dynamic Endpoints on page 242
- Configuring IPsec Rules on page 243
- Configuring IPsec Rule Sets on page 249
- Configuring Dynamic Endpoints for IPsec Tunnels on page 250
- Tracing IPsec Operations on page 255
- Configuring IPSec on the Services SDK on page 257
- Examples: Configuring IPsec Services on page 258

Minimum Security Association Configurations

The following sections show the minimum configurations necessary to set up security associations (SAs) for IPsec services:

- Minimum Manual SA Configuration on page 221
- Minimum Dynamic SA Configuration on page 221

Minimum Manual SA Configuration

To define a manual SA configuration, you must include at least the following statements at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction (inbound | outbound | bidirectional) {
  authentication {
    algorithm (hmac-md5-96 | hmac-sha1-96);
    key (ascii-text key | hexadecimal key);
  }
  encryption {
    algorithm algorithm;
    key (ascii-text key | hexadecimal key);
  }
  protocol (ah | esp | bundle);
  spi spi-value;
}
```

Minimum Dynamic SA Configuration

To define a dynamic SA configuration, you must include at least the following statements at the **[edit services ipsec-vpn]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
[edit services ipsec-vpn]
ike {
  proposal proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256);
    authentication-method pre-shared-keys;
    dh-group (group1 | group2 | group5 |group14);
    encryption-algorithm algorithm;
  }
  policy policy-name {
    proposals [ ike-proposal-names ];
    pre-shared-key (ascii-text key | hexadecimal key);
    version (1 | 2);
    mode (aggressive | main);
  }
}
ipsec {
  policy policy-name {
    proposals [ ipsec-proposal-names ];
  }
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    encryption-algorithm algorithm;
    protocol (ah | esp | bundle);
  }
}
```

NOTE:

- Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. The **version** statement under the **[edit services ipsec-vpn ike policy *name*]** hierarchy allows you to configure the specific IKE version to be supported.
- The **mode** statement under the **[edit services ipsec-vpn ike policy *name*]** hierarchy is required only if the **version** option is set to 1.

You must also include the **ipsec-policy** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then dynamic]** hierarchy level.

Configuring Security Associations

To use IPsec services, you create an SA between hosts. An SA is a simplex connection that allows two hosts to communicate with each other securely by means of IPsec. You can configure two types of SAs:

- Manual—Requires no negotiation; all values, including the keys, are static and specified in the configuration. As a result, each peer must have the same configured options for communication to take place. For information about how to configure a manual SA, see "Configuring Manual Security Associations" on page 223.
- Dynamic—Specifies proposals to be negotiated with the tunnel peer. The keys are generated as part of the negotiation and therefore do not need to be specified in the configuration. The dynamic SA includes one or more **proposal** statements, which allow

 Copyright © 2013, Juniper Networks, Inc.

you to prioritize a list of protocols and algorithms to be negotiated with the peer. For information about how to configure a dynamic SA, see "Configuring Dynamic Security Associations" on page 227.

This section includes the following topics:

- Configuring Manual Security Associations on page 223
- Configuring Dynamic Security Associations on page 227
- Clearing Security Associations on page 228



NOTE: Both OSPFv2 and OSPFv3 support IPsec authentication. However, dynamic or tunnel mode IPsec SAs are not supported for OSPFv3. If you add SAs into OSPFv3 by including the **ipsec-sa** statement at the **[edit protocols ospf3 area *area-number* interface *interface-name*]** hierarchy level, your configuration fails to commit. For more information about OSPF authentication and other OSPF properties, see the Junos OS Routing Protocols Configuration Guide.

Configuring Manual Security Associations

Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. As a result, each peer must have the same configured options for communication to take place.

To configure a manual IPsec security association, include statements at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction (inbound | outbound | bidirectional) {
  authentication {
    algorithm (hmac-md5-96 | hmac-sha1-96);
    key (ascii-text key | hexadecimal key);
  }
  auxiliary-spi auxiliary-spi-value;
  encryption {
    algorithm algorithm;
    key (ascii-text key | hexadecimal key);
  }
  protocol (ah | esp | bundle);
  spi spi-value;
}
```

To configure manual SA statements, do the following:

- Configuring the Direction for IPsec Processing on page 224
- Configuring the Protocol for a Manual IPsec SA on page 225
- Configuring the Security Parameter Index on page 225
- Configuring the Auxiliary Security Parameter Index on page 225

Copyright © 2013, Juniper Networks, Inc.

- Configuring Authentication for a Manual IPsec SA on page 225
- Configuring Encryption for a Manual IPsec SA on page 226

Configuring the Direction for IPsec Processing

The **direction** statement specifies inbound or outbound IPsec processing. If you want to define different algorithms, keys, or security parameter index (SPI) values for each direction, you configure the **inbound** and **outbound** options. If you want the same attributes in both directions, use the **bidirectional** option.

To configure the direction of IPsec processing, include the **direction** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
   direction (inbound | outbound | bidirectional) {
   ...
}
```

Example: Using Different Configuration for the Inbound and Outbound Directions

Define different algorithms, keys, and security parameter index values for each direction:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction inbound {
   protocol esp;
   spi 16384;
   encryption {
      algorithm 3des-cbc;
      key ascii-text 2345678901234567890123 4;
   }
}
direction outbound {
   protocol esp;
   spi 24576;
   encryption {
      algorithm 3des-cbc;
      key ascii-text 12345678901234567890abcd;
   }
}
```

Example: Using the Same Configuration for the Inbound and Outbound Directions

Define one set of algorithms, keys, and security parameter index values that is valid in both directions:

```
[edit services ipsec-vpn rule rule-name term term-name then manual]
direction bidirectional {
   protocol ah;
   spi 20001;
   authentication {
      algorithm hmac-md5-96;
      key ascii-text 123456789012abcd;
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

Configuring the Protocol for a Manual IPsec SA

IPsec uses two protocols to protect IP traffic: Encapsulating Security Payload (ESP) and authentication header (AH). The AH protocol is used for strong authentication. A third option, **bundle**, uses AH authentication and ESP encryption; it does not use ESP authentication because AH provides stronger authentication of IP packets.

To configure the IPsec protocol, include the **protocol** statement and specify the **ah**, **esp**, or **bundle** option at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
protocol (ah | bundle | esp);
```

Configuring the Security Parameter Index

An SPI is an arbitrary value that uniquely identifies which SA to use at the receiving host. The sending host uses the SPI to identify and select which SA to use to secure every packet. The receiving host uses the SPI to identify and select the encryption algorithm and key used to decrypt packets.



NOTE: Each manual SA must have a unique SPI and protocol combination. Use the auxiliary SPI when you configure the **protocol** statement to use the **bundle** option.

To configure the SPI, include the **spi** statement and specify a value (from 256 through 16,639) at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
spi spi-value;
```

Configuring the Auxiliary Security Parameter Index

Use the auxiliary SPI when you configure the **protocol** statement to use the **bundle** option.



NOTE: Each manual SA must have a unique SPI and protocol combination.

To configure the auxiliary SPI, include the **auxiliary-spi** statement and specify a value (from 256 through 16,639) at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
auxiliary-spi auxiliary-spi-value;
```

Configuring Authentication for a Manual IPsec SA

To configure an authentication algorithm, include the **authentication** statement and specify an authentication algorithm and a key at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
  authentication {
  algorithm (hmac-md5-96 | hmac-sha1-96);
  key (ascii-text key | hexadecimal key);
}
```

The algorithm can be one of the following:

- **hmac-md5-96**—Hash algorithm that authenticates packet data. It produces a 128-bit authenticator value and a 96-bit digest.
- **hmac-sha1-96**—Hash algorithm that authenticates packet data. It produces a 160-bit authenticator value and a 96-bit digest.

The key can be one of the following:

- **ascii-text**—ASCII text key. With the **hmac-md5-96** option, the key contains 16 ASCII characters. With the **hmac-sha1-96** option, the key contains 20 ASCII characters.
- **hexadecimal**—Hexadecimal key. With the **hmac-md5-96** option, the key contains 32 hexadecimal characters. With the **hmac-sha1-96** option, the key contains 40 hexadecimal characters.

Configuring Encryption for a Manual IPsec SA

To configure IPsec encryption, include the **encryption** statement and specify an algorithm and key at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]
  encryption {
  algorithm algorithm;
  key (ascii-text key | hexadecimal key);
}
```

The algorithm can be one of the following:

- **des-cbc**—Encryption algorithm that has a block size of 8 bytes; its key size is 64 bits long.
- **3des-cbc**—Encryption algorithm that has a block size of 24 bytes; its key size is 192 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

Copyright © 2013, Juniper Networks, Inc.

NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semiweak keys, see RFC 2409, *The Internet Key Exchange (IKE)*. The AES encryption algorithms use a software implementation that has much lower throughput, so DES remains the recommended option. For reference information on AES encryption, see RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*.

For **3des-cbc**, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

If you configure an authentication proposal but do not include the **encryption** statement, the result is NULL encryption. Certain applications expect this result. If you configure no specific authentication or encryption values, the Junos OS uses the default values of **sha1** for the authentication and **3des-cbc** for the encryption.

The key can be one of the following:

- **ascii-text**—ASCII text key. With the **des-cbc** option, the key contains 8 ASCII characters. With the **3des-cbc** option, the key contains 24 ASCII characters.
- **hexadecimal**—Hexadecimal key. With the **des-cbc** option, the key contains 16 hexadecimal characters. With the **3des-cbc** option, the key contains 48 hexadecimal characters.

NOTE: You cannot configure encryption when you use the AH protocol.

Configuring Dynamic Security Associations

You configure dynamic SAs with a set of proposals that are negotiated by the security gateways. The keys are generated as part of the negotiation and therefore do not need to be specified in the configuration. The dynamic SA includes one or more proposals, which allow you to prioritize a list of protocols and algorithms to be negotiated with the peer.

To enable a dynamic SA, follow these steps:

1. Configure Internet Key Exchange (IKE) proposals and IKE policies associated with these proposals.
2. Configure IPsec proposals and an IPsec policy associated with these proposals.
3. Associate an SA with an IPsec policy by configuring the **dynamic** statement.

For more information about IKE policies and proposals, see "Configuring IKE Policies" on page 232 and "Configuring IKE Proposals" on page 228. For more information about IPsec policies and proposals, see "Configuring IPsec Policies" on page 240.

Copyright © 2013, Juniper Networks, Inc.

To configure a dynamic SA, include the **dynamic** statement and specify an IPsec policy name at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level. The **ike-policy** statement is optional unless you use the preshared key authentication method.

```
[edit services ipsec-vpn rule rule-name term term-name then]
dynamic {
   ike-policy policy-name;
   ipsec-policy policy-name;
}
```



NOTE: If you want to establish a dynamic SA, the attributes in at least one configured IPsec and IKE proposal must match those of its peer.

Clearing Security Associations

You can set up the router software to clear IKE or IPsec SAs automatically when the corresponding services PIC restarts or is taken offline. To configure this property, include the **clear-ike-sas-on-pic-restart** or **clear-ipsec-sas-on-pic-restart** statement at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
clear-ike-sas-on-pic-restart;
clear-ipsec-sas-on-pic-restart;
```

After you add this statement to the configuration, all the IKE or IPsec SAs corresponding to the tunnels in the PIC will be cleared when the PIC restarts or goes offline.

Configuring IKE Proposals

Dynamic security associations (SAs) require IKE configuration. With dynamic SAs, you configure IKE first, and then the SA. IKE creates the dynamic SAs and negotiates them for IPsec. The IKE configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway.

You can configure one or more IKE proposals. Each proposal is a list of IKE attributes to protect the IKE connection between the IKE host and its peer.



NOTE: In IKEv1, the authentication method for SAs is negotiated with the remote peer based on the type of authentication method configured in the IKE proposal. In IKEv2, such a negotiation is not performed with the remote peer. Instead, each IKE peer uses the authentication method that is locally configured for them.

For SAs in IKEv2, the authentication method is either the default value as IKEv1 (if another authentication method is not configured in the IKE proposal) or all IKEv2 proposals in the IKE policy must be configured with the same authentication method value.

Copyright © 2013, Juniper Networks, Inc.

To configure an IKE proposal, include the **proposal** statement and specify a name at the **[edit services ipsec-vpn ike]** hierarchy level:

```
[edit services ipsec-vpn ike]
proposal proposal-name {
  authentication-algorithm (md5 | sha1 | sha-256);
  authentication-method (dsa-signatures | pre-shared-key | rsa-signatures);
  dh-group (group1 | group2 | group5 |group14);
  encryption-algorithm algorithm;
  lifetime-seconds seconds;
}
```

This section includes the following topics:

- Configuring the Authentication Algorithm for an IKE Proposal on page 229
- Configuring the Authentication Method for an IKE Proposal on page 229
- Configuring the Diffie-Hellman Group for an IKE Proposal on page 230
- Configuring the Encryption Algorithm for an IKE Proposal on page 230
- Configuring the Lifetime for an IKE SA on page 231
- Example: Configuring an IKE Proposal on page 232

Configuring the Authentication Algorithm for an IKE Proposal

To configure the authentication algorithm for an IKE proposal, include the **authentication-algorithm** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
authentication-algorithm (md5 | sha1 | sha-256);
```

The authentication algorithm can be one of the following:

- **md5**—Produces a 128-bit digest.
- **sha1**—Produces a 160-bit digest.
- **sha-256**—Produces a 256-bit digest.



NOTE: For reference information on Secure Hash Algorithms (SHAs), see Internet draft **draft-eastlake-sha2-02.txt**, *Secure Hash Algorithms (SHA and HMAC-SHA)* (expires July 2006).

Configuring the Authentication Method for an IKE Proposal

To configure the authentication method for an IKE proposal, include the **authentication-method** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
```

The authentication method can be one of the following:

Copyright © 2013, Juniper Networks, Inc.

- **dsa-signatures**—Digital Signature Algorithm
- **pre-shared-keys**—A key derived from an out-of-band mechanism; the key authenticates the exchanges
- **rsa-signatures**—Public key algorithm (supports encryption and digital signatures)

Configuring the Diffie-Hellman Group for an IKE Proposal

Diffie-Hellman is a public-key cryptography scheme that allows two parties to establish a shared secret over an insecure communications channel. It is also used within IKE to establish session keys.

To configure the Diffie-Hellman group for an IKE proposal, include the **dh-group** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
dh-group (group1 | group2 | group5 |group14);
```

The group can be one of the following:

- **group1**—Specifies that IKE use the 768-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group2**—Specifies that IKE use the 1024-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group5**—Specifies that IKE use the 1536-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group14**—Specifies that IKE use the 2048-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.

Using a Diffie-Hellman group based on a greater number of bits results a more secure IKE tunnel than using a group based on fewer bits. However, this additional security entails additional processing time.

Configuring the Encryption Algorithm for an IKE Proposal

To configure the encryption algorithm for an IKE proposal, include the **encryption-algorithm** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
encryption-algorithm algorithm;
```

The encryption algorithm can be one of the following:

- **3des-cbc**—Cipher block chaining encryption algorithm with a key size of 24 bytes; its key size is 192 bits long.
- **des-cbc**—Cipher block chaining encryption algorithm with a key size of 8 bytes; its key size is 56 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.

Copyright © 2013, Juniper Networks, Inc.

- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semiweak keys, see RFC 2409, *The Internet Key Exchange (IKE)*. The AES encryption algorithms use a software implementation that has much lower throughput, so DES remains the recommended option.

For **3des-cbc**, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

If you configure an authentication proposal but do not include the **encryption** statement, the result is NULL encryption. Certain applications expect this result. If you configure no specific authentication or encryption values, the Junos OS uses the default values of **sha1** for the authentication and **3des-cbc** for the encryption.

Configuring the Lifetime for an IKE SA

The **lifetime-seconds** statement sets the lifetime of an IKE SA. When the IKE SA expires, it is replaced by a new SA (and SPI) or the IPsec connection is terminated.

To configure the lifetime for an IKE SA, include the **lifetime-seconds** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ike proposal proposal-name]
lifetime-seconds seconds;
```

By default, the IKE SA lifetime is 3600 seconds. The range is from 180 through 86,400 seconds.

NOTE: In IKEv1, the lifetime for SAs is negotiated with the remote peer based on the type of lifetime configured in the IKE proposal. In IKEv2, such a negotiation is not performed with the remote peer. Instead, each IKE peer uses the lifetime that is locally configured for them.

For SAs in IKEv2, the lifetime is either the default value as IKEv1 (if another lifetime is not configured in the IKE proposal) or all IKEv2 proposals in the IKE policy must be configured with the same lifetime value.

NOTE: For IKE proposals, there is only one SA lifetime value, specified by the Junos OS. IPsec proposals use a different mechanism; for more information, see "Configuring the Lifetime for an IPsec SA" on page 239.

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring an IKE Proposal

Configure an IKE proposal:

```
[edit services ipsec-vpn ike]
proposal ike-proposal {
   authentication-method pre-shared-keys;
   dh-group group1;
   authentication-algorithm sha1;
   encryption-algorithm 3des-cbc;
}
```

Configuring IKE Policies

An IKE policy defines a combination of security parameters (IKE proposals) to be used during IKE negotiation. It defines a peer address and the proposals needed for that connection. Depending on which authentication method is used, it defines the preshared key for the given peer or the local certificate. During the IKE negotiation, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

A match is made when both policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used. The configured preshared key must also match its peer.

Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. You can configure the specific IKE phase to be supported for the negotiation. However, if only IKEv1 is supported, the Junos OS rejects IKEv2 negotiations. Similarly, if only IKEv2 is supported, the Junos OS rejects all IKEv1 negotiations.

The key management process (kmd) daemon determines which version of IKE is used in a negotiation. If kmd is the IKE initiator, it uses IKEv1 by default and retains the configured version for negotiations. If kmd is the IKE responder, it accepts connections from both IKEv1 and IKEv2.

You can create multiple, prioritized proposals at each peer to ensure that at least one proposal matches a remote peer's proposal.

First, you configure one or more IKE proposals; then you associate these proposals with an IKE policy. You can also prioritize a list of proposals used by IKE in the **policy** statement by listing the proposals you want to use, from first to last.

To configure an IKE policy, include the **policy** statement and specify a policy name at the **[edit services ipsec-vpn ike]** hierarchy level:

```
[edit services ipsec-vpn ike]
policy policy-name {
   description description;
   local-certificate identifier;
   local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
```


Copyright © 2013, Juniper Networks, Inc.

```
    version (1 | 2);
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
    remote-id {
      any-remote-id;
      ipv4_addr [ values ];
      ipv6_addr [ values ];
      key_id [ values ];
    }
  }
```

This section includes the following topics:

- Configuring the IKE Phase on page 233
- Configuring the Mode for an IKE Policy on page 233
- Configuring the Proposals in an IKE Policy on page 234
- Configuring the Preshared Key for an IKE Policy on page 234
- Configuring the Local Certificate for an IKE Policy on page 234
- Configuring the Description for an IKE Policy on page 236
- Configuring Local and Remote IDs for IKE Phase 1 Negotiation on page 236
- Example: Configuring an IKE Policy on page 236

For an example of an IKE policy configuration, see "Example: Configuring an IKE Policy" on page 236.

Configuring the IKE Phase

Starting with Junos OS Release 11.4, both IKEv1 and IKEv2 are supported by default on all M Series, MX Series, and T Series routers. You can configure the specific IKE phase to be supported for the negotiation. However, if only IKEv1 is supported, the Junos OS rejects IKEv2 negotiations. Similarly, if only IKEv2 is supported, the Junos OS rejects all IKEv1 negotiations.

To configure the IKE phase used, include the **version** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
version (1 | 2);
```

Configuring the Mode for an IKE Policy

IKE policy has two modes: aggressive and main. By default, main mode is enabled. Main mode uses six messages, in three exchanges, to establish the IKE SA. (These three steps are IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer.) Main mode also allows a peer to hide its identity.

Aggressive mode also establishes an authenticated IKE SA and keys. However, aggressive mode uses half the number of messages, has less negotiation power, and does not provide identity protection. The peer can use the aggressive or main mode to start IKE negotiation; the remote peer accepts the mode sent by the peer.

Copyright © 2013, Juniper Networks, Inc.



NOTE: The mode configuration is required only if the version option is set to 1.

To configure the mode for an IKE policy, include the **mode** statement and specify **aggressive** or **main** at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
mode (aggressive | main);
```

Configuring the Proposals in an IKE Policy

The IKE policy includes a list of one or more proposals associated with an IKE policy.

To configure the proposals in an IKE policy, include the **proposals** statement and specify one or more proposal names at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
proposals [ proposal-names ];
```

Configuring the Preshared Key for an IKE Policy

When you include the **authentication-method pre-shared-keys** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level, IKE policy preshared keys authenticate peers; for more information, see "Configuring the Authentication Method for an IKE Proposal" on page 229. You must manually configure a preshared key, which must match that of its peer. The preshared key can be an ASCII text (alphanumeric) key or a hexadecimal key.

To configure the preshared key in an IKE policy, include the **pre-shared-key** statement and a key at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
pre-shared-key (ascii-text key | hexadecimal key);
```

The key can be one of the following:

- **ascii-text**—ASCII text key. With the **des-cbc** option, the key contains 8 ASCII characters. With the **3des-cbc** option, the key contains 24 ASCII characters.
- **hexadecimal**—Hexadecimal key. With the **des-cbc** option, the key contains 16 hexadecimal characters. With the **3des-cbc** option, the key contains 48 hexadecimal characters.

Configuring the Local Certificate for an IKE Policy

When you include the **authentication-method rsa-signatures** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level, public key infrastructure (PKI) digital certificates authenticate peers; for more information, see "Configuring the Authentication Method for an IKE Proposal" on page 229. You must identify a local certificate that is sent to the peer during the IKE authentication phase.

Copyright © 2013, Juniper Networks, Inc.

To configure the local certificate for an IKE policy, include the **local-certificate** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
local-certificate identifier;
```

The **local-certificate** statement specifies the identifier used to obtain the end entity's certificate from the certification authority. Configuring it in an IKE policy allows you the flexibility of using a separate certificate with each remote peer if that is needed. You must also specify the identity of the certification authority by configuring the **ca-profile** statement at the **[edit security pki]** hierarchy level; for more information, see the Junos OS System Basics Configuration Guide. For complete examples of digital certificate configuration, see the Junos OS Feature Guides.

You can use the configured profiles to establish a set of trusted certification authorities for use with a particular service set. This enables you to configure separate service sets for individual clients to whom you are providing IP services; the distinct service sets provide logical separation of one set of IKE sessions from another, using different local gateway addresses, or *virtualization*. To configure the set of trusted certification authorities, include the **trusted-ca** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level:

```
[edit services service-set service-set-name ipsec-vpn-options]
trusted-ca ca-profile;
```

For more information, see "Configuring IPsec Service Sets" on page 473.

Configuring a Certificate Revocation List

A certificate revocation list (CRL) contains a list of digital certificates that have been cancelled before their expiration date. When a participating peer uses a digital certificate, it checks the certificate signature and validity. It also acquires the most recently issued CRL and checks that the certificate serial number is not on that CRL.



NOTE: By default, certificate revocation list verification is enabled. You can disable CRL verification by including the **disable** statement at the **[edit security pki ca-profile *ca-profile-name* revocation-check]** hierarchy level.

By default, if the router either cannot access the Lightweight Directory Access Protocol (LDAP) URL or retrieve a valid certificate revocation list, certificate verification fails and the IPsec tunnel is not established. To override this behavior and permit the authentication of the IPsec peer when the CRL is not downloaded, include the **disable on-download-failure** statement at the **[edit security pki ca-profile *ca-profile-name* revocation-check crl]** hierarchy level.

To use the CA certificate revocation list, you include statements at the **[edit security pki ca-profile *ca-profile-name* revocation-check]** hierarchy level. For details, see the Junos OS System Basics Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

Configuring the Description for an IKE Policy

To specify an optional text description for an IKE policy, include the **description** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
description description;
```

Configuring Local and Remote IDs for IKE Phase 1 Negotiation

You can optionally specify local identifiers for use in IKE phase 1 negotiation. If the **local-id** statement is omitted, the local gateway address is used.

To specify one or more local IDs, include the **local-id** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
```

You can also specify remote gateway identifiers for which the IKE policy is used. The remote gateway address in which this policy is defined is added by default.

To specify one or more remote IDs, include the **remote-id** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ike policy policy-name]
  remote-id {
  any-remote-id;
  ipv4_addr [ values ];
  ipv6_addr [ values ];
  key_id [ values ];
}
```

The **any-remote-id** option allows any remote address to connect. This option is supported only in dynamic endpoints configurations and cannot be configured along with specific values. For more information about dynamic endpoint configurations, see "Configuring Dynamic Endpoints for IPsec Tunnels" on page 250.

Example: Configuring an IKE Policy

Define two IKE policies: **policy 10.1.1.2** and **policy 10.1.1.1**. Each policy is associated with **proposal-1** and **proposal-2**. The following configuration uses only IKEv1 for negotiation.

```
[edit services ipsec-vpn]
ike {
  proposal proposal-1 {
    authentication-method pre-shared-keys;
    dh-group group1;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 1000;
  }
  proposal proposal-2 {
    authentication-method pre-shared-keys;
    dh-group group2;
```


Copyright © 2013, Juniper Networks, Inc.

```
    authentication-algorithm md5;
    encryption-algorithm des-cbc;
    lifetime-seconds 10000;
  }
  proposal proposal-3 {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm md5;
    encryption-algorithm des-cbc;
    lifetime-seconds 10000;
  }
  policy 10.1.1.2 {
    mode main;
    proposals [ proposal-1 proposal-2 ];
    pre-shared-key ascii-text example-pre-shared-key;
  }
  policy 10.1.1.1 {
    local-certificate certificate-file-name;
    local-key-pair private-public-key-file;
    mode aggressive;
    proposals [ proposal-2 proposal-3 ]
    pre-shared-key hexadecimal 0102030abbcd;
  }
}
```

NOTE: Updates to the current IKE proposal and policy configuration are not applied to the current IKE SA; updates are applied to new IKE SAs.

If you want the new updates to take immediate effect, you must clear the existing IKE security associations so that they will be reestablished with the changed configuration. For information about how to clear the current IKE security association, see the Junos OS Operational Mode Commands.

Configuring IPsec Proposals

An IPsec proposal lists protocols and algorithms (security services) to be negotiated with the remote IPsec peer.

To configure an IPsec proposal, include the **proposal** statement and specify an IPsec proposal name at the **[edit services ipsec-vpn ipsec]** hierarchy level:

```
[edit services ipsec-vpn ipsec]
proposal proposal-name {
  authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
  description description;
  encryption-algorithm algorithm;
  lifetime-seconds seconds;
  protocol (ah | esp | bundle);
}
```

Copyright © 2013, Juniper Networks, Inc.

This section discusses the following topics:

- Configuring the Authentication Algorithm for an IPsec Proposal on page 238
- Configuring the Description for an IPsec Proposal on page 238
- Configuring the Encryption Algorithm for an IPsec Proposal on page 238
- Configuring the Lifetime for an IPsec SA on page 239
- Configuring the Protocol for a Dynamic SA on page 240

Configuring the Authentication Algorithm for an IPsec Proposal

To configure the authentication algorithm for an IPsec proposal, include the **authentication-algorithm** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
```

The authentication algorithm can be one of the following:

- **hmac-md5-96**—Hash algorithm that authenticates packet data. It produces a 128-bit digest. Only 96 bits are used for authentication.
- **hmac-sha1-96**—Hash algorithm that authenticates packet data. It produces a 160-bit digest. Only 96 bits are used for authentication.

Configuring the Description for an IPsec Proposal

To specify an optional text description for an IPsec proposal, include the **description** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
description description;
```

Configuring the Encryption Algorithm for an IPsec Proposal

To configure encryption algorithm for an IPsec proposal, include the **encryption-algorithm** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
encryption-algorithm algorithm;
```

The encryption algorithm can be one of the following:

- **3des-cbc**—Encryption algorithm that has a block size of 24 bytes; its key size is 192 bits long.
- **des-cbc**—Encryption algorithm that has a block size of 8 bytes; its key size is 48 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

Copyright © 2013, Juniper Networks, Inc.



NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semiweak keys, see RFC 2409, *The Internet Key Exchange (IKE)*. The AES encryption algorithms use a software implementation that has much lower throughput, so DES remains the recommended option.

For **3des-cbc**, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

If you configure an authentication proposal but do not include the **encryption** statement, the result is NULL encryption. Certain applications expect this result. If you configure no specific authentication or encryption values, the Junos OS uses the default values of **sha1** for the authentication and **3des-cbc** for the encryption.

Configuring the Lifetime for an IPsec SA

When a dynamic IPsec SA is created, two types of lifetimes are used: hard and soft. The hard lifetime specifies the lifetime of the SA. The soft lifetime, which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire. This allows the key management system to negotiate a new SA before the hard lifetime expires.



NOTE: In IKEv1, the lifetime for SAs is negotiated with the remote peer based on the type of lifetime configured in the IPsec proposal. In IKEv2, such a negotiation is not performed with the remote peer. Instead, each IKE peer uses the lifetime that is locally configured for them.

For SAs in IKEv2, the lifetime is either the default value as IKEv1 (if another lifetime is not configured in the IPsec proposal) or all IKEv2 proposals in the IPsec policy must be configured with the same lifetime value.

To configure the hard lifetime value, include the **lifetime-seconds** statement and specify the number of seconds at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
lifetime-seconds seconds;
```

The default lifetime is 28,800 seconds. The range is from 180 through 86,400 seconds.

The soft lifetime values are as follows:

- Initiator: Soft lifetime = Hard lifetime – 135 seconds.
- Responder: Soft lifetime = Hard lifetime – 90 seconds.

Copyright © 2013, Juniper Networks, Inc.

Configuring the Protocol for a Dynamic SA

The **protocol** statement sets the protocol for a dynamic SA. IPsec uses two protocols to protect IP traffic: ESP and AH. The ESP protocol can support authentication, encryption, or both. The AH protocol is used for strong authentication. AH also authenticates the IP packet. The **bundle** option uses AH authentication and ESP encryption; it does not use ESP authentication because AH provides stronger authentication of IP packets.

To configure the protocol for a dynamic SA, include the **protocol** statement and specify the **ah**, **esp**, or **bundle** option at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec proposal proposal-name]
protocol (ah | esp | bundle);
```

Configuring IPsec Policies

An IPsec policy defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines Perfect Forward Secrecy (PFS) and the proposals needed for the connection. During the IPsec negotiation, IPsec looks for a proposal that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

A match is made when both policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used.

You can create multiple, prioritized IPsec proposals at each peer to ensure that at least one proposal matches a remote peer's proposal.

First, you configure one or more IPsec proposals; then you associate these proposals with an IPsec policy. You can prioritize a list of proposals used by IPsec in the **policy** statement by listing the proposals you want to use, from first to last.

To configure an IPsec policy, include the **policy** statement, and specify the policy name and one or more proposals to associate with the policy, at the **[edit services ipsec-vpn ipsec]** hierarchy level:

```
[edit services ipsec-vpn ipsec]
policy policy-name {
   description description;
   perfect-forward-secrecy {
      keys (group1 | group2 | group5 | group14);
   }
   proposals [ proposal-names ];
}
```

This section includes the following topics related to configuring an IPsec policy:

- Configuring the Description for an IPsec Policy on page 241
- Configuring Perfect Forward Secrecy on page 241

 Copyright © 2013, Juniper Networks, Inc.

- Configuring the Proposals in an IPsec Policy on page 241
- Example: Configuring an IPsec Policy on page 242

Configuring the Description for an IPsec Policy

To specify an optional text description for an IPsec policy, include the **description** statement at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec policy policy-name]
description description;
```

Configuring Perfect Forward Secrecy

PFS provides additional security by means of a Diffie-Hellman shared secret value. With PFS, if one key is compromised, previous and subsequent keys are secure because they are not derived from previous keys. This statement is optional.

To configure PFS, include the **perfect-forward-secrecy** statement and specify a Diffie-Hellman group at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec policy policy-name]
perfect-forward-secrecy {
  keys (group1 | group2 | group5 | group14);
}
```

The key can be one of the following:

- **group1**—Specifies that IKE use the 768-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group2**—Specifies that IKE use the 1024-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group5**—Specifies that IKE use the 1536-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- **group14**—Specifies that IKE use the 2048-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.

The higher numbered groups provide more security than the lowered numbered groups,, but require more processing time.

Configuring the Proposals in an IPsec Policy

The IPsec policy includes a list of one or more proposals associated with an IPsec policy.

To configure the proposals in an IPsec policy, include the **proposals** statement and specify one or more proposal names at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level:

```
[edit services ipsec-vpn ipsec policy policy-name]
proposals [ proposal-names ];
```

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring an IPsec Policy

Define an IPsec policy, **dynamic policy-1**, that is associated with two proposals (**dynamic-1** and **dynamic-2**):

```
[edit services ipsec-vpn ipsec]
proposal dynamic-1 {
   protocol esp;
   authentication-algorithm hmac-md5-96;
   encryption-algorithm 3des-cbc;
   lifetime-seconds 6000;
}
proposal dynamic-2 {
   protocol esp;
   authentication-algorithm hmac-sha1-96;
   encryption-algorithm 3des-cbc;
   lifetime-seconds 6000;
}
policy dynamic-policy-1 {
   perfect-forward-secrecy {
      keys group1;
   }
   proposals [ dynamic-1 dynamic-2 ];
}
```



NOTE: Updates to the current IPsec proposal and policy configuration are not applied to the current IPsec SA; updates are applied to new IPsec SAs.

If you want the new updates to take immediate effect, you must clear the existing IPsec security associations so that they will be reestablished with the changed configuration. For information about how to clear the current IPsec security association, see the Junos OS Operational Mode Commands.

IPsec Policy for Dynamic Endpoints

An IPsec policy for dynamic endpoints defines a combination of security parameters (IPsec proposals) used during IPsec negotiation between dynamic peer security gateways, in which the remote ends of tunnels do not have a statically assigned IP address.

During the IPsec negotiation, theIPsec policy looks for an IPsec proposal that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match. A match is made when the policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used.

If no policy is set, any policy proposed by the dynamic peer is accepted.

For more information about configuring IPsec policy, see "Configuring IPsec Policies" on page 240.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring IPsec Policies on page 240

Configuring IPsec Rules

To configure an IPsec rule, include the **rule** statement and specify a rule name at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      destination-address address;
      ipsec-inside-interface interface-name;
      source-address address;
    }
    then {
      anti-replay-window-size bits;
      backup-remote-gateway address;
      clear-dont-fragment-bit;
      dynamic {
        ike-policy policy-name;
        ipsec-policy policy-name;
      }
      initiate-dead-peer-detection;
      manual {
        direction (inbound | outbound | bidirectional) {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
          }
          auxiliary-spi spi-value;
          encryption {
            algorithm algorithm;
            key (ascii-text key | hexadecimal key);
          }
          protocol (ah | bundle | esp);
          spi spi-value;
        }
      }
      no-anti-replay;
      remote-gateway address;
      syslog;
      tunnel-mtu bytes;
    }
  }
}
```

Each IPsec rule consists of a set of terms, similar to a firewall filter. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.

Copyright © 2013, Juniper Networks, Inc.

- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how to configure the components of IPsec rules:

- Configuring Match Direction for IPsec Rules on page 244
- Configuring Match Conditions in IPsec Rules on page 244
- Configuring Actions in IPsec Rules on page 246

Configuring Match Direction for IPsec Rules

Each rule must include a **match-direction** statement that specifies whether the match is applied on the input or output side of the interface. To configure where the match is applied, include the **match-direction (input | output)** statement at the **[edit services ipsec-vpn rule *rule-name*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name]
match-direction (input | output);
```

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output. For more information on inside and outside interfaces, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

On the AS or Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that match the packet direction are considered.

Configuring Match Conditions in IPsec Rules

To configure the match conditions in an IPsec rule, include the **from** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name]
from {
  destination-address address;
  ipsec-inside-interface interface-name;
  source-address address;
}
```

You can use either the source address or the destination address as a match condition, in the same way that you would configure a firewall filter; for more information, see the Routing Policy Configuration Guide.

IPsec services support both IPv4 and IPv6 address formats. If you do not specifically configure either the source address or destination address, the default value **0.0.0.0/0**

Copyright © 2013, Juniper Networks, Inc.

(IPv4 ANY) is used. To use IPv6 ANY (**0::0/128**) as either source or destination address, you must configure it explicitly.

For next-hop-style service sets only, the **ipsec-inside-interface** statement allows you to assign a logical interface to the tunnels established as a result of this match condition. The **inside-service-interface** statement that you can configure at the **[edit services service-set *name* next-hop-service]** hierarchy level allows you to specify **.1** and **.2** as inside and outside interfaces. However, you can configure multiple adaptive services logical interfaces with the **service-domain inside** statement and use one of them to configure the **ipsec-inside-interface** statement. For more information, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468 and Interface Properties.

The Junos OS evaluates the criteria you configure in the **from** statement. If multiple link-type tunnels are configured within the same next-hop-style service set, the **ipsec-inside-interface** value enables the rule lookup module to distinguish a particular tunnel from other tunnels in case the source and destination addresses for all of them are **0.0.0.0/0** (ANY-ANY).



NOTE: When you configure the ipsec-inside-interface statement, interface-style service sets are not supported.

A special situation is provided by a term containing an "any-any" match condition (usually because the **from** statement is omitted). If there is an any-any match in a tunnel, a flow is not needed, because all flows within this tunnel use the same security association (SA) and packet selectors do not play a significant role. As a result, these tunnels will use packet-based IPsec. This strategy saves some flow resources on the PIC, which can be used for other tunnels that need a flow-based service.

The following configuration example shows an any-any tunnel configuration with no **from** statement in **term-1**. Missing selectors in the **from** clause result in a packet-based IPsec service.

```
services {
   ipsec-vpn {
      rule rule-1 {
         term term-1 {
            then {
               remote-gateway 10.1.0.1;
               dynamic {
                  ike-policy ike_policy;
                  ipsec-policy ipsec_policy;
               }
            }
         }
         match-direction input;
      }
      .....
   }
```

Flowless IPsec service is provided to link-type tunnels with an any-any matching, as well as to dynamic tunnels with any-any matching in both dedicated and shared mode.

Copyright © 2013, Juniper Networks, Inc.

For link-type tunnels, a mixture of flowless and flow-based IPsec is supported within a service set. If a service set includes some terms with any-any matching and some terms with selectors in the **from** clause, packet-based service is provided for the any-any tunnels and flow-based service is provided for the other tunnels with selectors.

For non link-type tunnels, if a service set contains both any-any terms and selector-based terms, flow-based service is provided to all the tunnels.

Configuring Actions in IPsec Rules

To configure actions in an IPsec rule, include the **then** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name]
then {
    anti-replay-window-size bits;
    backup-remote-gateway address;
    clear-dont-fragment-bit;
    dynamic {
        ike-policy policy-name;
        ipsec-policy policy-name;
    }
    initiate-dead-peer-detection;
    manual {
        direction (inbound | outbound | bidirectional) {
            authentication {
                algorithm (hmac-md5-96 | hmac-sha1-96);
                key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi spi-value;
            encryption {
                algorithm algorithm;
                key (ascii-text key | hexadecimal key);
            }
            protocol (ah | bundle | esp);
            spi spi-value;
        }
    }
    no-anti-replay;
    remote-gateway address;
    syslog;
    tunnel-mtu bytes;
}
```

The principal IPsec actions are to configure a dynamic or manual SA:

- You configure a dynamic SA by including the **dynamic** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level and referencing policies you have configured at the **[edit services ipsec-vpn ipsec]** and **[edit services ipsec-vpn ike]** hierarchy levels; for more information, see "Configuring Dynamic Security Associations" on page 227.
- You configure a manual SA by including the **manual** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level; for more information, see "Configuring Manual Security Associations" on page 223.

Copyright © 2013, Juniper Networks, Inc.

You can configure the following additional properties:

- Enabling IPsec Packet Fragmentation on page 247
- Configuring Destination Addresses for Dead Peer Detection on page 247
- Configuring or Disabling IPsec Anti-Replay on page 248
- Enabling System Log Messages on page 249
- Specifying the MTU for IPsec Tunnels on page 249

Enabling IPsec Packet Fragmentation

To enable fragmentation of IP version 4 (IPv4) packets in IPsec tunnels, include the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
clear-dont-fragment-bit;
```

Setting the **clear-dont-fragment-bit** statement clears the Don't Fragment (DF) bit in the packet header, regardless of the packet size. If the packet size exceeds the tunnel maximum transmission unit (MTU) value, the packet is fragmented before encapsulation. For IPsec tunnels, the default MTU value is 1500 regardless of the interface MTU setting.

Configuring Destination Addresses for Dead Peer Detection

To specify the remote address to which the IPsec traffic is directed, include the **remote-gateway** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
remote-gateway address;
```

To specify a backup remote address, include the **backup-remote-gateway** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
backup-remote-gateway address;
```

These two statements support both IPv4 and IPv6 address formats.

Configuring the **backup-remote-gateway** statement enables the dead peer detection (DPD) protocol, which monitors the tunnel state and remote peer availability. When the primary tunnel defined by the **remote-gateway** statement is active, the backup tunnel is in standby mode. If the DPD protocol determines that the primary remote gateway address is no longer reachable, a new tunnel is established to the backup address.

If there is no incoming traffic from a peer during a defined interval of 10 seconds, the router detects a tunnel as inactive. A global timer polls all tunnels every 10 seconds and the Adaptive Services (AS) or Multiservices Physical Interface Card (PIC) sends a message listing any inactive tunnels. If a tunnel becomes inactive, the router takes the following steps to failover to the backup address:

Copyright © 2013, Juniper Networks, Inc.

1. The adaptive services message triggers the DPD protocol to send a hello message to the peer.
2. If no acknowledgment is received, two retries are sent at 2-second intervals, and then the tunnel is declared dead.
3. Failover takes place if the tunnel is declared dead or there is an IPsec Phase 1 negotiation timeout. The primary tunnel is put in standby mode and the backup becomes active.
4. If the negotiation to the backup tunnel times out, the router switches back to the primary tunnel. If both peers are down, it tries the failover six times. It then stops failing over and reverts to the original configuration, with the primary tunnel active and the backup in standby mode.

You can also enable triggering of DPD Hello messages without configuring a backup remote gateway by including the **initiate-dead-peer-detection** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
initiate-dead-peer-detection;
```

The monitoring behavior is the same as described for the **backup-remote-gateway** statement. This configuration enables the router to initiate DPD Hellos when a backup IPsec gateway does not exist and clean up the IKE and IPsec SAs in case the IKE peer is not reachable.

If the DPD protocol determines that the primary remote gateway address is no longer reachable, a new tunnel is established to the backup address. However, when you configure **initiate-dead-peer-detection** without a backup remote gateway address and the DPD protocol determines that the primary remote gateway address is no longer reachable, the tunnel is declared dead and IKE and IPsec SAs are cleaned up.

For more information on the DPD protocol, see RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*.

Configuring or Disabling IPsec Anti-Replay

To configure the size of the IPsec antireplay window, include the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
anti-replay-window-size bits;
```

anti-replay-window-size can take values in the range from 64 through 4096 bits. The default value is 64 bits for AS PICs and 128 bits for Multiservices PICs and DPCs. AS PICs can support a maximum replay window size of 1024 bits, whereas Multiservices PICs and DPCs can support a maximum replay window size of 4096 bits. When the software is committing an IPsec configuration , the key management process (kmd) is unable to differentiate between the service interface types. As a result, if the maximum antireplay window size exceeds 1024 for AS PICs, the commit succeeds and no error message is produced. However, the software internally sets the antireplay window size for AS PICs to 1024 bits even if the configured value of the **anti-replay-window-size** is larger.

Copyright © 2013, Juniper Networks, Inc.

To disable the IPsec antireplay feature, include the **no-anti-replay** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
no-anti-replay;
```

By default, antireplay service is enabled. Occasionally this can cause interoperability issues with other vendors' equipment.

Enabling System Log Messages

To record an alert in the system logging facility, include the **syslog** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
syslog;
```

Specifying the MTU for IPsec Tunnels

To configure a specific maximum transmission unit (MTU) value for IPsec tunnels, include the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services ipsec-vpn rule rule-name term term-name then]
tunnel-mtu bytes;
```



NOTE: The **tunnel-mtu** setting is the only place you need to configure an MTU value for IPsec tunnels. Inclusion of an **mtu** setting at the **[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number* family inet]** hierarchy level is not supported.

Configuring IPsec Rule Sets

The **rule-set** statement defines a collection of IPsec rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services ipsec-vpn]** hierarchy level with a **rule** statement for each rule:

```
[edit services ipsec-vpn]
rule-set rule-set-name {
    rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Copyright © 2013, Juniper Networks, Inc.

Configuring Dynamic Endpoints for IPsec Tunnels

IPsec tunnels can also be established using *dynamic peer* security gateways, in which the remote ends of tunnels do not have a statically assigned IP address. Since the remote address is not known and might be pulled from an address pool each time the remote host reboots, establishment of the tunnel relies on using IKE **main** mode with either preshared global keys or digital certificates that accept any remote identification value. For more information on IKE policy modes, see "Configuring the Mode for an IKE Policy" on page 233. Both policy-based and link-type tunnels are supported:

- Policy-based tunnels used shared mode.
- Link-type or routed tunnels use dedicated mode. Each tunnel allocates a service interface from a pool of interfaces configured for the dynamic peers. Routing protocols can be configured to run on these service interfaces to learn routes over the IPsec tunnel that is used as a link in this scenario.

This section includes the following topics:

- Authentication Process on page 250
- Implicit Dynamic Rules on page 251
- Reverse Route Insertion on page 251
- Configuring an IKE Access Profile on page 252
- Referencing the IKE Access Profile in a Service Set on page 253
- Configuring the Interface Identifier on page 254
- Default IKE and IPsec Proposals on page 254

Authentication Process

The remote (dynamic peer) initiates the negotiations with the local (Juniper Networks) router. The local router uses the default IKE and IPsec policies to match the proposals sent by the remote peer to negotiate the security association (SA) values. Implicit proposals contain a list of all the supported transforms that the local router expects from all the dynamic peers.

If preshared key authentication is used, the preshared key is global for a service set. When seeking the preshared key for the peer, the local router matches the peer's source address against any explicitly configured preshared keys in that service set. If a match is not found, the local router uses the global preshared key for authentication. This key is the one configured in the IKE access profile referenced by the service set.

Phase 2 of the authentication matches the *proxy identities* of the protected hosts and networks sent by the peer against a list of configured proxy identities. The accepted proxy identity is used to create the dynamic rules for encrypting the traffic. You can configure proxy identities by including the **allowed-proxy-pair** statement in the IKE access profile. If no entry matches, the negotiation is rejected.

If you do not configure the **allowed-proxy-pair** statement, the default value **ANY(0.0.0.0/0)-ANY** is applied, and the local router accepts any proxy identities sent

Copyright © 2013, Juniper Networks, Inc.

by the peer. Both IPv4 and IPv6 addresses are accepted, but you must configure all IPv6 addresses manually.

Once the phase 2 negotiation completes successfully, the router builds the dynamic rules and inserts the reverse route into the routing table using the accepted proxy identity.

Implicit Dynamic Rules

After successful negotiation with the dynamic peer, the key management process (kmd) creates a dynamic rule for the accepted phase 2 proxy and applies it on the local AS or Multiservices PIC. The source and destination addresses are specified by the accepted proxy. This rule is used to encrypt traffic directed to one of the end hosts in the phase 2 proxy identity.

The dynamic rule includes an **ipsec-inside-interface** value, which is the interface name assigned to the dynamic tunnel. The **source-address** and **destination-address** values are accepted from the proxy ID. The **match-direction** value is **input** for next-hop-style service sets.



NOTE: **You do not configure this rule; it is created by the key management process (kmd).**

Rule lookup for static tunnels is unaffected by the presence of a dynamic rule; it is performed in the order configured. When a packet is received for a service set, static rules are always matched first.

Dynamic rules are matched after the rule match for static rules has failed.

Response to dead peer detection (DPD) hello messages takes place the same way with dynamic peers as with static peers. Initiating DPD hello messages from dynamic peers is not supported. For more information on DPD, see "Configuring Destination Addresses for Dead Peer Detection" on page 247.

Reverse Route Insertion

Static routes are automatically inserted into the route table for those networks and hosts protected by a remote tunnel endpoint. These protected hosts and networks are known as remote proxy identities.

Each route is created based on the remote proxy network and mask sent by the peer and is inserted in the relevant route table after successful phase 1 and phase 2 negotiations.

The route preference for each static reverse route is 1. This value is necessary to avoid conflict with similar routes that might be added by the routing protocol process (rpd).

No routes are added if the accepted remote proxy address is the default (**0.0.0.0/0**). In this case you can run routing protocols over the IPsec tunnel to learn routes and add static routes for the traffic you want to be protected over this tunnel.

For next-hop-style service sets, the reverse routes include next hops pointing to the locations specified by the **inside-service-interface** statement.

Copyright © 2013, Juniper Networks, Inc.

The route table in which to insert these routes depends on where the **inside-service-interface** location is listed. If these interfaces are present in a VPN routing and forwarding (VRF) instance, then routes are added to the corresponding VRF table; otherwise, the routes are added to **inet.0**.



NOTE: Reverse route insertion takes place only for tunnels to dynamic peers. These routes are added only for next-hop-style service sets.

Configuring an IKE Access Profile

You can configure only one tunnel profile per service set for all dynamic peers. The configured preshared key in the profile is used for IKE authentication of all dynamic peers terminating in that service set. Alternatively, you can include the **ike-policy** statement to reference an IKE policy you define with either specific identification values or a wildcard (the **any-remote-id** option). You configure the IKE policy at the **[edit services ipsec-vpn ike]** hierarchy level; for more information, see "Configuring IKE Policies" on page 232.

The IKE tunnel profile specifies all the information needed to complete the IKE negotiation. Each protocol has its own statement hierarchy within the client statement to configure protocol-specific attribute value pairs, but only one client configuration is allowed for each profile. The following is the configuration at the **[edit access]** hierarchy level; for more information on access profiles, see the Junos OS System Basics Configuration Guide.

```
[edit access]
profile profile-name {
   client * {
      ike {
         allowed-proxy-pair {
            remote remote-proxy-address local local-proxy-address;
         }
         pre-shared-key (ascii-text key-string | hexadecimal key-string);
         ike-policy policy-name;
         interface-id <string-value>;
         ipsec-policy ipsec-policy;
      }
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

NOTE: For dynamic peers, the Junos OS supports the IKE main mode with either the preshared key method of authentication or an IKE access profile that uses a local digital certificate.

- In preshared key mode, the IP address is used to identify a tunnel peer to get the preshared key information. The client value * (wildcard) means that configuration within this profile is valid for all dynamic peers terminating within the service set accessing this profile.
- In digital certificate mode, the IKE policy defines which remote identification values are allowed; for more information, see "Configuring IKE Policies" on page 232.

The following statements make up the IKE profile:

- **allowed-proxy-pair**—During phase 2 IKE negotiation, the remote peer supplies its network address (**remote**) and its peer's network address (**local**). Since multiple dynamic tunnels are authenticated through the same mechanism, this statement must include the list of possible combinations. If the dynamic peer does not present a valid combination, the phase 2 IKE negotiation fails.

  By default, **remote 0.0.0.0/0 local 0.0.0.0/0** is used if no values are configured. Both IPv4 and IPv6 address formats are supported in this configuration, but there are no default IPv6 addresses. You must specify even **0::0/0**.

- **pre-shared-key**—Key used to authenticate the dynamic peer during IKE phase 1 negotiation. This key is known to both ends through an out-of-band secure mechanism. You can configure the value either in **hexadecimal** or **ascii-text** format. It is a mandatory value.

- **ike-policy**—Policy that defines the remote identification values corresponding to the allowed dynamic peers; can contain a wildcard value **any-remote-id** for use in dynamic endpoint configurations only.

- **interface-id**—Interface identifier, a mandatory attribute used to derive the logical service interface information for the session.

- **ipsec-policy**—Name of the IPsec policy that defines the IPsec policy information for the session. You define the IPsec policy at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level. If no policy is set, any policy proposed by the dynamic peer is accepted.

Referencing the IKE Access Profile in a Service Set

To complete the configuration, you need to reference the IKE access profile configured at the **[edit access]** hierarchy level. To do this, include the **ike-access-profile** statement at the **[edit services service-set *name* ipsec-vpn-options]** hierarchy level:

```
[edit services service-set name]
ipsec-vpn-options {
    local-gateway address;
    ike-access-profile profile-name;
```

Copyright © 2013, Juniper Networks, Inc.

```
}
next-hop-service {
  inside-service-interface interface-name;
  outside-service-interface interface-name;
}
```

The **ike-access-profile** statement must reference the same name as the **profile** statement you configured for IKE access at the **[edit access]** hierarchy level. You can reference only one access profile in each service set. This profile is used to negotiate IKE and IPsec security associations with dynamic peers only.



NOTE: If you configure an IKE access profile in a service set, no other service set can share the same **local-gateway** address.

Also, you must configure a separate service set for each VRF instance. All interfaces referenced by the **ipsec-inside-interface** statement within a service set must belong to the same VRF instance.

Configuring the Interface Identifier

You can configure an interface identifier for a group of dynamic peers, which specifies which adaptive services logical interface(s) take part in the dynamic IPsec negotiation. By assigning the same interface identifier to multiple logical interfaces, you can create a pool of interfaces for this purpose. To configure an interface identifier, include the **ipsec-interface-id** statement and the **dedicated** or **shared** statement at the **[edit interfaces** *interface-name* **unit** *logical-unit-number* **dial-options]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number dial-options]
ipsec-interface-id identifier;
(dedicated | shared);
```

Specifying the interface identifier in the **dial-options** statement makes this logical interface part of the pool identified by the **ipsec-interface-id** statement.



NOTE: Only one interface identifier can be specified at a time. You can include the **ipsec-interface-id** statement or the **l2tp-interface-id** statement, but not both.

If you configure **shared** mode, it enables one logical interface to be shared across multiple tunnels. The **dedicated** statement specifies that the logical interface is used in a dedicated mode, which is necessary when you are configuring an IPsec link-type tunnel. You must include the **dedicated** statement when you specify an **ipsec-interface-id** value.

Default IKE and IPsec Proposals

The software includes implicit default IKE and IPsec proposals to match the proposals sent by the dynamic peers. The values are shown in Table 15 on page 255; if more than one value is shown, the first value is the default. For more information on IKE proposals,

 Copyright © 2013, Juniper Networks, Inc.

see "Configuring IKE Proposals" on page 228; for more information on IPsec proposals, see "Configuring IPsec Proposals" on page 237.

NOTE: RSA certificates are not supported with dynamic endpoint configuration.

Table 15: Default IKE and IPsec Proposals for Dynamic Negotiations

Statement Name	Values
Implicit IKE Proposal	
authentication-method	pre-shared keys
dh-group	group1, group2, group5, group14
authentication-algorithm	sha1, md5, sha-256
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256
lifetime-seconds	3600 seconds
Implicit IPsec Proposal	
protocol	esp, ah, bundle
authentication-algorithm	hmac-sha1-96, hmac-md5-96
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256
lifetime-seconds	28,800 seconds (8 hours)

Tracing IPsec Operations

Trace operations track IPsec events and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/kmd**.

To trace IPsec operations, include the **traceoptions** statement at the **[edit services ipsec-vpn]** hierarchy level:

```
[edit services ipsec-vpn]
traceoptions {
    file <filename> <files number> <match regular-expression> <size bytes> <world-readable |
        no-world-readable>;
    flag flag;
    level level;
    no-remote-trace;
}
```

Copyright © 2013, Juniper Networks, Inc.

You can specify the following IPsec tracing flags:

- **all**—Trace everything.
- **certificates**—Trace certificates events.
- **database**—Trace security associations database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **snmp**—Trace SNMP operations.
- **timer**—Trace internal timer events.

The **level** statement sets the key management process (kmd) tracing level. The following values are supported:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**–Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

Disabling IPsec Tunnel Endpoint in Traceroute

If you include the **no-ipsec-tunnel-in-traceroute** statement at the **[edit services ipsec-vpn]** hierarchy level, the IPsec tunnel is not treated as a next hop and TTL is not decremented. Also, if the TTL reaches zero, an ICMP time exceeded message is not generated.

```
[edit services ipsec-vpn]
no-ipsec-tunnel-in-traceroute;
```



NOTE: This functionality is also provided by the **passive-mode-tunneling** statement described in "Configuring IPsec Service Sets" on page 473. You can use the **no-ipsec-tunnel-in-traceroute** statement in specific scenarios in which the IPsec tunnel should not be treated as a next hop and passive mode is not desired.

Tracing IPsec PKI Operations

Trace operations track IPsec PKI events and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/pkid**.

Copyright © 2013, Juniper Networks, Inc.

To trace IPsec PKI operations, include the **traceoptions** statement at the **[edit security pki]** hierarchy level:

```
[edit security pki]
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag (all | certificate-verification | enrollment | online-crl-check);
}
```

You can specify the following PKI tracing flags:

- **all**—Trace everything.
- **certificates**—Trace certificates events.
- **database**—Trace security associations database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **snmp**—Trace SNMP operations.
- **timer**—Trace internal timer events.

Configuring IPSec on the Services SDK

Starting with Junos OS Release 11.4, IPSec is supported by the Services SDK. IPSec on the Services SDK is supported on all M Series, T Series and MX Series routers with Multiservices 100, Multiservices 400 PICs, and Multiservices DPCs.

IPSec on the Services SDK has the following limitations:

- IPSec on the Services SDK supports only policies negotiated between dynamic peer security gateways in which the remote ends of tunnels do not have a statically assigned IP address (Dynamic Endpoints).
- Encapsulating Security Payload (ESP) is the only protocol that is supported for protecting IP traffic.
- IPSec on the Services SDK does not support IPv6.

To enable IPSec for the Services SDK on the adaptive services interface, configure the **object-cache-size**, **policy-db-size**, and **package** statements at the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]** hierarchy level. For the IPSec plugin on the Services SDK, ***package-name*** in the **package *package-name*** statement is **jservices-ipsec**.

For more information about the Services SDK, see the *SDK Applications Configuration Guide and Command Reference*.

Copyright © 2013, Juniper Networks, Inc.

The following example shows how to enable IPSec for the Services SDK on the adaptive services interface:

```
chassis fpc 1 {
  pic 2 {
    adaptive-services {
      service-package {
        extension-provider {
          control-cores 1;
          data-cores 7;
          object-cache-size 1280;
          policy-db-size 64;
          package jservices-crypto-base;
          package jservices-ipsec;
        }
      }
    }
  }
}
```

Configure the inside and outside interfaces for next-hop-style service sets:

```
service-set abc {
  next-hop-service {
    inside-service-interface ms-0/2/0.1; # Name and logical unit number of the service
      interface associated with the service set applied inside the network.
    outside-service-interface ms-0/2/0.2; # Name and logical unit number of the service
      interface associated with the service set applied outside the network.
  }
}
```

Examples: Configuring IPsec Services

See the following sections:

- Example: Configuring Statically Assigned Tunnels on page 258
- Example: Configuring Dynamically Assigned Tunnels on page 261
- Multitask Example: Configuring IPsec Services on page 265

Example: Configuring Statically Assigned Tunnels

Following is the configuration of the provider edge (PE) router, demonstrating the usage of next-hop service sets and dynamic SA configuration:

```
[edit interfaces]
so-0/0/0 {
  no-keepalives;
  encapsulation cisco-hdlc;
  unit 0 {
    family inet {
      address 10.6.6.6/32;
    }
  }
}
so-2/2/0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
    description "teller so-0/2/0";
    no-keepalives;
    encapsulation cisco-hdlc;
    unit 0 {
      family inet {
        address 10.21.1.1/16;
      }
    }
}
sp-3/1/0 {
    unit 0 {
      family inet {
        address 10.7.7.7/32;
      }
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
}
[edit policy-options]
policy-statement vpn-export {
    then {
      community add vpn-comm;
      accept;
    }
}
policy-statement vpn-import {
    term a {
      from community vpn-comm;
      then accept;
    }
}
community vpn-comm members target:100:20;
[edit routing-instances]
vrf {
    instance-type vrf;
    interface sp-3/1/0.1; # Inside sp interface
    interface so-0/0/0.0;
    route-distinguisher 192.168.0.1:1;
    vrf-import vpn-import;
    vrf-export vpn-export;
    routing-options {
      static {
        route 10.0.0.0/0 next-hop so-0/0/0.0;
        route 10.11.11.1/32 next-hop so-0/0/0.0;
        route 10.8.8.1/32 next-hop sp-3/1/0.1;
      }
    }
}
[edit services]
ipsec-vpn {
```

Copyright © 2013, Juniper Networks, Inc.

```
    rule rule-1 {
      term term-1 {
        then {
          remote-gateway 10.21.2.1;
          dynamic {
            ike-policy ike-policy;
          }
        }
      }
      match-direction input;
    }
    ike {
      policy ike-policy {
        pre-shared-key ascii-text "$9$ExmcSeMWxdVYBI";
      }
    }
  }
  service-set service-set-1 {
    ipsec-vpn {
      local-gateway 10.21.1.1;
    }
    ipsec-vpn-rules rule-1;
    next-hop-service {
      inside-service-interface sp-3/1/0.1;
      outside-service-interface sp-3/1/0.2;
    }
  }
```

Following is an example for configuring multiple link-type tunnels to static peers using a single next-hop style service set:

```
  services ipsec-vpn {
    rule demo-rule {
      term term-0 {
        from {
          ipsec-inside-interface sp-0/0/0.1;
        }
        then {
          remote-gateway 10.2.2.2;
          dynamic {
            ike-policy demo-ike-policy;
          }
        }
      }
      term term-1 {
        from {
          ipsec-inside-interface sp-0/0/0.3;
        }
        then {
          remote-gateway 10.3.3.3;
          dynamic {
            ike-policy demo-ike-policy;
          }
        }
      }
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
    match-direction input;
}
services {
    service-set demo-service-set {
        next-hop-service {
            inside-service-interface sp-0/0/0.1;
            outside-service-interface sp-0/0/0.2;
        }
        ipsec-vpn-options {
            local-gateway 10.1.1.1;
        }
        ipsec-rules demo-rule;
    }
}
interfaces sp-0/0/0 {
    unit 0 {
        family inet;
    }
    unit 1 {
        family inet;
        service-domain inside;
    }
    unit 2 {
        family inet;
        service-domain outside;
    }
    unit 3 {
        family inet;
        service-domain inside;
    }
    unit 4 {
        family inet;
        service-domain inside;
    }
}
```

Example: Configuring Dynamically Assigned Tunnels

The following examples are based on this network configuration (see Figure 6 on page 262):

- A local network N-1 behind security gateway SG-1, a Juniper Networks router terminating static as well as dynamic peer endpoints. The tunnel termination address on SG-1 is 10.1.1.1 and the local network address is 172.16.1.0/24.
- Two remote peer routers that obtain addresses from an ISP pool and run RFC-compliant IKE. Remote network N-2 has address 172.16.2.0/24 and resides behind security gateway SG-2 with tunnel termination address 10.2.2.2. Remote network N-3 has address 172.16.3.0/24 and resides behind security gateway SG-3 with tunnel termination address 10.3.3.3.

Copyright © 2013, Juniper Networks, Inc.

Figure 6: IPsec Dynamic Endpoint Tunneling Topology

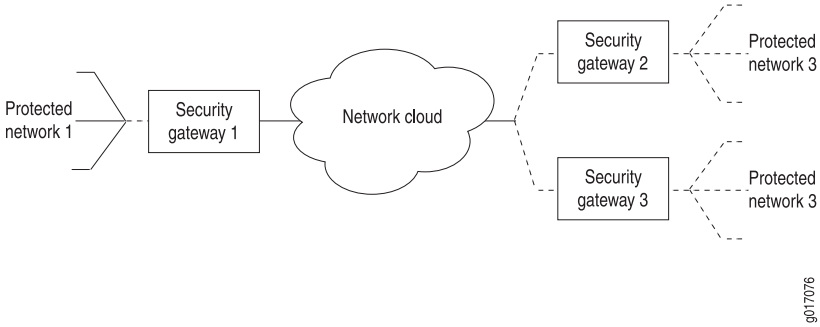

The examples in this section show the following configurations:

- Configuring a Next-Hop Style Service Set with Link-Type Tunnels on page 262
- Configuring a Next-Hop Style Service-Set with Policy-Based Tunnels on page 264

NOTE: All the configurations are given for the Juniper Networks router terminating dynamic endpoint connections.

Configuring a Next-Hop Style Service Set with Link-Type Tunnels

```
access {
   profile demo-access-profile client * {
      ike {
         allowed-proxy-pair {
            remote 0.0.0.0/0 local 0.0.0.0/0; # ANY to ANY
         }
         pre-shared-key {
            ascii-text keyfordynamicpeers;
         }
         interface-id demo-ipsec-interface-id;
      }
   }
   services {
      service-set demo-service-set {
         next-hop-service {
            inside-service-interface sp-1/0/0.1;
            outside-service-interface sp-1/0/0.2;
         }
         ipsec-vpn-options {
            local-gateway 10.1.1.1;
            ike-access-profile demo-ike-access-profile;
         }
      }
   }
}
```

NOTE: Including the ike-access-profile statement enables the software to incorporate implicit proposals for dynamic endpoint authentication. You do not need to configure IKE or IPsec proposals explicitly.

Copyright © 2013, Juniper Networks, Inc.

```
interfaces {
  sp-0/0/0 {
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
    unit 3 {
      family inet;
      service-domain inside;
      dial-options {
        ipsec-interface-id demo-ipsec-interface-id;
        dedicated;
      }
    }
    unit 4 {
      family inet;
      service-domain inside;
      dial-options {
        ipsec-interface-id demo-ipsec-interface-id;
        dedicated;
      }
    }
  }
}
```

The following results are obtained:

- Reverse routes inserted after successful negotiation:

  None

- Routes learned by routing protocol:

  172.16.2.0/24

  172.16.3.0/24

- Dynamic implicit rules created after successful negotiation:

```
rule: junos-dynamic-rule-0
term: term-0
  local-gateway-address : 10.1.1.1 #Tunnel termination address on SG-1
  remote-gateway-address: 10.2.2.2 #Tunnel termination address on SG-2
  source-address : 0.0.0.0/0
  destination-address : 0.0.0.0/0
ipsec-inside-interface: sp-0/0/0.3
term: term-1
  local-gateway-address : 10.1.1.1 #Tunnel termination address on SG-1
  remote-gateway-address: 10.3.3.3 #Tunnel termination address on SG-3
  source-address : 0.0.0.0/0
  destination-address : 0.0.0.0/0
```

Copyright © 2013, Juniper Networks, Inc.

```
    ipsec-inside-interface: sp-0/0/0.4
    match-direction: input
```

Configuring a Next-Hop Style Service-Set with Policy-Based Tunnels

```
access {
  profile demo-access-profile client * {
    ike {
      allowed-proxy-pair {
        remote 172.16.2.0/24 local 172.16.1.0/24; #N-2 <==> #N-1
        remote 172.16.3.0/24 local 172.16.1.0/24; #N-3 <==> #N-1
      }
      pre-shared-key {
        ascii-text keyfordynamicpeers;
      }
      interface-id demo-ipsec-interface-id;
    }
  }
}
services {
  service-set demo-service-set {
    next-hop-service {
      inside-service-interface sp-1/0/0.1;
      outside-service-interface sp-1/0/0.2;
    }
    ipsec-vpn-options {
      local-gateway 10.1.1.1;
    }
    ike-access-profile demo-ike-access-profile;
  }
}
```



NOTE: Including the **ike-access-profile** statement enables the software to incorporate implicit proposals for dynamic endpoint authentication. You do not need to configure IKE or IPsec proposals explicitly.

```
interfaces {
  sp-0/0/0 {
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet;
      service-domain inside;
    }
    unit 2 {
      family inet;
      service-domain outside;
    }
    unit 3 {
      family inet;
      service-domain inside;
      dial-options {
        ipsec-interface-id demo-ipsec-interface-id;
        mode shared;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
  }
  # VRF configuration, if not inet.0
  routing-instances {
    demo-vrf {
      instance-type vrf;
      interface sp-0/0/0.1;
      interface sp-0/0/0.3;
      .....
    }
  }
```

The following results are obtained:

- Reverse routes injected after successful negotiation:

```
demo-vrf.inet.0: .... # Routing instance
172.11.0.0/24 *[Static/1]..
   > via sp-0/0/0.3
   172.12.0.0/24 *[Static/1]..
> via sp-0/0/0.3
```

- Dynamic implicit rules created after successful negotiation:

```
rule: junos-dynamic-rule-0
term: term-0
   local-gateway-address : 10.1.1.1 #Tunnel termination address on SG-1
      remote-gateway-address: 10.2.2.2 #Tunnel termination address on SG-2
      source-address : 172.16.1.0/24
      destination-address : 172.16.2.0/24
   ipsec-inside-interface: sp-0/0/0.3
term: term-1
   local-gateway-address : 10.1.1.1 #Tunnel termination address on SG-1
      remote-gateway-address: 10.3.3.3 #Tunnel termination address on SG-3
      source-address : 172.16.1.0/24
      destination-address : 172.16.3.0/24
   ipsec-inside-interface: sp-0/0/0.3
match-direction: input
```

Multitask Example: Configuring IPsec Services

The following example-based instructions show how to configure IPsec services. The configuration involves defining an IKE policy, an IPsec policy, IPsec rules, trace options, and service sets.

This topic includes the following tasks:

1. Configuring the IKE Proposal on page 266
2. Configuring the IKE Policy (and Referencing the IKE Proposal) on page 266
3. Configuring the IPsec Proposal on page 267
4. Configuring the IPsec Policy (and Referencing the IPsec Proposal) on page 268
5. Configuring the IPsec Rule (and Referencing the IKE and IPsec Policies) on page 268

Copyright © 2013, Juniper Networks, Inc.

6. Configuring IPsec Trace Options on page 269
7. Configuring the Access Profile (and Referencing the IKE and IPsec Policies) on page 270
8. Configuring the Service Set (and Referencing the IKE Profile and the IPsec Rule) on page 271

Configuring the IKE Proposal

The IKE proposal configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway. For more information about IKE proposals, see "Configuring IKE Proposals" on page 228.

To define the IKE proposal:

1. In configuration mode, go to the following hierarchy level:

   user@host# edit services ipsec-vpn

2. Configure the authentication method, which is **pre-shared keys** in this example:

   [edit services ipsec-vpn]
   user@host# set ike proposal test-IKE-proposal authentication-method pre-shared-keys

3. Configure the Diffie-Hellman Group and specify a name—for example, **group1**:

   [edit services ipsec-vpn]
   user@host# set ike proposal test-IKE-proposal dh-group group1

4. Configure the authentication algorithm, which is **sha1** in this example:

   [edit services ipsec-vpn]
   user@host# set ike proposal test-IKE-proposal authentication-algorithm sha1

5. Configure the encryption algorithm, which is **aes-256-cbc** in this example:

   [edit services ipsec-vpn]
   user@host# set ike proposal test-IKE-proposal encryption-algorithm aes-256-cbc

The following sample output shows the configuration of the IKE proposal:

```
[edit services ipsec-vpn]
user@host# show ike
proposal test-IKE-proposal {
    authentication-method pre-shared-keys;
    dh-group group1;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
}
```

Configuring the IKE Policy (and Referencing the IKE Proposal)

The IKE policy configuration defines the proposal, mode, addresses, and other security parameters used during IKE negotiation. For more information about IKE policies, see "Configuring IKE Policies" on page 232.

To define the IKE policy and reference the IKE proposal:

1. In configuration mode, go to the following hierarchy level:

   user@host# edit services ipsec-vpn

Copyright © 2013, Juniper Networks, Inc.

2. Configure the IKE first phase mode—for example, **main**:

   ```
   [edit services ipsec-vpn]
   user@host# set ike policy test-IKE-policy mode main
   ```

3. Configure the proposal, which is **test-IKE-proposal** in this example:

   ```
   [edit services ipsec-vpn]
   user@host# set ike policy test-IKE-policy proposals test-IKE-proposal
   ```

4. Configure the local identification with an IPv4 address—for example, **192.168.255.2**:

   ```
   [edit services ipsec-vpn]
   user@host# set ike policy test-IKE-policy local-id ipv4_addr 192.168.255.2
   ```

5. Configure the preshared key in ASCII text format, which is **TEST** in this example:

   ```
   [edit services ipsec-vpn]
   user@host# set ike policy test-IKE-policy pre-shared-key ascii-text TEST
   ```

The following sample output shows the configuration of the IKE policy:

```
[edit services ipsec-vpn]
user@host# show ike
policy test-IKE-policy {
    mode main;
    proposals test-IKE-proposal;
    local-id ipv4_addr 192.168.255.2;
    pre-shared-key ascii-text TEST;
}
```

Configuring the IPsec Proposal

The IPsec proposal configuration defines the protocols and algorithms (security services) that are required to negotiate with the remote IPsec peer. For more information about IPsec proposals, see "Configuring IPsec Proposals" on page 237.

To define the IPsec proposal:

1. In configuration mode, go to the following hierarchy level:

   ```
   user@host# edit services ipsec-vpn
   ```

2. Configure the IPsec protocol for the proposal—for example, **esp**:

   ```
   [edit services ipsec-vpn]
   user@host# set ipsec proposal test-IPsec-proposal protocol esp
   ```

3. Configure the authentication algorithm for the proposal, which is **hmac-sha1-96** in this example:

   ```
   [edit services ipsec-vpn]
   user@host# set ipsec proposal test-IPsec-proposal authentication-algorithm
     hmac-sha1-96
   ```

4. Configure the encryption algorithm for the proposal, which is **aes-256-cbc** in this example:

   ```
   [edit services ipsec-vpn]
   ```

Copyright © 2013, Juniper Networks, Inc.

user@host# **set ipsec proposal test-IPsec-proposal encryption-algorithm aes-256-cbc**

The following sample output shows the configuration of the IPsec proposal:

```
[edit services ipsec-vpn]
user@host# show ike
proposal test-IPsec-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-256-cbc;
}
```

Configuring the IPsec Policy (and Referencing the IPsec Proposal)

The IPsec policy configuration defines a combination of security parameters (IPsec proposals) used during IPsec negotiation. It defines PFS and the proposals needed for the connection. For more information about IPsec policies, see "Configuring IPsec Policies" on page 240.

To define the IPsec policy and reference the IPsec proposal:

1. In configuration mode, go to the following hierarchy level:

   user@host# **edit services ipsec-vpn**

2. Configure the keys for perfect forward secrecy in the IPsec policy—for example, **group1**:

   [edit services ipsec-vpn]
   user@host# **set ipsec policy test-IPsec-policy perfect-forward-secrecy keys group1**

3. Configure a set of IPsec proposals in the IPsec policy—for example, **test-IPsec-proposal**:

   [edit services ipsec-vpn]
   user@host# **set ipsec policy test-IPsec-policy proposals test-IPsec-proposal**

The following sample output shows the configuration of the IPsec policy:

```
[edit services ipsec-vpn]
user@host# show ipsec policy test-IPsec-policy
perfect-forward-secrecy {
    keys group1;
}
proposals test-IPsec-proposal;
```

Configuring the IPsec Rule (and Referencing the IKE and IPsec Policies)

The IPsec rule configuration defines the direction that specifies whether the match is applied on the input or output side of the interface. The configuration also consists of a set of terms that specify the match conditions and applications that are included and excluded and also specify the actions and action modifiers to be performed by the router software. For more information about IPsec rules, see "Configuring IPsec Rules" on page 243.

To define the IPsec rule and reference the IKE and IPsec policies:

1. In configuration mode, go to the following hierarchy level:

   user@host# **edit services ipsec-vpn**

 Copyright © 2013, Juniper Networks, Inc.

2. Configure the IP destination address for the IPsec term in the IPsec rule—for example, **192.168.255.2/32**:

   ```
   [edit services ipsec-vpn]
   user@host# set rule test-IPsec-rule term 10 from destination-address 192.168.255.2/32
   ```

3. Configure the remote gateway address for the IPsec term in the IPsec rule—for example, **0.0.0.0**:

   ```
   [edit services ipsec-vpn]
   user@host# set rule test-IPsec-rule term 10 then remote-gateway 0.0.0.0
   ```

4. Configure a dynamic security association for IKE policy for the IPsec term in the IPsec rule, which is **test-IKE-policy** in this example:

   ```
   [edit services ipsec-vpn]
   user@host# set rule test-IPsec-rule term 10 then dynamic ike-policy test-IKE-policy
   ```

5. Configure a dynamic security association for IKE proposal for the IPsec term in the IPsec rule, which is **test-IPsec-proposal** in this example:

   ```
   [edit services ipsec-vpn]
   user@host# set rule test-IPsec-rule term 10 then dynamic ipsec-policy test-IPsec-policy
   ```

6. Configure a direction for which the rule match is being applied in the IPsec rule—for example, **input**:

   ```
   [edit services ipsec-vpn]
   user@host# set rule test-IPsec-rule match-direction input
   ```

The following sample output shows the configuration of the IPsec rule:

```
[edit services ipsec-vpn]
user@host# show rule test-IPsec-rule
term 10 {
    from {
        destination-address {
            192.168.255.2/32;
        }
    }
    then {
        remote-gateway 0.0.0.0;
        dynamic {
            ike-policy test-IKE-policy;
            ipsec-policy test-IPsec-policy;
        }
    }
}
match-direction input;
```

Configuring IPsec Trace Options

The IPsec trace options configuration tracks IPsec events and records them in a log file in the **/var/log** directory. By default, this file is named **/var/log/kmd**. For more information about IPsec rules, see "Tracing IPsec Operations" on page 255.

To define the IPsec trace options:

1. In configuration mode, go to the following hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit services ipsec-vpn
```

2. Configure the trace file, which is **ipsec.log** in this example:

```
[edit services ipsec-vpn]
user@host# set traceoptions file ipsec.log
```

3. Configure all the tracing parameters with the option **all** in this example:

```
[edit services ipsec-vpn]
user@host# set traceoptions flag all
```

The following sample output shows the configuration of the IPsec trace options:

```
[edit services ipsec-vpn]
user@host# show traceoptions
file ipsec.log;
flag all;
```

Configuring the Access Profile (and Referencing the IKE and IPsec Policies)

The access profile configuration defines the access profile and references the IKE and IPsec policies. For more information about access profile, see Configuring an IKE Access Profile.

To define the access profile and reference the IKE and IPsec policies:

1. In configuration mode, go to the following hierarchy level:

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

2. Configure the list of local and remote proxy identity pairs with the **allowed-proxy-pair** option. In this example, **10.0.0.0/24** is the IP address for local proxy identity and **10.0.1.0/24** is the IP address for remote proxy identity:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike allowed-proxy-pair local
  10.0.0.0/24 remote 10.0.1.0/24
```

3. Configure the IKE policy—for example, **test-IKE-policy**:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike ike-policy test-IKE-policy
```

4. Configure the IPsec policy—for example, **test-IPsec-policy**:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike ipsec-policy test-IPsec-policy
```

5. Configure the identity of logical service interface pool, which is **TEST-intf** in this example:

```
[edit access]
user@host# set profile IKE-profile-TEST client * ike interface-id TEST-intf
```

The following sample output shows the configuration of the access profile:

```
[edit access]
user@host# show
profile IKE-profile-TEST {
```


Copyright © 2013, Juniper Networks, Inc.

```
    client * {
        ike {
            allowed-proxy-pair local 10.0.0.0/24 remote 10.0.1.0/24;
            ike-policy test-IKE-policy;
            ipsec-policy test-IPsec-policy; # new statement
            interface-id TEST-intf;
        }
    }
}
```

Configuring the Service Set (and Referencing the IKE Profile and the IPsec Rule)

The service set configuration defines IPsec service sets that require additional specifications and references the IKE profile and the IPsec rule. For more information about IPsec service sets, see "Configuring IPsec Service Sets" on page 473.

To define the service set configuration with the next-hop service sets and IPsec VPN options:

1. In configuration mode, go to the following hierarchy level:

   user@host# [edit services]

2. Configure a service set with parameters for next hop service interfaces for the inside network—for example, **sp-1/2/0.1**:

   [edit services]
   user@host# set service-set TEST next-hop-service inside-service-interface sp-1/2/0.1

3. Configure a service set with parameters for next hop service interfaces for the outside network—for example, **sp-1/2/0.2**:

   [edit services]
   user@host# set service-set TEST next-hop-service outside-service-interface sp-1/2/0.2

4. Configure the IPsec VPN options with the address and routing instance for the local gateway—for example, **192.168.255.2**:

   [edit services]
   user@host# set service-set TEST ipsec-vpn-options local-gateway 192.168.255.2

5. Configure the IPsec VPN options with the IKE access profile for dynamic peers, which is **IKE-profile-TEST** in this example:

   [edit services]
   user@host# set service-set TEST ipsec-vpn-options ike-access-profile IKE-profile-TEST

6. Configure a service set with IPsec VPN rules, which is **test-IPsec-rule** in this example:

   [edit services]
   user@host# set service-set TEST ipsec-vpn-rules test-IPsec-rule

The following sample output shows the configuration of the service set configuration referencing the IKE profile and the IPsec rule:

```
[edit services]user@host# show service-set TEST
next-hop-service {
    inside-service-interface sp-1/2/0.1;
    outside-service-interface sp-1/2/0.2;
}
```

Copyright © 2013, Juniper Networks, Inc.

```
ipsec-vpn-options {
    local-gateway 192.168.255.2;
    ike-access-profile IKE-profile-TEST;
}
ipsec-vpn-rules test-IPsec-rule;
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 15

Summary of IPsec Services Configuration Statements

The following sections explain each of the IP Security (IPsec) services statements. The statements are organized alphabetically.

anti-replay-window-size (Services IPsec VPN)

Syntax anti-replay-window-size *bits*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the size of the IPsec antireplay window.

Options ***bits***—Size of the antireplay window, in bits.
Default: 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)
Range: 64 through 4096 bits

Usage Guidelines See "Configuring or Disabling IPsec Anti-Replay" on page 248.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

authentication (Services IPsec VPN)

Syntax

```
authentication {
  algorithm (hmac-md5-96 | hmac-sha1-96);
  key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure IPsec authentication parameters for a manual security association (SA).

Options

algorithm—Hash algorithm that authenticates packet data. The algorithm can be one of the following:

- **hmac-md5-96**—Produces a 128-bit digest.
- **hmac-sha1-96**—Produces a 160-bit digest.

key—Type of authentication key. The key can be one of the following:

- **ascii-text key**—ASCII text key. For **hmac-md5-96**, the key is 16 ASCII characters; for **hmac-sha1-96**, the key is 20 ASCII characters.
- **hexadecimal key**—Hexadecimal key. For **hmac-md5-96**, the key is 32 hexadecimal characters; for **hmac-sha1-96**, the key is 40 hexadecimal characters.

Usage Guidelines See "Configuring Authentication for a Manual IPsec SA" on page 225.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

authentication-algorithm

See the following sections:

- authentication-algorithm (Services IKE) on page 275
- authentication-algorithm (Services IPsec) on page 275

authentication-algorithm (Services IKE)

Syntax authentication-algorithm (md5 | sha1 | sha-256);

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.
sha-256 option added in Junos OS Release 7.6.

Description Configure the Internet Key Exchange (IKE) hash algorithm that authenticates packet data.

Options **md5**—Produces a 128-bit digest.

sha1—Produces a 160-bit digest.

sha-256—Produces a 256-bit digest.

Usage Guidelines See "Configuring the Authentication Algorithm for an IKE Proposal" on page 229.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

authentication-algorithm (Services IPsec)

Syntax authentication-algorithm (hmac-md5-96 | hmac-sha1-96);

Hierarchy Level [edit services ipsec-vpn ipsec proposal *ipsec-proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the IPsec hash algorithm that authenticates packet data.

Options **hmac-md5-96**—Produces a 128-bit digest.

hmac-sha1-96—Produces a 160-bit digest.

Usage Guidelines See "Configuring the Authentication Algorithm for an IPsec Proposal" on page 238.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

authentication-method (Services IPsec VPN)

Syntax authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an IKE authentication method.

Options **dsa-signatures**—Digital signature algorithm (DSA).

rsa-signatures—Public key algorithm (supports encryption and digital signatures).

pre-shared-keys—A key derived from an out-of-band mechanism; the key authenticates the exchange.

Usage Guidelines See "Configuring the Authentication Method for an IKE Proposal" on page 229.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

auxiliary-spi (Services IPsec VPN)

Syntax auxiliary-spi *spi-value*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an auxiliary Security Parameter Index (SPI) for a manual SA. Use the auxiliary SPI when you configure the **protocol** statement to use the **bundle** option.

Options ***spi-value***—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

Range: 256 through 16,639

Usage Guidelines See "Configuring the Auxiliary Security Parameter Index" on page 225. For information about SPI, see "Configuring the Security Parameter Index" on page 225 and **spi**.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

backup-remote-gateway

Syntax backup-remote-gateway *address*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the backup remote address to which the IPsec traffic is directed when the primary remote gateway is down. Configuring this statement also enables the dead peer detection (DPD) protocol.

Options ***address***—Backup remote IPv4 or IPv6 address.

Usage Guidelines See "Configuring Destination Addresses for Dead Peer Detection" on page 247.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

clear-dont-fragment-bit (Services IPsec VPN)

Syntax clear-dont-fragment-bit;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation.

Usage Guidelines See "Configuring Actions in IPsec Rules" on page 246.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

clear-ike-sas-on-pic-restart

Syntax clear-ike-sas-on-pic-restart;

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced in Junos OS Release 8.5.

Description Clear IKE security associations (SAs) when the corresponding PIC restarts or is taken offline.

Usage Guidelines See "Clearing Security Associations" on page 228.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

clear-ipsec-sas-on-pic-restart

Syntax clear-ipsec-sas-on-pic-restart;

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced in Junos OS Release 9.2.

Description Clear IPsec security associations (SAs) when the corresponding PIC restarts or is taken offline.

Usage Guidelines See "Clearing Security Associations" on page 228.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

description (Services IPsec VPN)

Syntax description *description*;

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*],
[edit services ipsec-vpn ike proposal *proposal-name*],
[edit services ipsec-vpn ipsec (Services IPsec VPN) policy *policy-name*],
[edit services ipsec-vpn ipsec (Services IPsec VPN) proposal *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the text description for an IKE or IPsec policy or proposal.

Usage Guidelines See "Configuring the Description for an IKE Policy" on page 236, "Configuring the Description for an IPsec Proposal" on page 238, and "Configuring the Description for an IPsec Policy" on page 241.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

destination-address (Services IPsec VPN)

Syntax destination-address *address*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the destination address for rule matching.

Options ***address***—Destination IP address.

Usage Guidelines See "Configuring Match Conditions in IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

dh-group

Syntax dh-group (group1 | group2 | group5 |group14);

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the IKE Diffie-Hellman prime modulus group to use for performing the new Diffie-Hellman exchange.

Options

group1—768-bit.

group2—1024-bit.

group5—1536-bit.

group14—2048-bit.

Usage Guidelines See "Configuring the Diffie-Hellman Group for an IKE Proposal" on page 230.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

direction

Syntax

```
direction (inbound | outbound | bidirectional) {
  protocol (ah | bundle | esp);
  spi spi-value;
  auxiliary-spi spi-value;
  authentication (Services IPsec VPN) {
    algorithm (hmac-md5-96 | hmac-sha1-96);
    key (ascii-text key | hexadecimal key);
  }
  encryption {
    algorithm algorithm;
    key (ascii-text key | hexadecimal key);
  }
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then manual]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the direction in which manual SAs are applied.

Options

bidirectional—Apply the SA in both directions.

inbound—Apply the SA on inbound traffic.

outbound—Apply the SA on outbound traffic.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 244.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

dynamic

Syntax

```
dynamic {
  ike-policy policy-name;
  ipsec-policy policy-name;
}
```

Hierarchy Level [edit services (IPsec VPN) ipsec-vpn rule (Services IPsec VPN) *rule-name* term (Services IPsec VPN) *term-name* then (Services IPsec VPN)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a dynamic IPsec SA.

Options

ike-policy *policy-name*—Name of the IKE policy. This statement is optional for the non-preshared-key authentication method. For digital signature-based authentication, this statement is optional and the default policy is used if none is supplied.

ipsec-policy *policy-name*—Name of the IPsec policy. This statement is optional and the default policy is used if none is supplied.

Usage Guidelines See "Configuring Dynamic Security Associations" on page 227.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

encryption

Syntax

```
encryption {
   algorithm algorithm;
   key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]

Release Information Statement introduced before Junos OS Release 7.4.
aes-128-cbc, **aes-192-cbc**, and **aes-256-cbc** options added in Junos OS Release 7.6.

Description Configure an encryption algorithm and key for manual SA.

Options **algorithm**—Type of encryption algorithm. The algorithm can be one of the following:

- **des-cbc**—Has a block size of 8 bytes (64 bits); the key size is 48 bits long.
- **3des-cbc**—Has a block size of 8 bytes (64 bits); the key size is 192 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
- **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.



NOTE: For 3des-cbc, the first 8 bytes should differ from the second 8 bytes, and the second 8 bytes should be the same as the third 8 bytes.

key—Type of encryption key. The key can be one of the following:

- **ascii-text**—ASCII text key. Following are the key lengths, in ASCII characters, for the different encryption options:
  - **des-cbc** option, 8 ASCII characters
  - **3des-cbc** option, 24 ASCII characters
  - **aes-128-cbc** option, 16 ASCII characters
  - **aes-192-cbc** option, 24 ASCII characters
  - **aes-256-cbc** option, 32 ASCII characters
- **hexadecimal**—Hexadecimal key. Following are the key lengths, in hexadecimal characters, for the different encryption options:
  - **des-cbc** option, 16 hexadecimal characters
  - **3des-cbc** option, 48 hexadecimal characters
  - **aes-128-cbc** option, 32 hexadecimal characters
  - **aes-192-cbc** option, 48 hexadecimal characters

Copyright © 2013, Juniper Networks, Inc.

- **aes-256-cbc** option, 64 hexadecimal characters

Usage Guidelines See "Configuring Encryption for a Manual IPsec SA" on page 226.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

encryption-algorithm (Services IPsec VPN)

Syntax encryption-algorithm *algorithm*;

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*],
[edit services ipsec-vpn ipsec proposal *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.
aes-128-cbc, **aes-192-cbc**, and **aes-256-cbc** options added in Junos OS Release 7.6.

Description Configure an IKE or IPsec encryption algorithm.

Options **3des-cbc**—Has a block size of 24 bytes; the key size is 192 bits long.

des-cbc—Has a block size of 8 bytes; the key size is 48 bits long.

aes-128-cbc—Advanced Encryption Standard (AES) 128-bit encryption algorithm.

aes-192-cbc—Advanced Encryption Standard (AES) 192-bit encryption algorithm.

aes-256-cbc—Advanced Encryption Standard (AES) 256-bit encryption algorithm.

Usage Guidelines See "Configuring the Encryption Algorithm for an IKE Proposal" on page 230 and "Configuring the Encryption Algorithm for an IPsec Proposal" on page 238.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

from (Services IPsec VPN)

Syntax

```
from {
  destination-address address;
  ipsec-inside-interface interface-name;
  source-address address;
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify input conditions for the IPsec term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ike

Syntax

```
ike {
  proposal proposal-name {
    authentication-algorithm (md5 | sha1 | sha-256);
    authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2 | group5 |group14);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
  }
  policy policy-name {
    description description;
    local-certificate identifier;
    local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
    version (1 | 2);
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
    remote-id {
      any-remote-id;
      ipv4_addr [ values ];
      ipv6_addr [ values ];
      key_id [ values ];
    }
  }
}
```

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure IKE.

The statements are explained separately.

Usage Guidelines See "Configuring IKE Proposals" on page 228 and "Configuring IKE Policies" on page 232.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

initiate-dead-peer-detection

Syntax initiate-dead-peer-detection;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable triggering of dead peer detection (DPD) Hello messages to the remote peer for the specified tunnel.

Usage Guidelines See "Configuring Destination Addresses for Dead Peer Detection" on page 247.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- backup-remote-gateway on page 277

ipsec (Services IPsec VPN)

Syntax

```
ipsec {
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    description description;
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
  }
  policy policy-name {
    description description;
    perfect-forward-secrecy {
      keys (group1 | group2);
    }
    proposals [ proposal-names ];
  }
}
```

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure IPsec.

The statements are explained separately.

Usage Guidelines See "Configuring Security Associations" on page 222.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ipsec-inside-interface

Syntax ipsec-inside-interface *interface-name*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 7.4.

Description Specify the interface name for next-hop-style service sets. This value is also implicitly generated in dynamic endpoint tunneling.

Options ***interface-name***—Service interface for internal network.

Usage Guidelines See "Configuring Match Conditions in IPsec Rules" on page 244 or "Configuring Dynamic Endpoints for IPsec Tunnels" on page 250.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

lifetime-seconds (Services IPsec VPN)

Syntax lifetime-seconds *seconds*;

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*],
[edit services ipsec-vpn ipsec proposal *proposal-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the lifetime of an IKE or IPsec SA. This statement is optional.

Options ***seconds***—Lifetime

Default: 3600 seconds (IKE); 28,800 seconds (IPsec)

Range: 180 through 86,400

Usage Guidelines See "Configuring the Lifetime for an IKE SA" on page 231 and "Configuring the Lifetime for an IPsec SA" on page 239.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

local-certificate (Services IPsec VPN)

Syntax local-certificate *identifier*;

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*]

Release Information Statement introduced in Junos OS Release 7.5.

Description Name of the certificate that needs to be sent to the peer during the IKE authentication phase.

Options ***identifier***—Name of certificate.

Usage Guidelines See "Configuring the Local Certificate for an IKE Policy" on page 234.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

local-id

Syntax local-id (ipv4_addr *ipv4-address* | ipv6-addr *ipv6-address* | key-id *identifier*);

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*]

Release Information Statement introduced before Junos OS Release 7.4.
ipv6_addr option added in Junos OS Release 7.6.

Description Specify local identifiers for IKE Phase 1 negotiation. This statement is optional.

Options **ipv4_addr** ***ipv4-address***—IPv4 address identification value.

ipv6_addr ***ipv6-address***—IPv6 address identification value.

key_id ***identifier***—Key identification value.

Usage Guidelines See "Configuring Local and Remote IDs for IKE Phase 1 Negotiation" on page 236.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

manual

Syntax

```
manual {
  direction (inbound | outbound | bidirectional) {
    authentication {
      algorithm (hmac-md5-96 | hmac-sha1-96);
      key (ascii-text key | hexadecimal key);
    }
    auxiliary-spi spi-value;
    encryption {
      algorithm algorithm;
      key (ascii-text key | hexadecimal key);
    }
    spi spi-value;
    protocol (ah | esp | bundle);
  }
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a manual IPsec SA.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Manual Security Associations" on page 223.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

match-direction (Services IPsec VPN)

Syntax match-direction (input | output);

Hierarchy Level [edit services ipsec-vpn rule *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the direction in which the rule match is applied.

Options **input**—Apply the rule match on input.

output—Apply the rule match on output.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

mode (Services IPsec VPN)

Syntax mode (aggressive | main);

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IKE policy mode.

Default main

Options
- **aggressive**—Takes half the number of messages of main mode, has less negotiation power, and does not provide identity protection.
- **main**—Uses six messages, in three peer-to-peer exchanges, to establish the IKE SA. These three steps include the IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer. Also provides identity protection.

Usage Guidelines See "Configuring the Mode for an IKE Policy" on page 233.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

no-anti-replay (Services IPsec VPN)

Syntax no-anti-replay;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Disable IPsec antireplay service, which occasionally causes interoperability issues for security associations.

Usage Guidelines See "Configuring or Disabling IPsec Anti-Replay" on page 248.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

no-ipsec-tunnel-in-traceroute

Syntax no-ipsec-tunnel-in-traceroute;

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced in Junos OS Release 10.0.

Description Disables displaying the IPsec tunnel endpoint in the trace route output. The IPsec tunnel is not treated as a next hop and TTL is not decremented. If the TTL becomes zero, the ICMP time exceeded message will not be generated.

Usage Guidelines See "Configuring or Disabling IPsec Anti-Replay" on page 248.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

perfect-forward-secrecy (Services IPsec VPN)

Syntax
```
perfect-forward-secrecy {
  keys (group1 | group2 |group5 |group14);
}
```

Hierarchy Level [edit services ipsec-vpn ipsec policy *policy-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define Perfect Forward Secrecy (PFS). Creates single-use keys. This statement is optional.

Options **keys**—Type of Diffie-Hellman prime modulus group that IKE uses when performing the new Diffie-Hellman exchange. The key can be one of the following:

- **group1**—768-bit.
- **group2**—1024-bit.
- **group5**—1536-bit.
- **group14**—2048-bit.

Usage Guidelines See "Configuring Perfect Forward Secrecy" on page 241.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

policy

See the following sections:

- policy (Services IKE) on page 293
- policy (Services IPsec VPN) on page 294

policy (Services IKE)

Syntax

```
policy policy-name {
  description description;
  local-certificate identifier;
  local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
  version (1 | 2);
  mode (aggressive | main);
  pre-shared-key (ascii-text key | hexadecimal key);
  proposals [ proposal-names ];
  remote-id {
    any-remote-id;
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
  }
}
```

Hierarchy Level [edit services ipsec-vpn ike]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IKE policy.

Options *policy-name*—IKE policy name.

The remaining statements are explained separately.

Usage Guidelines See "Configuring IKE Policies" on page 232.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

policy (Services IPsec VPN)

Syntax

```
policy policy-name {
  description description;
  perfect-forward-secrecy {
    keys (group1 | group 14 | group2 | group 5);
  }
  proposals [ proposal-names ];
}
```

Hierarchy Level [edit services ipsec-vpn ipsec]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IPsec policy.

Options *policy-name*—IPsec policy name.

The remaining statements are explained separately.

Usage Guidelines See "Configuring IPsec Policies" on page 240.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

pre-shared-key (Services IKE)

Syntax pre-shared-key (ascii-text *key* | hexadecimal *key*);

Hierarchy Level [edit services ike policy *policy-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a preshared key for an IKE policy.

Options *key*—Value of preshared key. The key can be one of the following:

- **ascii-text**—ASCII text key.
- **hexadecimal**—Hexadecimal key.

Usage Guidelines See "Configuring IKE Policies" on page 232.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

proposal

See the following sections:

- proposal (Services IKE) on page 295
- proposal (Services IPsec VPN) on page 296

proposal (Services IKE)

Syntax

```
proposal proposal-name {
  authentication-algorithm (md5 | sha1 | sha-256);
  authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
  description description;
  dh-group (group1 | group2 | group5 |group14);
  encryption-algorithm algorithm;
  lifetime-seconds seconds;
}
```

Hierarchy Level [edit services ipsec-vpn ike]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IKE proposal for a dynamic SA.

Options *proposal-name*—IKE proposal name.

The remaining statements are explained separately.

Usage Guidelines See "Configuring IKE Proposals" on page 228.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

proposal (Services IPsec VPN)

Syntax

```
proposal proposal-name {
  authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
  description description;
  encryption-algorithm algorithm;
  lifetime-seconds seconds;
  protocol (ah | esp | bundle);
}
```

Hierarchy Level [edit services ipsec-vpn ipsec]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IPsec proposal for a dynamic SA.

Options *proposal-name*—IPsec proposal name.

The remaining statements are explained separately.

Usage Guidelines See "Configuring IPsec Proposals" on page 237.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

proposals

Syntax proposals [*proposal-names*];

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*],
[edit services ipsec-vpn ipsec policy *policy-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a list of proposals to include in the IKE or IPsec policy.

Options *proposal-names*—List of IKE or IPsec proposal names.

Usage Guidelines See "Configuring the Proposals in an IKE Policy" on page 234 and "Configuring the Proposals in an IPsec Policy" on page 241.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

protocol

Syntax protocol (ah | esp | bundle);

Hierarchy Level [edit services ipsec-vpn ipsec proposal *proposal-name*],
[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define an IPsec protocol for a dynamic or manual SA.

Options **ah**—Authentication Header protocol.

esp—Encapsulating Security Payload protocol.

bundle—AH and ESP protocol.

Usage Guidelines See "Configuring the Protocol for a Manual IPsec SA" on page 225.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

remote-gateway

Syntax remote-gateway *address*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the remote address to which the IPsec traffic is directed.

Options ***address***—Remote IPv4 or IPv6 address.

Usage Guidelines See "Configuring Actions in IPsec Rules" on page 246.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

remote-id

Syntax

```
remote-id {
    any-remote-id;
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
}
```

Hierarchy Level [edit services ipsec-vpn ikepolicy *policy-name*]

Release Information Statement introduced before Junos OS Release 7.4.
ipv6_addr option added in Junos OS Release 7.6.
any-remote-id option added in Junos OS Release 8.2.

Description Define the remote identification values to which the IKE policy applies.

Options

any-remote-id—Allow any remote address to connect. This option is supported only in dynamic endpoints configurations and cannot be configured along with specific values.

ipv4_addr [*values*]—Define one or more IPv4 address identification values.

ipv6_addr [*values*]—Define one or more IPv6 address identification values.

key_id [*values*]—Define one or more key identification values.

Usage Guidelines See "Configuring Local and Remote IDs for IKE Phase 1 Negotiation" on page 236.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rule (Services IPsec VPN)

Syntax

```
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      destination-address address;
      ipsec-inside-interface interface-name;
      source-address address;
    }
    then {
      anti-replay-window-size bits;
      backup-remote-gateway address;
      clear-dont-fragment-bit;
      dynamic {
        ike-policy policy-name;
        ipsec-policy policy-name;
      }
      initiate-dead-peer-detection;
      manual {
        direction (inbound | outbound | bidirectional) {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
          }
          auxiliary-spi spi-value;
          encryption {
            algorithm algorithm;
            key (ascii-text key | hexadecimal key);
          }
          protocol (ah | bundle | esp);
          spi spi-value;
        }
      }
      no-anti-replay;
      remote-gateway address;
      syslog;
      tunnel-mtu bytes;
    }
  }
}
```

Hierarchy Level

[edit services ipsec-vpn],
[edit services ipsec-vpn rule-set *rule-set-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that comprise this rule.

The remaining statements are explained separately.

Copyright © 2013, Juniper Networks, Inc.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

rule-set (Services IPsec VPN)

Syntax

```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

Usage Guidelines See "Configuring IPsec Rule Sets" on page 249.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

services (IPsec VPN)

Syntax services ipsec-vpn { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service rules to be applied to traffic.

Options **ipsec-vpn**—IPsec set of rules statements.

Usage Guidelines See IPsec Properties.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

source-address (Services IPsec VPN)

Syntax source-address *address*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* from]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the source address for rule matching.

Options ***address***—Source IP address.

Usage Guidelines See "Configuring Match Conditions in IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

spi

Syntax spi *spi-value*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the SPI for an SA.

Options ***spi-value***—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).

Range: 256 through 16,639



NOTE: Use the auxiliary SPI when you configure the protocol statement to use the bundle option.

Usage Guidelines See "Configuring the Security Parameter Index" on page 225.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

syslog (Services IPsec VPN)

Syntax syslog;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable system logging. The system log information for the Adaptive Services or Multiservices Physical Interface Card (PIC) is passed to the kernel for logging in the **/var/log** directory.

Usage Guidelines See "Configuring Actions in IPsec Rules" on page 246.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

term (Services IPsec VPN)

Syntax

```
term term-name {
  from {
    destination-address address;
    ipsec-inside-interface interface-name;
    source-address address;
  }
  then {
    anti-replay-window-size bits;
    backup-remote-gateway address;
    clear-dont-fragment-bit;
    dynamic {
      ike-policy policy-name;
      ipsec-policy policy-name;
    }
    initiate-dead-peer-detection;
    manual {
      direction (inbound | outbound | bidirectional) {
        authentication {
          algorithm (hmac-md5-96 | hmac-sha1-96);
          key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi spi-value;
        encryption {
          algorithm algorithm;
          key (ascii-text key | hexadecimal key);
        }
        protocol (ah | bundle | esp);
        spi spi-value;
      }
    }
    no-anti-replay;
    remote-gateway address;
    syslog;
    tunnel-mtu bytes;
  }
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IPsec term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

then (Services IPsec VPN)

Syntax

```
then {
  anti-replay-window-size bits;
  backup-remote-gateway address;
  clear-dont-fragment-bit;
  dynamic {
    ike-policy policy-name;
    ipsec-policy policy-name;
  }
  initiate-dead-peer-detection;
  manual {
    direction (inbound | outbound | bidirectional) {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      auxiliary-spi spi-value;
      encryption {
        algorithm algorithm;
        key (ascii-text key | hexadecimal key);
      }
      protocol (ah | bundle | esp);
      spi spi-value;
    }
  }
  no-anti-replay;
  remote-gateway address;
  syslog;
  tunnel-mtu bytes;
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the IPsec term actions.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 244.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Services IPsec VPN)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size bytes> <world-readable |
    no-world-readable>;
  flag flag;
  level level;
  no-remote-trace;
}
```

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced in Junos OS Release 7.5.
level option added in Junos OS Release 10.0.

Description Configure IPsec tracing operations. By default, messages are written to **/var/log/kmd**.

Options

files *number*—Maximum number of trace data files.

Range: 2 through 1000

flag *flag*—Tracing operation to perform:

- **all**—Trace everything.
- **certificates**—Trace certificates that apply to the IPsec service set.
- **database**—Trace security associations database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **snmp**—Trace SNMP operations.
- **timer**—Trace internal timer events.

level *level*—Key management process (kmd) tracing level. The following values are supported:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**–Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

Copyright © 2013, Juniper Networks, Inc.

size *bytes*—Maximum trace file size.

Usage Guidelines See "Tracing IPsec Operations" on page 255.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

traceoptions (PKI)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
}
```

Hierarchy Level [edit security pki]

Description Configure security public key infrastructure (PKI) trace options. To specify more than one trace option, include multiple **flag** statements. Trace option output is recorded in the **/var/log/pkid** file.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the **file** statement, you must specify a filename.

files *number*—(Optional) Maximum number of trace files. When a trace file (for example, **pkid**) reaches its maximum size, it is renamed **pkid.0**, then **pkid.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 2 files

flag—Trace operation to perform. To specify more than one trace operation, include multiple flag statements:

all—Trace with all flags enabled.

certificate-verification—Trace PKI certificate verification events.

online-crl-check—Trace PKI online certificate revocation list (CRL) events.

enrollment—PKI certificate enrollment tracing.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

size *maximum-file-size*—(Optional) Maximum size of each trace file, in kilobytes (KB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files *number*** option.

Default: 1024 KB

world-readable | no-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

tunnel-mtu (Services IPsec VPN)

Syntax
tunnel-mtu *bytes*;

Hierarchy Level
[edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information
Statement introduced in Junos OS Release 7.5.

Description
Maximum transmission unit (MTU) size for IPsec tunnels.

Options
bytes—MTU size.
Default: 1500 bytes
Range: 256 through 9192 bytes

Usage Guidelines
See "Specifying the MTU for IPsec Tunnels" on page 249.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- mtu on page 1205

version (IKE)

Syntax
version (1 | 2);

Hierarchy Level
[edit services ipsec-vpn ike policy *policy-name*],

Release Information
Statement introduced in Junos OS Release 11.4.

Description
Configure the Internet Key Exchange (IKE) version that is used to negotiate dynamic SAs for IPSec.

Options
1—Uses IKEv1.

2—Uses IKEv2.

Usage Guidelines
See "Configuring IKE Policies" on page 232.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 16

Layer 2 Tunneling Protocol Services Configuration Guidelines

The Layer 2 Tunneling Protocol (L2TP) enables you to set up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented. Multiple L2TP PPP sessions can share the same remote peer IP address, which enables you to set up redundant sessions between the same links.

- If you configure Multilink PPP, the same remote IP address can be shared across multiple bundles, because the IP address negotiation takes place on the bundle rather than on each session.
- If Multilink PPP is not configured, multiple sessions can share the same remote IP address.

The last session or bundle to come up accomplishes the traffic transfer. When this session or bundle goes down, the traffic switches to the next-to-last session or bundle to come up. For example, if four sessions or bundles labeled A, B, C, and D share the same remote IP address and come up in alphabetical order, D initially handles the data transfer. If D goes down, traffic switches over to C, and so forth. If another session or bundle E subsequently comes up and has the same address, the traffic switches over to it.

To configure L2TP services, include the **l2tp** statement at the **[edit services]** hierarchy level:

```
[edit services]
l2tp {
    ip-reassembly (L2TP);
    tunnel-group group-name {
        hello-interval seconds;
        hide-avps;
        l2tp-access-profile profile-name;
        local-gateway address address;
        maximum-send-window packets;
        ppp-access-profile profile-name;
        receive-window packets;
        retransmit-interval seconds;
        service-interface interface-name;
        syslog {
            host hostname {
                services severity-level;
```

Copyright © 2013, Juniper Networks, Inc.

```
            facility-override facility-name;
            log-prefix prefix-value;
          }
        }
        tunnel-timeout seconds;
      }
      traceoptions {
        debug-level level;
        filter {
          protocol name;
        }
        flag flag;
        interfaces interface-name {
          debug-level level;
          flag flag;
        }
      }
    }
```



NOTE: L2TP configurations on Adaptive Services and Multiservices PICs are supported only on M7i, M10i, and M120 routers. For information about L2TP LAC and LNS configurations on MX Series routers for subscriber access, see L2TP for Subscriber Access Overview in the *Junos Subscriber Access Configuration Guide*.

You configure other components of this feature at the **[edit access]** and **[edit interfaces]** hierarchy levels. Those configurations are summarized in this chapter; for more information, see the Junos OS System Basics Configuration Guide or the Junos® OS Network Interfaces.

This chapter contains the following sections:

- L2TP Services Configuration Overview on page 311
- L2TP Minimum Configuration on page 312
- Configuring L2TP Tunnel Groups on page 314
- Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 318
- AS PIC Redundancy for L2TP Services on page 320
- IP Packet Fragment Reassembly for L2TP Overview on page 320
- Configuring IP Inline Reassembly for L2TP on page 321
- Tracing L2TP Operations on page 322
- Examples: Configuring L2TP Services on page 324

Copyright © 2013, Juniper Networks, Inc.

L2TP Services Configuration Overview

The statements for configuring L2TP services are found at the following hierarchy levels:

- **[edit services l2tp tunnel-group** *group-name***]**

  The L2TP **tunnel-group** statement identifies an L2TP instance or L2TP server. Associated statements specify the local gateway address on which incoming tunnels and sessions are accepted, the Adaptive Services (AS) Physical Interface Card (PIC) that processes data for the sessions in this tunnel group, references to L2TP and PPP access profiles, and other attributes for configuring window sizes and timer values.

- **[edit interfaces sp-***fpc/pic/port* **unit** *logical-unit-number* **dial-options]**

  The **dial-options** statement includes configuration for the **l2tp-interface-id** statement and the **shared**/**dedicated** flag. The interface identifier associates a user session with a logical interface. Sessions can use either shared or dedicated logical interfaces. To run routing protocols, a session must use a dedicated logical interface.

- **[edit access profile** *profile-name* **client** *name* **l2tp]**

  Tunnel profiles are defined at the **[edit access]** hierarchy level. Tunnel clients are defined with authentication, multilink negotiation and fragmentation, and other L2TP attributes in these profiles.

- **[edit access profile** *profile-name* **client** *name* **ppp]**

  User profiles are defined at the **[edit access]** hierarchy level. User clients are defined with authentication and other PPP attributes in these profiles. These client profiles are used when local authentication is specified.

- **[edit access radius-server** *address***]**

  When you configure **authentication-order radius** at the **[edit access profile** *profile-name***]** hierarchy level, you must configure a RADIUS service at the **[edit access radius-server]** hierarchy level.



NOTE: For more information about configuring properties at the **[edit access]** hierarchy level, see the Junos OS System Basics Configuration Guide. For information about L2TP LAC and LNS configurations on MX Series routers for subscriber access, see L2TP for Subscriber Access Overview in the *Junos Subscriber Access Configuration Guide*.

Copyright © 2013, Juniper Networks, Inc.

L2TP Minimum Configuration

To configure L2TP services, you must perform at least the following tasks:

- Define a tunnel group at the **[edit services l2tp]** hierarchy level with the following attributes:
  - **l2tp-access-profile**—Profile name for the L2TP tunnel.
  - **ppp-access-profile**—Profile name for the L2TP user.
  - **local-gateway**—Address for the L2TP tunnel.
  - **service-interface**—AS PIC interface for the L2TP service.
  - Optionally, you can configure **traceoptions** for debugging purposes.

  The following example shows a minimum configuration for a tunnel group with trace options:

```
[edit services l2tp]
tunnel-group finance-lns-server {
   l2tp-access-profile westcoast_bldg_1_tunnel;
   ppp-access-profile westcoast_bldg_1;
   local-gateway {
      address 10.21.255.129;
   }
   service-interface sp-1/3/0;
}
traceoptions {
   flag all;
   filter {
      protocol udp;
      protocol l2tp;
      protocol ppp;
      protocol radius;
   }
}
```

- At the **[edit interfaces]** hierarchy level:
  - Identify the physical interface at which L2TP tunnel packets enter the router, for example **ge-0/3/0**.
  - Configure the AS PIC interface with **unit 0 family inet** defined for IP service, and configure another logical interface with **family inet** and the **dial-options** statement.

  The following example shows a minimum interfaces configuration for L2TP:

```
[edit interfaces]
ge-0/3/0 {
   unit 0 {
      family inet {
         address 10.58.255.129/28;
      }
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

```
sp-1/3/0 {
  unit 0 {
    family inet;
  }
  unit 20 {
    dial-options {
      l2tp-interface-id test;
      shared;
    }
    family inet;
  }
}
```

- At the **[edit access]** hierarchy level:
  - Configure a tunnel profile. Each client specifies a unique L2TP Access Concentrator (LAC) name with an **interface-id** value that matches the one configured on the AS PIC interface unit; **shared-secret** is authentication between the LAC and the L2TP Network Server (LNS).
  - Configure a user profile. If RADIUS is used as the authentication method, it needs to be defined.
  - Define the RADIUS server with an IP address, port, and authentication data shared between the router and the RADIUS server.

    

    NOTE: When the L2TP Network Server (LNS) is configured with RADIUS authentication, the default behavior is to accept the preferred RADIUS-assigned IP address. Previously, the default behavior was to accept and install the nonzero peer IP address that came into the IP-Address option of the IPCP Configuration Request packet.

  - Optionally, you can define a group profile for common attributes, for example **keepalive 0** to turn off keepalive messages.

  The following example shows a minimum profiles configuration for L2TP:

```
[edit access]
group-profile westcoast_users {
  ppp {
    keepalive 0;
  }
}
profile westcoast_bldg_1_tunnel {
  client production {
    l2tp {
      interface-id test;
      shared-secret "$9$n8HX6A01RhlvL1R"; # SECRET-DATA
    }
    user-group-profile westcoast_users;
  }
}
profile westcoast_bldg_1 {
  authentication-order radius;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    radius-server {
      192.168.65.63 {
        port 1812;
        secret "$9$Vyb4ZHkPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
      }
    }
```

Configuring L2TP Tunnel Groups

To establish L2TP service on a router, you need to identify an L2TP tunnel group and specify a number of values that define which access profiles, interface addresses, and other properties to use in creating a tunnel. To identify the tunnel group, include the **tunnel-group** statement at the **[edit services l2tp]** hierarchy level:

```
tunnel-group group-name {
  hello-interval seconds;
  hide-avps;
  l2tp-access-profile profile-name;
  local-gateway address address;
  maximum-send-window packets;
  ppp-access-profile profile-name;
  receive-window packets;
  retransmit-interval seconds;
  service-interface interface-name;
  syslog {
    host hostname {
      services severity-level;
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
  tunnel-timeout seconds;
}
```



NOTE: If you delete a tunnel group or mark it inactive, all L2TP sessions in that tunnel group are terminated. If you change the value of the **local-gateway address** or the **service-interface** statement, all L2TP sessions using those settings are terminated. If you change or delete other statements at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level, new tunnels you establish will use the updated values but existing tunnels and sessions are not affected.

This following sections explain how to configure L2TP tunnel groups:

- Configuring Access Profiles for L2TP Tunnel Groups on page 315
- Configuring the Local Gateway Address and PIC on page 315
- Configuring Window Size for L2TP Tunnels on page 316
- Configuring Timers for L2TP Tunnels on page 316

Copyright © 2013, Juniper Networks, Inc.

- Hiding Attribute-Value Pairs for L2TP Tunnels on page 316
- Configuring System Logging of L2TP Tunnel Activity on page 317

Configuring Access Profiles for L2TP Tunnel Groups

To validate L2TP connections and session requests, you set up access profiles by configuring the **profile** statement at the **[edit access]** hierarchy level. You need to configure two types of profiles:

- L2TP tunnel access profile, which validates all L2TP connection requests to the specified local gateway address
- PPP access profile, which validates all PPP session requests through L2TP tunnels established to the local gateway address

For more information on configuring the profiles, see the Junos OS System Basics Configuration Guide. A profile example is included in "Examples: Configuring L2TP Services" on page 324.

To associate the profiles with a tunnel group, include the **l2tp-access-profile** and **ppp-access-profile** statements at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

```
l2tp-access-profile profile-name;
ppp-access-profile profile-name;
```

Configuring the Local Gateway Address and PIC

When you configure an L2TP group, you must also define a local address for the L2TP tunnel connections and the AS PIC that processes the requests:

- To configure the local gateway IP address, include the **local-gateway** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  ```
  local-gateway address address;
  ```

- To configure the AS PIC, include the **service-interface** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  ```
  service-interface sp-fpc/pic/port;
  ```

You can optionally specify the logical unit number along with the service interface. If specified, the unit is used as a logical interface representing PPP sessions negotiated using this profile.



NOTE: If you change the local gateway address or the service interface configuration, all L2TP sessions using those settings are terminated.

Dynamic class-of-service (CoS) functionality is supported on L2TP LNS sessions or L2TP sessions with ATM VCs, as long as the L2TP session is configured to use an IQ2 PIC on the egress interface. For more information, see the Junos OS Class of Service Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

Configuring Window Size for L2TP Tunnels

You can configure the maximum window size for packet processing at each end of the L2TP tunnel:

- The receive window size limits the number of concurrent packets the server processes. By default, the maximum is 16 packets. To change the window size, include the **receive-window** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  receive-window *packets*;

- The maximum-send window size limits the other end's receive window size. The information is transmitted in the receive window size attribute-value pair. By default, the maximum is 32 packets. To change the window size, include the **maximum-send-window** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  maximum-send-window *packets*;

Configuring Timers for L2TP Tunnels

You can configure the following timer values that regulate L2TP tunnel processing:

- Hello interval—If the server does not receive any messages within a specified time interval, the router software sends a hello message to the tunnel's remote peer. By default, the interval length is 60 seconds. If you configure a value of 0, no hello messages are sent. To configure a different value, include the **hello-interval** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  hello-interval *seconds*;

- Retransmit interval—By default, the retransmit interval length is 30 seconds. To configure a different value, include the **retransmit-interval** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  retransmit-interval *seconds*;

- Tunnel timeout—If the server cannot send any data through the tunnel within a specified time interval, it assumes that the connection with the remote peer has been lost and deletes the tunnel. By default, the interval length is 120 seconds. To configure a different value, include the **tunnel-timeout** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

  tunnel-timeout *seconds*;

Hiding Attribute-Value Pairs for L2TP Tunnels

Once an L2TP tunnel has been established and the connection authenticated, information is encoded by means of attribute-value pairs. By default, this information is not hidden. To hide the attribute-value pairs once the shared secret is known, include the **hide-avps** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

hide-avps;

Copyright © 2013, Juniper Networks, Inc.

Configuring System Logging of L2TP Tunnel Activity

You can specify properties that control how system log messages are generated for L2TP services.

To configure interface-wide default system logging values, include the **syslog** statement at the **[edit services l2tp tunnel-group *group-name*]** hierarchy level:

```
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
  }
}
```

Configure the **host** statement with a hostname or IP address that specifies the system log target server. The hostname **local** directs system log messages to the Routing Engine. For external system log servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify only one system logging hostname.

Table 16 on page 317 lists the severity levels that you can specify in configuration statements at the **[edit services l2tp tunnel-group *group-name* syslog host *hostname*]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 16: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

We recommend setting the system logging severity level to **error** during normal operation. To monitor PIC resource usage, set the level to **warning**. To gather information about an intrusion attack when an intrusion detection system error is detected, set the level to

Copyright © 2013, Juniper Networks, Inc.

notice for a specific service set. To debug a configuration or log Network Address Translation (NAT) events, set the level to **info**.

For more information about system log messages, see the *Junos OS System Log Messages Reference*.

To use one particular facility code for all logging to the specified system log host, include the **facility-override** statement at the **[edit services l2tp tunnel-group *group-name* syslog host *hostname*]** hierarchy level:

```
facility-override facility-name;
```

The supported facilities include: **authorization**, **daemon**, **ftp**, **kernel**, **user**, and **local0** through **local7**.

To specify a text prefix for all logging to this system log host, include the **log-prefix** statement at the **[edit services l2tp tunnel-group *group-name* syslog host *hostname*]** hierarchy level:

```
log-prefix prefix-text;
```

Configuring the Identifier for Logical Interfaces that Provide L2TP Services

You can configure L2TP services on adaptive services interfaces on M7i, M10i, and M120 routers only. You must configure the logical interface to be dedicated or shared. If a logical interface is dedicated, it can represent only one session at a time. A shared logical interface can have multiple sessions.

To configure the logical interface, include the **l2tp-interface-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* dial-options]** hierarchy level:

```
l2tp-interface-id name;
(dedicated | shared);
```

The **l2tp-interface-id** name configured on the logical interface must be replicated at the **[edit access profile *name*]** hierarchy level:

- For a user-specific identifier, include the **l2tp-interface-id** statement at the **[edit access profile *name* ppp]** hierarchy level.
- For a group identifier, include the **l2tp-interface-id** statement at the **[edit access profile *name* l2tp]** hierarchy level.

You can configure multiple logical interfaces with the same interface identifier, to be used as a pool for several users. For more information on configuring access profiles, see the Junos OS System Basics Configuration Guide.



NOTE: If you delete the dial-options statement settings configured on a logical interface, all L2TP sessions running on that interface are terminated.

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring Multilink PPP on a Shared Logical Interface

Multilink PPP is supported on either shared or dedicated logical interfaces. The following example can be used to configure many multilink bundles on a single shared interface:

```
interfaces {
    sp-1/3/0 {
        traceoptions {
            flag all;
        }
        unit 0 {
            family inet;
        }
        unit 20 {
            dial-options {
                l2tp-interface-id test;
                shared;
            }
            family inet;
        }
    }
}
access {
    profile t {
        client test {
            l2tp {
                interface-id test;
                multilink;
                shared-secret "$9$n8HX6A01RhlvL1R"; # SECRET-DATA
            }
        }
    }
    profile u {
        authentication-order radius;
    }
    radius-server {
        192.168.65.63 {
            port 1812;
            secret "$9$Vyb4ZHkPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
        }
    }
}
services {
    l2tp {
        tunnel-group 1 {
            l2tp-access-profile t;
            ppp-access-profile u;
            local-gateway {
                address 10.70.1.1;
            }
            service-interface sp-1/3/0;
        }
        traceoptions {
            flag all;
            debug-level packet-dump;
```

Copyright © 2013, Juniper Networks, Inc.

```
            filter {
                protocol l2tp;
                protocol ppp;
                protocol radius;
            }
        }
    }
}
```

AS PIC Redundancy for L2TP Services

L2TP services support AS PIC redundancy. To configure redundancy, you specify a redundancy services PIC (**rsp**) interface in which the primary AS PIC is active and a secondary AS PIC is on standby. If the primary AS PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary AS PIC is restored, it remains in standby and does not preempt the secondary AS PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the **show interfaces redundancy** command.



NOTE: On L2TP, the only service option supported is *warm standby*, in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected. The tunnels and sessions are torn down upon switchover and need to be restarted by the LAC and PPP client, respectively. However, configuration is preserved and available on the new active PIC, although the protocol state needs to be reestablished.

As with the other AS PIC services that support warm standby, you can issue the **request interfaces (revert | switchover)** command to manually switch between primary and secondary L2TP interfaces.

For more information, see "Configuring AS or Multiservices PIC Redundancy" on page 522. For an example configuration, see "Examples: Configuring L2TP Services" on page 324. For information on operational mode commands, see the Junos OS Operational Mode Commands.

IP Packet Fragment Reassembly for L2TP Overview

You can configure the service interfaces on the MX Series routers with modular port concentrators (MPCs) to support reassemble fragmented IP packets for an L2TP connection. When packets are transmitted over an L2TP connection, the packets may be fragmented during transmission and need to be reassembled before they are processed further. Efficient reassembly is important for network throughput, scalability, and graceful response to congestion.

Fragmentation of IP packets for transmission and the need to reassemble the IP packets at a destination is a feature of how Layer 2 (the frame layer) and Layer 3 (the packet layer) operate. The maximum size of a frame, set by the Maximum Transmission Unit (MTU) value, and the maximum size of a packet are determined independently. Typically

Copyright © 2013, Juniper Networks, Inc.

the packet size can far exceed the MTU size defined for the outgoing connection. If the packet size (data plus IP and other headers) exceeds the configured frame size (usually set by the transport medium limits), the packet must be fragmented and split across multiple frames for transmission. Frames are always processed immediately, when they arrive (if error-free), but packet fragments cannot be processed until the whole packet has been reassembled. Each packet fragment inside a frame series, except the last packet fragment, has the more fragments (MF) IP header bit set, indicating that this packet is part of a whole. The last packet fragment inside a frame does not have this MF bit set and therefore ends the fragment sequence. After all of the fragments of a packet have arrived, the entire packet can be reassembled.

In an L2TP connection, packets are transmitted between the L2TP Access Concentrator (LAC) and the L2TP Network Server (LNS). For an IP packet being transmitted over an L2TP connection, the packet is fragmented at the LAC, at an LNS, or at any intermediate router. IP reassembly parameters configured on the service interfaces of the LAC and the LNS determine how the fragments are reassembled at the service interfaces to ensure efficient reassembly over an L2TP connection.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 321
- Protocols and Applications Supported by MX240, MX480, MX960 Enhanced MPCs (MPCEs)
- ip-reassembly

Configuring IP Inline Reassembly for L2TP

This procedure shows how to configure a service interface on a LAC or LNS to reassembly fragmented IP packets. This example creates a service set that configures the IP reassembly parameters for L2TP fragments. The service set is then associated with the L2TP service.

Before you configure inline IP reassembly, be sure you have:

- Configured L2TP.
- Configured a valid service interface on the LAC or LNS.

To configure inline IP reassembly:

1. Configure the chassis-level bandwidth used by the inline services interface on the FPC and PIC slot for inline IP fragment reassembly.

   ```
   [edit chassis]
   user@host# set fpc 2 pic 1 inline-services bandwidth 10g
   ```

2. Configure the interface-level logical unit used by the inline services (**si-**) interface on the FPC and PIC slot for inline IP fragment reassembly.

   ```
   [edit interfaces]
   user@host# set si-2/1/0 unit 0 family inet
   user@host# set si-2/1/0 unit 0 service-domain inside
   ```

Copyright © 2013, Juniper Networks, Inc.

NOTE: This configuration is not unique to L2TP. However, you must configure the family (**inet**) and service domain (**inside**) as shown.

3. Configure the service set (**set1**) for IP reassembly in the input match direction. (The **local** option loops the reassembled packets back to the local interface.)

```
[edit services]
user@host# set service-set set1

[edit services service set ip-reassembly-set]
user@host# set ip-reassembly-rules ipr_rule1
user@host# set next-hop-service inside-service-interface si-9/1/0.0
user@host# set next-hop-service outside-service-interface-type local
```

NOTE: You must configure both inside (**si-** interface) and outside type (**local**) service interfaces statements. The reassembly rule is not formulated outside of the service set; this statement simply initiates the reassembly process.

4. Configure the IP reassembly rule parameter

```
[edit services ip-reassembly]
user@host# set rule ipr_rule1 match-direction input;
```

5. Configure the service set (**set1**) for IP reassembly to bind to the L2TP service.

NOTE:

- The service set must be defined at the **[edit services]** hierarchy level.
- You cannot delete a service set instance if it is associated with an L2TP service.

```
[edit services l2tp]
user@host# set ip-reassembly service-set set1
```

Related Documentation

- IP Packet Fragment Reassembly for L2TP Overview on page 320
- Protocols and Applications Supported by MX240, MX480, MX960 Enhanced MPCs (MPCEs)
- ip-reassembly

Tracing L2TP Operations

Tracing operations track all AS PIC operations and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/l2tpd**.

Copyright © 2013, Juniper Networks, Inc.



NOTE: This topic refers to tracing L2TP LNS operations on M Series routers. To trace L2TP LAC operations on MX Series routers, see Tracing L2TP Operations for Subscriber Access.

To trace L2TP operations, include the **traceoptions** statement at the **[edit services l2tp]** hierarchy level:

```
traceoptions {
    debug-level level;
    file <filename> <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
    filter {
        protocol name;
        user-name username;
    }
    flag flag;
    interfaces interface-name {
        debug-level severity;
        flag flag;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

You can specify the following L2TP tracing flags:

- **all**—Trace everything.
- **configuration**—Trace configuration events.
- **protocol**—Trace routing protocol events.
- **routing-socket**—Trace routing socket events.
- **rpd**—Trace routing protocol process events.

You can specify a trace level for PPP, L2TP, RADIUS, and User Datagram Protocol (UDP) tracing. To configure a trace level, include the **debug-level** statement at the **[edit services l2tp traceoptions]** hierarchy level and specify one of the following values:

- **detail**—Detailed debug information
- **error**—Errors only
- **packet-dump**—Packet decoding information

You can filter by protocol. To configure filters, include the **filter protocol** statement at the **[edit services l2tp traceoptions]** hierarchy level and specify one or more of the following protocol values:

- **ppp**
- **l2tp**

Copyright © 2013, Juniper Networks, Inc.

- **radius**
- **udp**

To implement filtering by protocol name, you must also configure either **flag protocol** or **flag all**.

You can also configure traceoptions for L2TP on a specific adaptive services interface. To configure per-interface tracing, include the **interfaces** statement at the **[edit services l2tp traceoptions]** hierarchy level:

```
interfaces interface-name {
    debug-level level;
    flag flag;
}
```



NOTE: Implementing traceoptions consumes CPU resources and affects the packet processing performance.

You can specify the **debug-level** and **flag** statements for the interface, but the options are slightly different from the general L2TP traceoptions. You specify the debug level as **detail**, **error**, or **extensive**, which provides complete PIC debug information. The following flags are available:

- **all**—Trace everything.
- **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
- **packet-dump**—Dump each packet's content based on debug level.
- **protocol**—Trace L2TP, PPP, and multilink handling.
- **system**—Trace packet processing on the PIC.

Examples: Configuring L2TP Services

Configure L2TP with multiple group and user profiles and a pool of logical interfaces for concurrent tunnel sessions:

```
[edit access]
address-pool customer_a {
    address 10.1.1.1/32;
}
address-pool customer_b {
    address-range low 10.2.2.1 high 10.2.3.2;
}
group-profile sunnyvale_users {
    ppp {
        framed-pool customer_a;
        idle-timeout 15;
        primary-dns 192.168.65.1;
        secondary-dns 192.168.65.2;
        primary-wins 192.168.65.3;
```


Copyright © 2013, Juniper Networks, Inc.

```
        secondary-wins 192.168.65.4;
        interface-id west;
    }
}
group-profile eastcoast_users {
    ppp {
        framed-pool customer_b;
        idle-timeout 20;
        primary-dns 192.168.65.5;
        secondary-dns 192.168.65.6;
        primary-wins 192.168.65.7;
        secondary-wins 192.168.65.8;
        interface-id east;
    }
}
group-profile sunnyvale_tunnel {
    l2tp {
        maximum-sessions-per-tunnel 100;
        interface-id west_shared;
    }
}
group-profile east_tunnel {
    l2tp {
        maximum-sessions-per-tunnel 125;
        interface-id east_shared;
    }
}
profile sunnyvale_bldg_1 {
    client white {
        chap-secret "$9$3s2690IeK8X7VKM7VwgaJn/Ctu1hclv87Ct87"; #       SECRET-DATA
        ppp {
            idle-timeout 22;
            primary-dns 192.168.65.1;
            framed-ip-address 10.12.12.12/32;
            interface-id east;
        }
        group-profile sunnyvale_users;
    }
    client blue {
        chap-secret "$9$eq1KWxbwgZUHNdjqmTF3uO1Rhr-dsoJDNd"; #       SECRET-DATA
        group-profile sunnyvale_users;
    }
    authentication-order password;
}
profile sunnyvale_bldg_1_tunnel {
    client test {
        l2tp {
            shared-secret "$9$r3HKvLg4ZUDkX7JGjif5p0BIRS8LN"; # SECRET-DATA
            maximum-sessions-per-tunnel 75;
            interface-id west_shared;
            ppp-authentication chap;
        }
        group-profile sunnyvale_tunnel;
    }
    client production {
        l2tp {
```

```
        shared-secret
          "$9$R2QErv8X-goGylVwg4jiTz36/t0BEleWFnRhrlXxbs2aJDHqf3nCP5";
        ppp-authentication chap;
      }
      group-profile sunnyvale_tunnel;
    }
}
[edit services]
l2tp {
  tunnel-group finance-lns-server {
    l2tp-access-profile sunnyvale_bldg_1_tunnel;
    ppp-access-profile sunnyvale_bldg_1;
    local-gateway {
      address 10.1.117.3;
    }
    service-interface sp-1/3/0;
    receive-window 1500;
    maximum-send-window 1200;
    retransmit-interval 5;
    hello-interval 15;
    tunnel-timeout 55;
  }
  traceoptions {
    flag all;
  }
}
[edit interfaces sp-1/3/0]
unit0 {
  family inet;
}
unit 10 {
  dial-options {
    l2tp-interface-id foo-user;
    dedicated;
  }
  family inet;
}
unit 11 {
  dial-options {
    l2tp-interface-id east;
    dedicated;
  }
  family inet;
}
unit 12 {
  dial-options {
    l2tp-interface-id east;
    dedicated;
  }
  family inet;
}
unit 21 {
  dial-options {
    l2tp-interface-id west;
    dedicated;
  }
```


Copyright © 2013, Juniper Networks, Inc.

```
    family inet;
  }
  unit 30 {
    dial-options {
      l2tp-interface-id west_shared;
      shared;
    }
    family inet;
  }
  unit 40 {
    dial-options {
      l2tp-interface-id east_shared;
      shared;
    }
    family inet;
  }
```

Configure L2TP redundancy:

```
  interfaces {
    rsp0 {
      redundancy-options {
        primary sp-0/0/0;
        secondary sp-1/3/0;
      }
      unit 0 {
        family inet;
      }
      unit 11 {
        dial-options {
          l2tp-interface-id east_shared;
          shared;
        }
        family inet;
      }
    }
  }
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 17

Summary of Layer 2 Tunneling Protocol Configuration Statements

The following sections explain each of the Layer 2 Tunneling Protocol (L2TP) statements. The statements are organized alphabetically.

facility-override

Syntax facility-override *facility-name*;

Hierarchy Level [edit services l2tp tunnel-group *group-name* syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Override the default facility for system log reporting.

Options ***facility-name***—Name of the facility that overrides the default assignment. Valid entries include:

- **authorization**
- **daemon**
- **ftp**
- **kernel**
- **local0** through **local7**
- **user**

Usage Guidelines See "Configuring System Logging of L2TP Tunnel Activity" on page 317.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the keepalive timer for L2TP tunnels.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options ***seconds***—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of **0** means that no hello messages are sent.

Default: 60 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Timers for L2TP Tunnels on page 316
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

Copyright © 2013, Juniper Networks, Inc.

hide-avps

Syntax hide-avps;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Hide L2TP attribute-value pairs if the secret shared between the two ends of the tunnel is known.



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Default Attribute-value pairs that can be hidden are exposed, even if the secret information is known.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Hiding Attribute-Value Pairs for L2TP Tunnels on page 316

host (L2TP)

Syntax
```
host hostname {
  services severity-level;
  facility-override facility-name;
  log-prefix prefix-value;
}
```

Hierarchy Level [edit services l2tp tunnel-group *group-name* syslog]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the hostname for the system logging utility.

Options ***hostname***—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.

The remaining statements are explained separately.

Usage Guidelines See "Configuring System Logging of L2TP Tunnel Activity" on page 317.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ip-reassembly (L2TP)

Syntax
```
ip-reassembly {
  service-set service-set-name;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 13.1.

Description Associate the reassembly service-set with the L2TP service.



NOTE: The service set must be defined at the [edit services] hierarchy level.

Options **service-set *service-set-name***—Identifies the service set to be associated with the L2TP service.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- IP Packet Fragment Reassembly for L2TP Overview on page 320
- Configuring IP Inline Reassembly for L2TP on page 321

l2tp-access-profile

Syntax l2tp-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the profile used to validate all L2TP connection requests to the local gateway address.

Options ***profile-name***—Identifier for the L2TP connection profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Access Profiles for L2TP Tunnel Groups on page 315
- Configuring an L2TP Access Profile on the LNS

Copyright © 2013, Juniper Networks, Inc.

local-gateway address

Syntax local-gateway address *address*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the local (LNS) IP address for L2TP tunnel.

Options ***address***—Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Local Gateway Address and PIC on page 315.
- Configuring L2TP Tunnel Groups on page 314
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

log-prefix (L2TP)

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit services l2tp tunnel-group *group-name* syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the system logging prefix value.

Options ***prefix-value***—System logging prefix value.

Usage Guidelines See "Configuring System Logging of L2TP Tunnel Activity" on page 317.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

maximum-send-window

Syntax maximum-send-window *packets*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the size of the send window for L2TP tunnels, which limits the remote end's receive window size.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options ***packets***—Maximum number of packets the send window can hold at one time.
Default: 32

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Window Size for L2TP Tunnels on page 316

ppp-access-profile

Syntax ppp-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the profile used to validate all Point-to-Point Protocol (PPP) session requests through L2TP tunnels established to the local gateway address.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options ***profile-name***—Identifier for the PPP profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Access Profiles for L2TP Tunnel Groups on page 315

Copyright © 2013, Juniper Networks, Inc.

receive-window

Syntax receive-window *packets*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the size of the receive window for L2TP tunnels, which limits the number of packets the server processes concurrently.



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options ***packets***—Maximum number of packets the receive window can hold at one time.
Default: 16

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Window Size for L2TP Tunnels on page 316

retransmit-interval (Services)

Syntax retransmit-interval *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the maximum retransmit interval for L2TP tunnels.



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options ***seconds***—Interval, in seconds, after which the server retransmits data if no acknowledgment is received.
Default: 30 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Timers for L2TP Tunnels on page 316

Copyright © 2013, Juniper Networks, Inc.

service-interface

Syntax service-interface *interface-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.
Option **si-*fpc/pic/port*** introduced in Junos OS Release 11.4.

Description Specify the service interface responsible for handling L2TP processing.



NOTE: On MX Series routers, the service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options *interface-name*—Name of the service interface. The interface type depends on the line card as follows:

- **sp-*fpc/pic/port***—On AS or Multiservices PICs on M7i, M10i, and M120 routers.
- **si-*fpc/pic/port***—On MPCs on MX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Local Gateway Address and PIC on page 315
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

Copyright © 2013, Juniper Networks, Inc.

services

See the following sections:

- services (Hierarchy) on page 337
- services (L2TP System Logging) on page 338

services (Hierarchy)

Syntax services l2tp { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service properties to be applied to traffic.

Options **l2tp**—Identifies the L2TP set of services statements.

Usage Guidelines See "L2TP Services Configuration Overview" on page 311.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

services (L2TP System Logging)

Syntax services *severity-level*;

Hierarchy Level [edit services l2tp tunnel-group *group-name* syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the system logging severity level.

Options ***severity-level***—Assigns a severity level to the facility. Valid entries include:

- **alert**—Conditions that should be corrected immediately.
- **any**—Matches any level.
- **critical**—Critical conditions.
- **emergency**—Panic conditions.
- **error**—Error conditions.
- **info**—Informational messages.
- **notice**—Conditions that require special handling.
- **warning**—Warning messages.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring System Logging of L2TP Tunnel Activity on page 317

Copyright © 2013, Juniper Networks, Inc.

syslog

Syntax

```
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
  }
}
```

Hierarchy Level [edit services l2tp tunnel-group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the generation of system log messages for L2TP services. System log information is passed to the kernel for logging in the **/var/log/l2tpd** directory.



NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options The remaining statements are described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring System Logging of L2TP Tunnel Activity on page 317

Copyright © 2013, Juniper Networks, Inc.

traceoptions (L2TP)

Syntax

```
traceoptions {
  debug-level level;
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    protocol name;
    user-name username;
  }
  flag flag;
  interfaces interface-name {
    debug-level level;
    flag flag;
  }
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define tracing operations for L2TP processes.

Options

debug-level *level*—Trace level for PPP, L2TP, RADIUS, and UDP; this option does not apply to L2TP on MX Series routers:

- **detail**—Trace detailed debug information.
- **error**—Trace error information.
- **packet-dump**—Trace packet decoding information.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

filter protocol *name*—Additional filter for the specified protocol; this option does not apply to L2TP on MX Series routers:

- **l2tp**
- **ppp**
- **radius**
- **udp**

Copyright © 2013, Juniper Networks, Inc.

filter user-name *username*—Additional filter for the specified username; this option does not apply to L2TP on MX Series routers.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **configuration**—Trace configuration events.
- **events**—Trace interface events.
- **general**—Trace general events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization.
- **ipc-rx**—Trace IPC receive events.
- **ipc-tx**—Trace IPC transmit events.
- **memory**—Trace memory management code.
- **message**—Trace message processing code.
- **packet-error**—Trace packet error events.
- **parse**—Trace parsing events.
- **protocol**—Trace L2TP events.
- **receive-packets**—Trace received L2TP packets.
- **routing-process**—Trace routing process interactions.
- **routing-socket**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **states**—Trace state machine events.
- **timer**—Trace timer events.
- **transmit-packets**—Trace transmitted L2TP packets.
- **tunnel**—Trace tunnel events.

Copyright © 2013, Juniper Networks, Inc.

interfaces ***interface-name***—Apply L2TP traceoptions to a specific services interface. This option does not apply to L2TP on MX Series routers.

- **debug-level** ***level***—Trace level for the interface; this option does not apply to L2TP on MX Series routers:
  - **detail**—Trace detailed debug information.
  - **error**—Trace error information.
  - **extensive**—Trace all PIC debug information.
- **flag** ***flag***—Tracing operation to perform for the interface. This option does not apply to L2TP on MX Series routers. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:
  - **all**—Trace everything.
  - **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
  - **packet-dump**—Dump each packet content based on debug level.
  - **protocol**—Trace L2TP, PPP, and multilink handling.
  - **system**—Trace packet processing on the PIC.

level—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

Copyright © 2013, Juniper Networks, Inc.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Tracing L2TP Operations on page 322
- Tracing L2TP Operations for Subscriber Access

Copyright © 2013, Juniper Networks, Inc.

tunnel-group

Syntax

```
tunnel-group group-name {
  aaa-access-profile profile-name;
  dynamic-profile profile-name;
  hello-interval seconds;
  hide-avps;
  l2tp-access-profile profile-name;
  local-gateway address address;
  maximum-send-window packets;
  ppp-access-profile profile-name;
  receive-window packets;
  retransmit-interval seconds;
  service-device-pool pool-name;
  service-interface interface-name;
  syslog {
    host hostname {
      services severity-level;
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
  tos-reflect;
  tunnel-timeout seconds;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before Junos OS Release 7.4.
Support for MX Series routers introduced in Junos OS Release 11.4

Description Specify the L2TP tunnel properties.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *group-name*—Identifier for the tunnel group.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- (M71, M10i, and M120 routers) Configuring L2TP Tunnel Groups on page 314
- MX Series routers) Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

Copyright © 2013, Juniper Networks, Inc.

tunnel-timeout

Syntax tunnel-timeout *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the maximum downtime for an L2TP tunnel, after which the tunnel is terminated because the connection is presumed to have been lost.

Options ***seconds***—Interval after which the tunnel is terminated if no data can be sent.

Default: 120 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Timers for L2TP Tunnels on page 316
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 18

Link Services IQ Interfaces Configuration Guidelines

You can configure link services intelligent queuing (IQ) (LSQ or **lsq-**) interfaces on the Adaptive Services (AS) PIC, the internal Adaptive Services Module in the M7i platform, the Link Services II PIC, and the Multiservices PIC. LSQ interfaces are similar to link services interfaces, which are described in "Multilink and Link Services Logical Interface Configuration Overview" on page 1136. The important difference is that LSQ interfaces fully support Junos class of service (CoS) components.

The AS or Multiservices PIC has a limit of 1023 logical interfaces. Each logical interface is a Multilink Point-to-Point Protocol (MLPPP) bundle, an FRF.15 bundle, or an FRF.16 DLCI.

This chapter describes the Layer 2 service package and the CoS and failure recovery capabilities of LSQ interfaces. For detailed information about Layer 3 services, see other chapters in this manual and the Junos OS Feature Guides.



NOTE: **The Link Services II PIC offers the same functionality as the Layer 2 service package on AS or Multiservices PICs.**

This chapter contains the following sections:

- Layer 2 Service Package Capabilities and Interfaces on page 348
- Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 350
- Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 352
- Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 352
- Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 361
- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 365
- Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 367
- Configuring Multiclass MLPPP on LSQ Interfaces on page 367
- Oversubscribing Interface Bandwidth on LSQ Interfaces on page 369

Copyright © 2013, Juniper Networks, Inc.

- Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 374
- Configuring Link Services and CoS on Services PICs on page 377
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 380
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 385
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 390
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 395
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 on page 402
- Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 403
- Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 404
- Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 406

Layer 2 Service Package Capabilities and Interfaces

As described in "Enabling Service Packages" on page 39, you can configure the AS or Multiservices PIC and the internal ASM in the M7i platform to use either the Layer 2 or the Layer 3 service package.

When you enable the Layer 2 service package, the AS or Multiservices PIC supports *link services*. On the AS or Multiservices PIC and the ASM, link services include the following:

- Junos CoS components—"Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 361 describes how the Junos CoS components work on link services IQ (**lsq**) interfaces. For detailed information about Junos CoS components, see the Junos OS Class of Service Configuration Guide.
- Data compression using the compressed Real-Time Transport Protocol (CRTP) for use in voice over IP (VoIP) transmission.



NOTE: On LSQ interfaces, all multilink traffic for a single bundle is sent to a single processor. If CRTP is enabled on the bundle, it adds overhead to the CPU. Because T3 network interfaces support only one link per bundle, make sure you configure a fragmentation map for compressed traffic on these interfaces and specify the **no-fragmentation** option. For more information, see "Configuring Delay-Sensitive Packet Interleaving" on page 424 and "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

- Link fragment interleaving (LFI) on Frame Relay links using FRF.12 end-to-end fragmentation—The standard for FRF.12 is defined in the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*.
- LFI on Multilink Point-to-Point Protocol (MLPPP) links.

Copyright © 2013, Juniper Networks, Inc.

- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—The standard for FRF.15 is defined in the specification FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*.
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16)—The standard for FRF.16 is defined in the specification FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*.
- MLPPP—The standard for MLPPP is defined in the specification RFC 1990, *The PPP Multilink Protocol (MP)*.
- Multiclass extension to MLPPP—The standard is defined in the specification RFC 2686, *The Multi-Class Extension to Multi-Link PPP*.

For the LSQ interface on the AS or Multiservices PIC, the configuration syntax is almost the same as for Multilink and Link Services PICs. The primary difference is the use of the interface-type descriptor **lsq** instead of **ml** or **ls**. When you enable the Layer 2 service package on the AS or Multiservices PIC, the following interfaces are automatically created:

gr-*fpc/pic/port*
ip-*fpc/pic/port*
lsq-*fpc/pic/port*
lsq-*fpc/pic/port*:0
...
lsq-*fpc/pic/port*:*N*
mt-*fpc/pic/port*
pd-*fpc/pic/port*
pe-*fpc/pic/port*
sp-*fpc/pic/port*
vt-*fpc/pic/port*

Interface types **gr**, **ip**, **mt**, **pd**, **pe**, and **vt** are standard tunnel interfaces that are available on the AS or Multiservices PIC whether you enable the Layer 2 or the Layer 3 service package. These tunnel interfaces function the same way for both service packages, except that the Layer 2 service package does not support some tunnel functions, as shown in Table 5 on page 24. For more information about tunnel interfaces, see Tunnel Properties.



NOTE: Interface type **sp** is created because it is needed by the Junos OS. For the Layer 2 service package, the **sp** interface is not configurable, but you should not disable it.

Interface type **lsq-*fpc/pic/port*** is the physical link services IQ interface (**lsq**). Interface types **lsq-*fpc/pic/port*:0** through **lsq-*fpc/pic/port*:*N*** represent FRF.16 bundles. These interface types are created when you include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level. For more information, see "Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 361.

Copyright © 2013, Juniper Networks, Inc.

NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the router does not use the bandwidth value if the interfaces are included in an MLPPP or MLFR bundle. The bandwidth is calculated internally according to the time slots, framing, and byte-encoding of the interface. For more information about these properties, see the Junos® OS Network Interfaces.

Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS

Link services IQ (**lsq-**) interfaces that are paired with SONET PICs can use the Automatic Protection Switching (APS) configuration already available on SONET networks to provide failure recovery. SONET APS provides stateless failure recovery, if it is configured on SONET interfaces in separate chassis and each SONET PIC is paired with an AS or Multiservices PIC in the same chassis. If one of the following conditions for APS failure is met, the associated SONET PIC triggers recovery to the backup circuit and its associated AS or Multiservices PIC. The failure conditions are:

- Failure of Link Services IQ PIC
- Failure of FPC that hosts the Link Services IQ PIC
- Failure of Packet Forwarding Engine
- Failure of chassis

The guidelines for configuring SONET APS are described in the Junos® OS Network Interfaces.

The following sections describe how to configure failover properties:

- Configuring the Association between LSQ and SONET Interfaces on page 350
- Configuring SONET APS Interoperability with Cisco Systems FRF.16 on page 351
- Restrictions on APS Redundancy for LSQ Interfaces on page 352

Configuring the Association between LSQ and SONET Interfaces

To configure the association between AS or Multiservices PICs hosting link services IQ interfaces and the SONET interfaces, include the **lsq-failure-options** statement at the **[edit interfaces]** hierarchy level:

```
lsq-fpc/pic/port {
  lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
  }
}
```

For example, consider the following network scenario:

- Primary router includes interfaces **oc3-0/2/0** and **lsq-1/1/0**.
- Backup router includes interfaces **oc3-2/2/0** and **lsq-3/2/0**.

Copyright © 2013, Juniper Networks, Inc.

Configure SONET APS, with **oc3-0/2/0** as the working circuit and **oc3-2/2/0** as the protect circuit. Include the **trigger-link-failure** statement to extend failure to the LSQ PICs:

```
interfaces lsq-1/1/0 {
  lsq-failure-options {
    trigger-link-failure oc3-0/2/0;
  }
}
```



NOTE: You must configure the **lsq-failure-options** statement on the primary router only. The configuration is not supported on the backup router.

To inhibit the router from sending PPP termination-request messages to the remote host if the Link Services IQ PIC fails, include the **no-termination-request** statement at the **[edit interfaces lsq-*fpc/pic/port* lsq-failure-options]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port lsq-failure-options]
no-termination-request;
```

This functionality is supported on link PICs as well. To inhibit the router from sending PPP termination-request messages to the remote host if a link PIC fails, include the **no-termination-request** statement at the **[edit interfaces *interface-name* ppp-options]** hierarchy level.

```
[edit interfaces interface-name ppp-options]
no-termination-request;
```

The **no-termination-request** statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only, on the following PICs:

- Channelized OC3 IQ PICs
- Channelized OC12 IQ PICs
- Channelized STM1 IQ PICs
- Channelized STM4 IQ PICs

Configuring SONET APS Interoperability with Cisco Systems FRF.16

Juniper Networks routers configured with APS might not interoperate correctly with Cisco FRF.16. To enable interoperation, include the **cisco-interoperability** statement at the **[edit interfaces lsq-*fpc/pic/port* mlfr-uni-nni-bundle-options]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port mlfr-uni-nni-bundle-options]
cisco-interoperability send-lip-remove-link-for-link-reject;
```

The **send-lip-remove-link-for-link-reject** option prompts the router to send a Link Integrity Protocol remove link when it receives an add-link rejection message.

Copyright © 2013, Juniper Networks, Inc.

Restrictions on APS Redundancy for LSQ Interfaces

The following restrictions apply to LSQ failure recovery:

- It applies only to Link Services IQ PICs installed in M Series routers, except for M320 routers.
- You must configure the **failure-options** statement on physical LSQ interfaces, not on MLFR channelized units.
- The Link Services IQ PICs must be associated with SONET link PICs. The paired PICs can be installed on different routers or in the same router; in other words, both interchassis and intrachassis recovery are supported
- Failure recovery is stateless; as a result, route flapping and loss of link state is expected in interchassis recovery, requiring PPP renegotiation. In intrachassis recovery, no impact on traffic is anticipated with Routing Engine failover, but PIC failover results in PPP renegotiation.
- The switchover is not revertive: when the original hardware is restored to service, traffic does not automatically revert back to it.
- Normal APS switchover and PIC-triggered APS switchover can be distinguished only by checking the system log messages.



NOTE: When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG_ERR is generated. This mechanism applies to both Layer 2 and Layer 3 service packages.

Configuring LSQ Interface Redundancy in a Single Router Using SONET APS

Stateless switchover from one Link Services IQ PIC to another within the same router can be configured by using the SONET APS mechanism described in "Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS" on page 350. Each Link Services IQ PIC must be associated with a specified SONET link PIC within the same router.



NOTE: For complete intrachassis recovery, including recovery from Routing Engine failover, graceful Routing Engine switchover (GRES) must be enabled on the router. For more information, see the Junos OS System Basics Configuration Guide.

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

You can configure failure recovery on M Series, MX Series, and T Series routers that have multiple AS or Multiservices PICs and DPCs with **lsq-** interfaces by specifying a virtual

Copyright © 2013, Juniper Networks, Inc.

LSQ redundancy (**rlsq**) interface in which the primary Link Services IQ PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all LSQ processing is transferred to it. To determine which PIC is currently active, issue the **show interfaces redundancy** command.



NOTE: This configuration does not require the use of SONET APS for failover. Network interfaces that do not support SONET can be used, such as T1 or E1 interfaces.

The following sections provide more information:

- Configuring Redundant Paired LSQ Interfaces on page 353
- Restrictions on Redundant LSQ Interfaces on page 354
- Configuring Link State Replication for Redundant Link PICs on page 356
- Examples: Configuring Redundant LSQ Interfaces for Failure Recovery on page 357

Configuring Redundant Paired LSQ Interfaces

The physical interface type **rlsq** specifies the pairings between primary and secondary **lsq** interfaces to enable redundancy. To configure a backup **lsq** interface, include the **redundancy-options** statement at the **[edit interfaces rlsq*number*]** hierarchy level:

```
[edit interfaces rlsqnumber]
redundancy-options {
   (hot-standby | warm-standby);
   primary lsq-fpc/pic/port;
   secondary lsq-fpc/pic/port;
}
```

For the **rlsq** interface, ***number*** can be from 0 through 1023. If the primary **lsq** interface fails, traffic processing switches to the secondary interface. The secondary interface remains active even after the primary interface recovers. If the secondary interface fails and the primary interface is active, processing switches to the primary interface.

The **hot-standby** option is used with one-to-one redundancy configurations, in which one working PIC is supported by one backup PIC. It is supported with MLPPP, CRTP, FRF.15, and FRF.16 configurations for the LSQ interface to achieve an uninterrupted LSQ service. It sets the requirement for the failure detection and recovery time to be less than 5 seconds. The behavior is revertive, but you can manually switch between the primary and secondary PICs by issuing the **request interfaces (revert | switchover) rlsq*number*** operational mode command. It also provides a switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16.

The **warm-standby** option is used with redundancy configurations in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected.

Certain combinations of **hot-standby** and **warm-standby** configuration are not permitted and result in a configuration error. The following examples are permitted:

Copyright © 2013, Juniper Networks, Inc.

- Interface **rlsq0** configured with **primary lsq-0/0/0** and **warm-standby**, in combination with interface **rlsq0:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:0**, in combination with interface **rlsq0:1** configured with **primary lsq-0/0/0:1**

The following example combinations are not permitted:

- Interface **rlsq0** configured with **primary lsq-0/0/0** and **hot-standby**, in combination with interface **rlsq0:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:0**, in combination with interface **rlsq1:0** configured with **primary lsq-0/0/0:0**
- Interface **rlsq0:0** configured with **primary lsq-0/0/0:1**, in combination with interface **rlsq1:1** configured with **primary lsq-0/0/0:1**
- Interface **rlsq0** configured with **primary lsq-0/0/0**, in combination with interface **rlsq1** configured with **primary lsq-0/0/0**

In addition, the same physical interface cannot be reused as the primary interface for more than one **rlsq** interface, nor can any of the associated logical interfaces. For example, primary interface **lsq-0/0/0** cannot be reused in another **rlsq** interface as **lsq-0/0/0:0**.

Restrictions on Redundant LSQ Interfaces

Link Services IQ PIC failure occurs under the following conditions:

- The primary PIC fails to boot. In this case, the **rlsq** interface does not come up and manual intervention is necessary to reboot or replace the PIC, or to rename the primary PIC to the secondary one in the **rlsq** configuration.
- When configuring an **rlsq** interface, ensure that:
  - The unit number allocated to the **rlsq** interface is less than the number of Multilink Frame Relay user-to-network interface network-to-network interface (UNI-NNI) (FRF.16) bundles allocated on the Link Services PIC.
  - Data-link connection identifier (DLCI) is configured for the **rlsq** interface.

  If these conditions are not met, the **rlsq** interface does not boot. When you issue the **show interfaces redundancy** command, the state of the **rlsq** interface is indicated as **Waiting for primary MS PIC**.
- The primary PIC becomes active and then fails. The secondary PIC automatically takes over processing.
- A failover to the secondary PIC takes place. The secondary PIC then fails. If the primary PIC has been restored to active state, processing switches to it.
- The FPC that contains the Link Services IQ PIC fails.

Copyright © 2013, Juniper Networks, Inc.

The following constraints apply to redundant LSQ configurations:

- We recommend that primary and secondary PICs be configured in two different FPCs (in chassis other than M10i routers).
- You cannot configure a Link Services IQ PIC with explicit bundle configurations and as a constituent of an **rlsq** interface.
- Redundant LSQ configurations provide full GRES support. (You must configure GRES at the **[edit chassis]** hierarchy level; see the Junos OS System Basics Configuration Guide.
- If you configure the **redundancy-options** statement with the **hot-standby** option, the configuration must include one **primary** interface value and one **secondary** interface value.
- Since the same interface name is used for **hot-standby** and **warm-standby**, if you modify the configuration to change this attribute, it is recommended that you first deactivate the interface, commit the new configuration, and then reactivate the interface.
- You cannot make changes to an active **redundancy-options** configuration. You must deactivate the **rlsq*number*** interface configuration, change it, and reactivate it.
- The **rlsq*number*** configuration becomes active only if the primary interface is active. When the configuration is first activated, the primary interface must be active; if not, the **rlsq** interface waits until the primary interface comes up.
- You cannot modify the configuration of **lsq** interfaces after they have been included in an active **rlsq** interface.
- All the operational mode commands that apply to **rsp** interfaces also apply to **rlsq** interfaces. You can issue **show** commands for the **rlsq** interface or the primary and secondary **lsq** interfaces. However, statistics on the link interfaces are not carried over following a Routing Engine switchover.
- The **rlsq** interfaces also support the **lsq-failure-options** configuration, discussed in "Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS" on page 350. If the primary and secondary Link Services IQ PICs fail and the **lsq-failure-options** statement is configured, the configuration triggers a SONET APS switchover.
- Redundant LSQ configurations that require MLPPP Multilink Frame Relay (FRF.15 and FRF.16) are supported only with the **warm-standby** option.
- Redundant LSQ support is extended to ATM network interfaces.
- Channelized interfaces are used with FRF-16 bundles, for example **rlsq0:0**. The **rlsq** number and its constituents, the **primary** and **secondary** interfaces, must match for the configuration to be valid: either all must be channelized, or none. For an example of an FRF.16 configuration, see "Configuring LSQ Interface Redundancy for an FRF.16 Bundle" on page 361.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Adaptive Services and Multiservices PICs in layer-2 mode (running Layer 2 services) are not rebooted when a MAC flow-control situation is detected.

Configuring Link State Replication for Redundant Link PICs

Link state replication, also called *interface preservation*, is an addition to the SONET Automatic Protection Switching (APS) functionality that helps promote redundancy of the link PICs used in LSQ configurations.

Link state replication provides the ability to add two sets of links, one from the active (working) SONET PIC and the other from the backup (protect) SONET PIC to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without causing a link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation. For more information about SONET APS configurations, see the Junos® OS Network Interfaces.

To configure link state replication, include the **preserve-interface** statement at the **[edit interfaces *interface-name* sonet-options aps]** hierarchy level on both network interfaces:

```
edit interfaces interface-name sonet-options aps]
preserve-interface;
```

The following constraints apply to link PIC redundancy:

- APS functionality must be available on the SONET PICs and the interface configurations must be identical on both ends of the link. Any configuration mismatch causes the commit operation to fail.
- This feature is supported only with LSQ and SONET APS-enabled link PICs, including Channelized OC3, Channelized OC12, and Channelized STM1 intelligent queuing (IQ) PICs.
- Link state replication supports MLPPP and PPP over Frame Relay (**frame-relay-ppp**) encapsulation, and fully supports GRES.
- Enabling the interface or protocol traceoptions with a large number of MLPPP links can trigger Link Control Protocol (LCP) renegotiation during the link switchover time.

  NOTE: This renegotiation is more likely to take place for configurations with back-to-back Juniper Networks routers than in networks in which a Juniper Networks router is connected to an add/drop multiplexer (ADM).

- In general, networks that connect a Juniper Networks router to an ADM allow faster MLPPP link switchover than those with back-to-back Juniper Networks routers. The MLPPP link switchover time difference may be significant, especially for networks with a large number of MLPPP links.
- An aggressive LCP keepalive timeout configuration can lead to LCP renegotiation during the MLPPP link switchover. By default, the LCP keepalive timer interval is 10 seconds

Copyright © 2013, Juniper Networks, Inc.

and the consecutive link down count is 3. The MLPPP links start LCP negotiation only after a timeout of 30 seconds. Lowering these configuration values may trigger one or more of the MLPPP links to renegotiate during the switchover time.



NOTE: LCP renegotiation is more likely to take place for configurations with back-to-back Juniper Networks routers than in networks in which a Juniper Networks router is connected to an ADM.

As an example, the following configuration shows the link state replication configuration between the ports **coc3-1/0/0** and c**oc3-2/0/0**.

```
interfaces {
  coc3-1/0/0 {
    sonet-options {
      aps {
        preserve-interface;
        working-circuit aps-group-1;
      }
    }
  }
  coc3-2/0/0 {
    sonet-options {
      aps {
        preserve-interface;
        protect-circuit aps-group-1;
      }
    }
  }
}
```

Examples: Configuring Redundant LSQ Interfaces for Failure Recovery

Configuring LSQ Interface Redundancy for MLPPP

The following configuration shows that **lsq-1/1/0** and **lsq-1/3/0** work as a pair and the redundancy type is **hot-standby**, which sets the requirement for the failure detection and recovery time to be less than 5 seconds:

```
interfaces rlsq0 {
  redundancy-options {
    primary lsq-1/1/0;
    secondary lsq-1/3/0;
    hot-standby; #either hot-standby or warm-standby is supported
  }
}
```

The following example shows a related MLPPP configuration:



NOTE: MLPPP protocol configuration is required for this configuration.

Copyright © 2013, Juniper Networks, Inc.

```
interfaces {
  t1-/1/2/0 {
    unit 0 {
      family mlppp {
        bundle rlsq0.0;
      }
    }
  }
  rlsq0 {
    unit 0 {
      family inet {
        address 30.1.1.2/24;
      }
    }
  }
}
```

The following example shows a related CoS configuration:

```
class-of-service {
  interfaces {
    rlsq0 {
      unit * {
        fragmentation-maps fr-map1;
      }
    }
  }
}
```

The following example shows a complete link state replication configuration for MLPPP. This example uses two bundles, each with four T1 links. The first four T1 links (**t1-*:1** through **t1-*:4**) form the first bundle and the last four T1 links (**t1-*:5** through **t1-*:8**) form the second bundle. To minimize the duplication in the configuration, this example uses the **[edit groups]** statement; for more information, see the Junos OS System Basics Configuration Guide. This type of configuration is not required; it simplifies the task and minimizes duplication.

```
groups {
  ml-partition-group {
    interfaces {
      <coc3-*> {
        partition 1 oc-slice 1 interface-type coc1;
      }
      <coc1-*> {
        partition 1-8 interface-type t1;
      }
    }
  }
  ml-bundle-group-1 {
    interfaces {
      <t1-*:"[1-4]"> {
        encapsulation ppp;
        unit 0 {
          family mlppp {
            bundle lsq-0/1/0.0;
          }
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
      }
      ml-bundle-group-2 {
        interfaces {
          <t1-*:"[5-8]"> {
            encapsulation ppp;
            unit 0 {
              family mlppp {
                bundle lsq-0/1/0.1;
              }
            }
          }
        }
      }
    }
    interfaces {
      lsq-0/1/0 {
        unit 0 {
          encapsulation multilink-ppp;
          family inet {
            address 1.1.1.1/32 {
              destination 1.1.1.2;
            }
          }
        }
        unit 1 {
          encapsulation multilink-ppp;
          family inet {
            address 1.1.2.1/32 {
              destination 1.1.2.2;
            }
          }
        }
      }
      coc3-1/0/0 {
        apply-groups ml-partition-group;
        sonet-options {
          aps {
            preserve-interface;
            working-circuit aps-group-1;
          }
        }
      }
      coc1-1/0/0:1 {
        apply-groups ml-partition-group;
      }
      t1-1/0/0:1:1 {
        apply-groups ml-bundle-group-1;
      }
      t1-1/0/0:1:2 {
        apply-groups ml-bundle-group-1;
      }
      t1-1/0/0:1:3 {
        apply-groups ml-bundle-group-1;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    t1-1/0/0:1:4 {
      apply-groups ml-bundle-group-1;
    }
    t1-1/0/0:1:5 {
      apply-groups ml-bundle-group-2;
    }
    t1-1/0/0:1:6 {
      apply-groups ml-bundle-group-2;
    }
    t1-1/0/0:1:7 {
      apply-groups ml-bundle-group-2;
    }
    t1-1/0/0:1:8 {
      apply-groups ml-bundle-group-2;
    }
    coc3-2/0/0 {
      apply-groups ml-partition-group;
      sonet-options {
        aps {
          preserve-interface;
          protect-circuit aps-group-1;
        }
      }
    }
    coc1-2/0/0:1 {
      apply-groups ml-partition-group;
    }
    t1-2/0/0:1:1 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:2 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:3 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:4 {
      apply-groups ml-bundle-group-1;
    }
    t1-2/0/0:1:5 {
      apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:6 {
      apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:7 {
      apply-groups ml-bundle-group-2;
    }
    t1-2/0/0:1:8 {
      apply-groups ml-bundle-group-2;
    }
  }
```


Copyright © 2013, Juniper Networks, Inc.

Configuring LSQ Interface Redundancy for an FRF.15 Bundle

The following example shows a configuration for an FRF.15 bundle:

```
interfaces rlsq0 {
  redundancy-options {
    primary lsq-1/2/0;
    secondary lsq-1/3/0;
    warm-standby; #either hot-standby or warm-standby is supported
  }
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 30.1.1.1/24;
    }
  }
}
```

Configuring LSQ Interface Redundancy for an FRF.16 Bundle

The following example shows a configuration for an FRF.16 bundle:

```
interfaces rlsq0:0 {
  dce;
  encapsulation multilink-frame-relay-uni-nni;
  redundancy-options {
    primary lsq-1/2/0:0;
    secondary lsq-1/3/0:0;
    warm-standby; #either hot-standby or warm-standby is supported
  }
  unit 0 {
    dlci 1000;
    family inet {
      address 50.1.1.1/24;
    }
  }
}
```

Configuring CoS Scheduling Queues on Logical LSQ Interfaces

For link services IQ (**lsq-**) interfaces, you can specify a scheduler map for each logical unit. A logical unit represents either an MLPPP bundle or a DLCI configured on a FRF.16 bundle. The scheduler is applied to the traffic sent to an AS or Multiservices PIC running the Layer 2 link services package.

If you configure a scheduler map on a bundle, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-*fpc/pic/port*]** hierarchy level. If you configure a scheduler map on an FRF.16 DLCI, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-*fpc/pic/port*:*channel*]** hierarchy level. For more information, see the Junos OS Class of Service Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

If you need latency guarantees for multiclass or LFI traffic, you must use channelized IQ PICs for the constituent links. With non-IQ PICs, because queueing is not done at the channelized interface level on the constituent links, latency-sensitive traffic might not receive the type of service that it should. Constituent links from the following PICs support latency guarantees:

- Channelized E1 IQ PIC
- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC
- Channelized T3 IQ PIC

For scheduling queues on a logical interface, you can configure the following scheduler map properties at the **[edit class-of-service schedulers]** hierarchy level:

- **buffer-size**—The queue size; for more information, see "Configuring Scheduler Buffer Size" on page 363.
- **priority**—The transmit priority (low, high, strict-high); for more information, see "Configuring Scheduler Priority" on page 363.
- **shaping-rate**—The subscribed transmit rate; for more information, see "Configuring Scheduler Shaping Rate" on page 363.
- **drop-profile-map**—The random early detection (RED) drop profile; for more information, see "Configuring Drop Profiles" on page 364.

When you configure MLPPP and FRF.12 on M Series and T Series routers, you should configure a single scheduler with non-zero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ interface (**lsq**) and to each constituent link.

When you configure FRF.16 on M Series and T Series routers, you can assign a single scheduler map to the link services IQ interface (**lsq**) and to each link services IQ DLCI, or you can assign different scheduler maps to the various DLCIs of the bundle, as shown in "Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16" on page 388. For the constituent links of an FRF.16 bundle, you do not need to configure a custom scheduler. Because LFI and multiclass are not supported for FRF.16, the traffic from each constituent link is transmitted from queue 0. This means you should allow most of the bundle bandwidth to be used by queue 0. The default scheduler transmission rate and buffer size percentages for queues 0 through 3 are 95, 0, 0, and 5 percent, respectively. This default scheduler sends all user traffic to queue 0 and all network-control traffic to queue 3, and therefore it is well suited to the behavior of FRF.16. You can configure a custom scheduler that explicitly replicates the 95, 0, 0, and 5 percent queuing behaviors, and apply it to the constituent links.



NOTE: On T Series and M320 routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

Copyright © 2013, Juniper Networks, Inc.

For link services IQ interfaces (**lsq**), these scheduling properties work as they do in other PICs, except as noted in the following sections.



NOTE: On T Series and M320 routers, lsq interfaces do not support DiffServ code point (DSCP) and DSCP-IPv6 rewrite markers.

Configuring Scheduler Buffer Size

You can configure the scheduler buffer size in three ways: as a temporal value, as a percentage, and as a remainder. On a single logical interface (MLPPP or a FRF.16 DLCI), each queue can have a different buffer size.

If you specify a temporal value, the queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This number is computed by multiplying logical interface speed by the temporal value. For MLPPP bundles, logical interface speed is equal to the bundle bandwidth, which is the sum of constituent link speeds minus link-layer overhead. For MLFR FRF.16 DLCIs, logical interface speed is equal to bundle bandwidth multiplied by the DLCI shaping rate. In all cases, the maximum temporal value is limited to 200 milliseconds.

Buffer size percentages are implicitly converted into temporal values by multiplying the percentage by 200 milliseconds. For example, buffer size specified as **buffer-size percent 20** is the same as a 40-millisecond temporal delay. The link services IQ implementation guarantees 200 milliseconds of buffer delay for all interfaces with T1 and higher speeds. For slower interfaces, it guarantees one second of buffer delay.

The queueing algorithm evenly distributes leftover bandwidth among all queues that are configured with the **buffer-size remainder** statement. The queuing algorithm guarantees enough space in the transmit buffer for two MTU-sized packets.

Configuring Scheduler Priority

The transmit priority of each queue is determined by the scheduler and the forwarding class. Each queue receives a guaranteed amount of bandwidth specified with the scheduler **transmit-rate** statement.

Configuring Scheduler Shaping Rate

You use the shaping rate to set the percentage of total bundle bandwidth that is dedicated to a DLCI. For link services IQ DLCIs, only percentages are accepted, which allows adjustments in response to dynamic changes in bundle bandwidth—for example, when a link goes up or down. This means that absolute shaping rates are not supported on FRF.16 bundles. Absolute shaping rates are allowed for MLPPP and MLFR bundles only.

For scheduling between DLCIs in a MLFR FRF.16 bundle, you can configure a shaping rate for each DLCI. A shaping rate is expressed as a percentage of the aggregate bundle bandwidth. Shaping rate percentages for all DLCIs within a bundle can add up to 100 percent or less. Leftover bandwidth is distributed equally to DLCIs that do not have the **shaping-rate** statement included at the **[edit class-of-service interfaces lsq-*fpc/pic/port*:*channel* unit *logical-unit-number*]** hierarchy level. If none of the DLCIs in

Copyright © 2013, Juniper Networks, Inc.

an MLFR FRF.16 bundle specify a DLCI scheduler, the total bandwidth is evenly divided across all DLCIs.



NOTE: For FRF.16 bundles on link services IQ interfaces, only shaping rates based on percentage are supported.

Configuring Drop Profiles

You can configure random early detection (RED) on LSQ interfaces as in other CoS scenarios. To configure RED, include one or more drop profiles and attach them to a scheduler for a particular forwarding class. For more information about RED profiles, see the Junos OS Class of Service Configuration Guide.

The LSQ implementation performs tail RED. It supports a maximum of 256 drop profiles per PIC. Drop profiles are configurable on a per-queue, per-loss-priority, and per-TCP-bit basis.

You can attach scheduler maps with configured RED drop profiles to any LSQ logical interface: an MLPPP bundle, an FRF.15 bundle, or an FRF.16 DLCI. Different queues (forwarding classes) on the same logical interface can have different associated drop profiles.

The following example shows how to configure a RED profile on an LSQ interface:

```
[edit]
class-of-service {
  drop-profiles {
    drop-low {
      # Configure suitable drop profile for low loss priority
      ...
    }
    drop-high {
      # Configure suitable drop profile for high loss priority
      ...
    }
  }
  scheduler-maps {
    schedmap {
      # Best-effort queue will use be-scheduler
      # Other queues may use different schedulers
      forwarding-class be scheduler be-scheduler;
      ...
    }
  }
  schedulers {
    be-scheduler {
      # Configure two drop profiles for low and high loss priority
      drop-profile-map loss-priority low protocol any drop-profile drop-low;
      drop-profile-map loss-priority high protocol any drop-profile drop-high;
      # Other scheduler parameters (buffer-size, priority,
      # and transmit-rate) are already supported.
      ...
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  interfaces {
    lsq-1/3/0.0 {
      # Attach a scheduler map (that includes RED drop profiles)
      # to a LSQ logical interface.
      scheduler-map schedmap;
    }
  }
}
```



NOTE: The RED profiles should be applied only on the LSQ bundles and not on the egress links that constitute the bundle.

Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces

For link services IQ (**lsq-**) interfaces, you can specify fragmentation properties for specific forwarding classes. Traffic on each forwarding class can be either multilink encapsulated (fragmented and sequenced) or nonencapsulated (hashed with no fragmentation). By default, traffic in all forwarding classes is multilink encapsulated.

When you do not configure fragmentation properties for the queues on MLPPP interfaces, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number* fragment-threshold]** hierarchy level is the fragmentation threshold for all forwarding classes within the MLPPP interface. For MLFR FRF.16 interfaces, the fragmentation threshold you set at the **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options fragment-threshold]** hierarchy level is the fragmentation threshold for all forwarding classes within the MLFR FRF.16 interface.

If you do not set a maximum fragment size anywhere in the configuration, packets are still fragmented if they exceed the smallest maximum transmission unit (MTU) or maximum received reconstructed unit (MRRU) of all the links in the bundle. A nonencapsulated flow uses only one link. If the flow exceeds a single link, then the forwarding class must be multilink encapsulated, unless the packet size exceeds the MTU/MRRU.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the MRRU by including the **mrru** statement at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

To configure fragmentation properties on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
```

Copyright © 2013, Juniper Networks, Inc.

```
        forwarding-class class-name {
          (fragment-threshold bytes | no-fragmentation);
          multilink-class number;
        }
      }
    }
```

To set a per-forwarding class fragmentation threshold, include the **fragment-threshold** statement in the fragmentation map. This statement sets the maximum size of each multilink fragment.

To set traffic on a queue to be nonencapsulated rather than multilink encapsulated, include the **no-fragmentation** statement in the fragmentation map. This statement specifies that an extra fragmentation header is not prepended to the packets received on this queue and that static link load balancing is used to ensure in-order packet delivery.

For a given forwarding class, you can include either the **fragment-threshold** or **no-fragmentation** statement; they are mutually exclusive.

You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For a given forwarding class, you can include either the **multilink-class** or **no-fragmentation** statement; they are mutually exclusive. For more information about MCML, see "Configuring Multiclass MLPPP on LSQ Interfaces" on page 367.

To associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI, include the **fragmentation-map** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces]
lsq-fpc/pic/port {
  unit logical-unit-number { # Multilink PPP
  fragmentation-map map-name;
}
lsq-fpc/pic/port:channel { # MLFR FRF.16
unit logical-unit-number {
  fragmentation-map map-name;
}
```

For configuration examples, see the following topics:

- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 380
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 385
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 390
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 395
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 on page 402
- Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 403

Copyright © 2013, Juniper Networks, Inc.

- Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 404
- Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 406

For Link Services PIC link services (**ls-**) interfaces, fragmentation maps are not supported. Instead, you enable LFI by including the **interleave-fragments** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. For more information, see "Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces" on page 1144.

Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces

Link-layer overhead can cause packet drops on constituent links because of bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information.

By default, 4 percent of the total bundle bandwidth is set aside for link-layer overhead. In most network environments, the average link-layer overhead is 1.6 percent. Therefore, we recommend 4 percent as a safeguard. For more information, see RFC 4814, *Hash and Stuffing: Overlooked Factors in Network Device Benchmarking*.

For link services IQ (**lsq-**) interfaces, you can configure the percentage of bundle bandwidth to be set aside for link-layer overhead. To do this, include the **link-layer-overhead** statement:

```
link-layer-overhead percent;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]**
- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You can configure the value to be from 0 percent through 50 percent.

Configuring Multiclass MLPPP on LSQ Interfaces

For link services IQ (**lsq-**) interfaces with MLPPP encapsulation, you can configure multiclass MLPPP (MCML). If you do not configure MCML, fragments from different classes cannot be interleaved. All fragments for a single packet must be sent before the fragments from another packet are sent. Nonfragmented packets can be interleaved between fragments of another packet to reduce latency seen by nonfragmented packets. In effect, latency-sensitive traffic is encapsulated as regular PPP traffic, and bulk traffic is encapsulated as multilink traffic. This model works as long as there is a single class of latency-sensitive traffic, and there is no high-priority traffic that takes precedence over latency-sensitive traffic. This approach to LFI, used on the Link Services PIC, supports only two levels of traffic priority, which is not sufficient to carry the four-to-eight forwarding classes that are supported by M Series and T Series routers. For more information about the Link Services PIC support of LFI, see "Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces" on page 1144.

Copyright © 2013, Juniper Networks, Inc.

For link services IQ interfaces only, you can configure MCML, as defined in RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. MCML makes it possible to have multiple classes of latency-sensitive traffic that are carried over a single multilink bundle with bulk traffic. In effect, MCML allows different classes of traffic to have different latency guarantees. With MCML, you can map each forwarding class into a separate multilink class, thus preserving priority and latency guarantees.



NOTE: Configuring both LFI and MCML on the same bundle is not necessary, nor is it supported, because multiclass MLPPP represents a superset of functionality. When you configure multiclass MLPPP, LFI is automatically enabled.

The Junos OS implementation of MCML does not support compression of common header bytes, which is referred to in RFC 2686 as "prefix elision."

MCML greatly simplifies packet ordering issues that occur when multiple links are used. Without MCML, all voice traffic belonging to a single flow is hashed to a single link to avoid packet ordering issues. With MCML, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about voice services support on link services IQ interfaces (**lsq**), see "Configuring Services Interfaces for Voice Services" on page 422.

To configure MCML on a link services IQ interface, you must specify how many multilink classes should be negotiated when a link joins the bundle, and you must specify the mapping of a forwarding class into an MCML class.

To specify how many multilink classes should be negotiated when a link joins the bundle, include the **multilink-max-classes** statement:

multilink-max-classes *number*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The number of multilink classes can be 1 through 8. The number of multilink classes for each forwarding class must not exceed the number of multilink classes to be negotiated.

To specify the mapping of a forwarding class into a MCML class, include the **multilink-class** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level:

[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]
multilink-class *number*;

The multilink class index number can be 0 through 7. The **multilink-class** statement and **no-fragmentation** statements are mutually exclusive.

Copyright © 2013, Juniper Networks, Inc.

To view the number of multilink classes negotiated, issue the **show interfaces lsq-*fpc/pic/port.logical-unit-number* detail** command.

Oversubscribing Interface Bandwidth on LSQ Interfaces

The term *oversubscribing interface bandwidth* means configuring shaping rates (peak information rates [PIRs]) so that their sum exceeds the interface bandwidth.

On Channelized IQ PICs, Gigabit Ethernet IQ PICs, and FRF.16 link services IQ (**lsq-**) interfaces on AS and Multiservices PICs, you can oversubscribe interface bandwidth. The logical interfaces (and DLCIs within an FRF.16 bundle) can be oversubscribed when there is leftover bandwidth. The oversubscription is limited to the configured PIR. Any unused bandwidth is distributed equally among oversubscribed logical interfaces or DLCIs.

For networks that are not likely to experience congestion, oversubscribing interface bandwidth improves network utilization, thereby allowing more customers to be provisioned on a single interface. If the actual data traffic does not exceed the interface bandwidth, oversubscription allows you to sell more bandwidth than the interface can support.

We recommend avoiding oversubscription in networks that are likely to experience congestion. Be careful not to oversubscribe a service by too much, because this can cause degradation in the performance of the router during congestion. When you configure oversubscription, some output queues can be starved if the actual data traffic exceeds the physical interface bandwidth. You can prevent degradation by using statistical multiplexing to ensure that the actual data traffic does not exceed the interface bandwidth.



NOTE: You cannot oversubscribe interface bandwidth when you configure traffic shaping using the method described in Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs.

When configuring oversubscription for FRF.16 bundle interfaces, you can assign traffic control profiles that apply on a physical interface basis. When you apply traffic control profiles to FRF.16 bundles at the *logical* interface level, member link interface bandwidth is underutilized when there is a small proportion of traffic or no traffic at all on an individual DLCI. Support for traffic control features on the FRF.16 bundle physical interface level addresses this limitation.

To configure oversubscription of an interface, perform the following steps:

1. Include the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   shaping-rate (percent percentage | rate);
   ```

Copyright © 2013, Juniper Networks, Inc.



NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **shaping-rate** as a percentage.

On LSQ interfaces, you can configure the shaping rate as a percentage.

On IQ and IQ2 interfaces, you can configure the shaping rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. However, with this configuration approach, you cannot independently control the delay-buffer rate, as described in Step 2.



NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

This restriction does not apply to Gigabit Ethernet IQ2 PICs or link services IQ (LSQ) interfaces on AS or Multiservices PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see "Configuring Guaranteed Minimum Rate on LSQ Interfaces" on page 374.

2. Optionally, you can base the delay buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:



NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **delay-buffer-rate** as a percentage.

```
[edit class-of-service traffic-control-profiles profile-name]
delay-buffer-rate (percent percentage | rate);
```

The delay-buffer rate overrides the shaping rate as the basis for the delay-buffer calculation. In other words, the shaping rate or scaled shaping rate is used for delay-buffer calculations only when the delay-buffer rate is not configured.

For LSQ interfaces, if you do not configure a delay-buffer rate, the guaranteed rate (CIR) is used to assign buffers. If you do not configure a guaranteed rate, the shaping rate (PIR) is used in the undersubscribed case, and the scaled shaping rate is used in the oversubscribed case.

Copyright © 2013, Juniper Networks, Inc.

On LSQ interfaces, you can configure the delay-buffer rate as a percentage.

On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

The actual delay buffer is based on the calculations described in the Junos OS Class of Service Configuration Guide. For an example showing how the delay-buffer rates are applied, see "Examples: Oversubscribing an LSQ Interface" on page 372.

Configuring large buffers on relatively low-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed.

This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of zero, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If you do not configure a delay-buffer rate or a guaranteed rate, the logical interface receives a delay-buffer rate in proportion to the shaping rate and the remaining delay-buffer rate available. In other words, the delay-buffer rate for each logical interface with no configured delay-buffer rate is equal to:

(remaining delay-buffer rate * shaping rate) / (sum of shaping rates)

The remaining delay-buffer rate is equal to:

(interface speed) – (sum of configured delay-buffer rates)

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   [edit class-of-service traffic-control-profiles *profile-name*]
   scheduler-map *map-name*;

   For information about configuring schedulers and scheduler maps, see the Junos OS Class of Service Configuration Guide.

4. Optionally, you can enable large buffer sizes to be configured. To do this, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

   [edit chassis fpc *slot-number* pic *pic-number*]
   q-pic-large-buffer;

   If you do not include this statement, the delay-buffer size is more restricted.
   We recommend restricted buffers for delay-sensitive traffic, such as voice traffic. For more information, see the Junos OS Class of Service Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit interfaces interface-name ]
   per-unit-scheduler;
   ```

   When you include this statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a two-port Gigabit Ethernet IQ PIC, the maximum number is 384.

6. To enable scheduling for FRF.16 bundles physical interfaces, include the **no-per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit interfaces interface-name]
   no-per-unit-scheduler;
   ```

7. To apply the traffic-scheduling profile to the logical interface, include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

   ```
   [edit class-of-service interfaces interface-name unit logical-unit-number]
   output-traffic-control-profile profile-name;
   ```

   You cannot include the **output-traffic-control-profile** statement in the configuration if any of the following statements are included in the logical interface configuration: **scheduler-map**, **shaping-rate**, **adaptive-shaper**, or **virtual-channel-group**.

   For a table that shows how the bandwidth and delay buffer are allocated in various configurations, see the Junos OS Class of Service Configuration Guide.

Examples: Oversubscribing an LSQ Interface

Oversubscribing an LSQ Interface with Scheduling Based on the Logical Interface

Apply a traffic-control profile to a logical interface representing a DLCI on an FRF.16 bundle.

```
interfaces {
  lsq-1/3/0:0 {
    per-unit-scheduler;
    unit 0 {
      dlci 100;
    }
    unit 1 {
      dlci 200;
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tc_0 {
      shaping-rate percent 100;
      guaranteed-rate percent 60;
      delay-buffer-rate percent 80;
    }
    tc_1 {
      shaping-rate percent 80;
      guaranteed-rate percent 40;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    interfaces {
      lsq-1/3/0 {
        unit 0 {
          output-traffic-control-profile tc_0;
        }
        unit 1 {
          output-traffic-control-profile tc_1;
        }
      }
    }
}
```

Oversubscribing an LSQ Interface with Scheduling Based on the Physical Interface

Apply a traffic-control profile to the physical interface representing an FRF.16 bundle:

```
interfaces {
  lsq-0/2/0:0 {
    no-per-unit-scheduler;
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      dlci 100;
      family inet {
        address 18.18.18.2/24;
      }
    }
  }
class-of-service {
  traffic-control-profiles {
    rlsq_tc {
      scheduler-map rlsq;
      shaping-rate percent 60;
      delay-buffer-rate percent 10;
    }
  }
  interfaces {
    lsq-0/2/0:0 {
      output-traffic-control-profile rlsq_tc;
    }
  }
}
scheduler-maps {
  rlsq {
    forwarding-class best-effort scheduler rlsq_scheduler;
    forwarding-class expedited-forwarding scheduler rlsq_scheduler1;
  }
}
schedulers {
  rlsq_scheduler {
    transmit-rate percent 20;
    priority low;
  }
  rlsq_scheduler1 {
    transmit-rate percent 40;
    priority high;
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

Configuring Guaranteed Minimum Rate on LSQ Interfaces

On Gigabit Ethernet IQ PICs, Channelized IQ PICs, and FRF.16 link services IQ (LSQ) interfaces on AS and Multiservices PICs, you can configure guaranteed bandwidth, also known as a committed information rate (CIR). This allows you to specify a guaranteed rate for each logical interface. The guaranteed rate is a minimum. If excess physical interface bandwidth is available for use, the logical interface receives more than the guaranteed rate provisioned for the interface.

You cannot provision the sum of the guaranteed rates to be more than the physical interface bandwidth, or the bundle bandwidth for LSQ interfaces. If the sum of the guaranteed rates exceeds the interface or bundle bandwidth, the commit operation does not fail, but the software automatically decreases the rates so that the sum of the guaranteed rates is equal to the available bundle bandwidth.

To configure a guaranteed minimum rate, perform the following steps:

1. Include the **guaranteed-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   guaranteed-rate (percent percentage | rate);
   ```

   On LSQ interfaces, you can configure the guaranteed rate as a percentage.

   On IQ and IQ2 interfaces, you can configure the guaranteed rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

   

   NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

   This restriction does not apply to Gigabit Ethernet IQ2 PICs or link services IQ (LSQ) interfaces on AS or Multiservices PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see the Junos OS Class of Service Configuration Guide.

2. Optionally, you can base the delay buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   delay-buffer-rate (percent percentage | rate);
   ```

   On LSQ interfaces, you can configure the delay-buffer rate as a percentage.

Copyright © 2013, Juniper Networks, Inc.

On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bits per second.

The actual delay buffer is based on the calculations described in tables in the Junos OS Class of Service Configuration Guide. For an example showing how the delay-buffer rates are applied, see "Example: Configuring Guaranteed Minimum Rate" on page 376.

If you do not include the **delay-buffer-rate** statement, the delay-buffer calculation is based on the guaranteed rate, the shaping rate if no guaranteed rate is configured, or the scaled shaping rate if the interface is oversubscribed.

If you do not specify a shaping rate or a guaranteed rate, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 4 MTU-sized packets.

You can configure a rate for the delay buffer that is higher than the guaranteed rate. This can be useful when the traffic flow might not require much bandwidth in general, but in some cases can be bursty and therefore needs a large buffer.

Configuring large buffers on relatively low-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed. This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of 0, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If the guaranteed rate of a logical interface cannot be implemented, that logical interface receives a delay-buffer rate of 0, even if the configured delay-buffer rate is within the interface speed. If at a later time the guaranteed rate of the logical interface can be met, the configured delay-buffer rate is reevaluated and if the delay-buffer rate is within the remaining bandwidth, it is implemented.

If any logical interface has a configured guaranteed rate, all other logical interfaces on that port that do not have a guaranteed rate configured receive a delay-buffer rate of 0. This is because the absence of a guaranteed rate configuration corresponds to a guaranteed rate of 0 and, consequently, a delay-buffer rate of 0.

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   scheduler-map map-name;
   ```

   For information about configuring schedulers and scheduler maps, see the Junos OS Class of Service Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

4. To enable large buffer sizes to be configured, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

   ```
   [edit chassis fpc slot-number pic pic-number]
   q-pic-large-buffer;
   ```

   If you do not include this statement, the delay-buffer size is more restricted. For more information, see the Junos OS Class of Service Configuration Guide.

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit interfaces interface-name ]
   per-unit-scheduler;
   ```

   When you include this statement, the maximum number of VLANs supported is 767 on a single-port Gigabit Ethernet IQ PIC. On a two-port Gigabit Ethernet IQ PIC, the maximum number is 383.

6. To apply the traffic-scheduling profile to the logical interface, include the output-traffic-control-profile statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

   ```
   [edit class-of-service interfaces interface-name unit logical-unit-number]
   output-traffic-control-profile profile-name;
   ```

Example: Configuring Guaranteed Minimum Rate

Two logical interface units, **0** and **1**, are provisioned with a guaranteed minimum of 750 Kbps and 500 Kbps, respectively. For logical unit **1**, the delay buffer is based on the guaranteed rate setting. For logical unit **0**, a delay-buffer rate of 500 Kbps is specified. The actual delay buffers allocated to each logical interface are 2 seconds of 500 Kbps. The 2-second value is based on the following calculation:

```
delay-buffer-rate < [8 x 64 Kbps]): 2 seconds of delay-buffer-rate
```

For more information about this calculation, see the Junos OS Class of Service Configuration Guide.

```
chassis {
    fpc 3 {
        pic 0 {
            q-pic-large-buffer;
        }
    }
}
interfaces {
    t1-3/0/1 {
        per-unit-scheduler;
    }
}
class-of-service {
    traffic-control-profiles {
        tc-profile3 {
            guaranteed-rate 750k;
            scheduler-map sched-map3;
            delay-buffer-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    tc-profile4 {
      guaranteed-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
      scheduler-map sched-map4;
    }
  }
  interface t1-3/0/1 {
    unit 0 {
      output-traffic-control-profile tc-profile3;
    }
    unit 1 {
      output-traffic-control-profile tc-profile4;
    }
  }
}
```

Configuring Link Services and CoS on Services PICs

To configure link services and CoS on an AS or Multiservices PIC, you must perform the following steps:

1. Enable the Layer 2 service package. You enable service packages per PIC, not per port. When you enable the Layer 2 service package, the entire PIC uses the configured package. To enable the Layer 2 service package, include the **service-package** statement at the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services]** hierarchy level, and specify **layer-2**:

   ```
   [edit chassis fpc slot-number pic pic-number adaptive-services]
   service-package layer-2;
   ```

   For more information about AS or Multiservices PIC service packages, see "Enabling Service Packages" on page 39 and "Layer 2 Service Package Capabilities and Interfaces" on page 348.

2. Configure a multilink PPP or FRF.16 bundle by combining constituent links into a virtual link, or bundle.

   ### Configuring an MLPPP Bundle

   To configure an MLPPP bundle, configure constituent links and bundle properties by including the following statements in the configuration:

   ```
   [edit interfaces interface-name unit logical-unit-number]
   encapsulation ppp;
   family mlppp {
     bundle lsq-fpc/pic/port.logical-unit-number;
   }
   [edit interfaces lsq-fpc/pic/port unit logical-unit-number]
   drop-timeout milliseconds;
   encapsulation multilink-ppp;
   fragment-threshold bytes;
   link-layer-overhead percent;
   minimum-links number;
   mrru bytes;
   short-sequence;
   family inet {
   ```

Copyright © 2013, Juniper Networks, Inc.

```
    address address;
  }
```

For more information about these statements, see the Link and Multilink Properties.

Configuring an MLFR FRF.16 Bundle

To configure an MLFR FRF.16 bundle, configure constituent links and bundle properties by including the following statements in the configuration:

```
[edit chassis fpc slot-number pic slot-number]
mlfr-uni-nni-bundles number;
[edit interfaces interface-name ]
encapsulation multilink-frame-relay-uni-nni;
unit logical-unit-number {
  family mlfr-uni-nni {
    bundle lsq-fpc/pic/port:channel;
  }
}
```

For more information about the **mlfr-uni-nni-bundles** statement, see the Junos OS System Basics Configuration Guide. MLFR FRF.16 uses channels as logical units.

For MLFR FRF.16, you must configure one end as data circuit-terminating equipment (DCE) by including the following statements at the **[edit interfaces lsq-*fpc/pic/port*:*channel*]** hierarchy level.

```
encapsulation multilink-frame-relay-uni-nni;
dce;
mlfr-uni-nni-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  hello-timer milliseconds;
  link-layer-overhead percent;
  lmi-type (ansi | itu);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 number;
  t392 number;
  yellow-differential-delay milliseconds;
}
unit logical-unit-number {
  dlci dlci-identifier;
  family inet {
    address address;
  }
}
```

For more information about MLFR UNI NNI properties, see Link and Multilink Properties.

Copyright © 2013, Juniper Networks, Inc.

3. To configure CoS components for each multilink bundle, enable per-unit scheduling on the interface, configure a scheduler map, apply the scheduler to each queue, configure a fragmentation map, and apply the fragmentation map to each bundle. Include the following statements:

```
[edit interfaces]
lsq-fpc/pic/port {
  per-unit-scheduler; # Enables per-unit scheduling on the bundle
}
[edit class-of-service]
interfaces {
  lsq-fpc/pic/port { # Multilink PPP
  unit logical-unit-number {
    scheduler-map map-name; # Applies scheduler map to each queue
  }
}
lsq-fpc/pic/port:channel { # MLFR FRF.16
unit logical-unit-number {
  # Scheduler map provides scheduling information for
  # the queues within a single DLCI.
  scheduler-map map-name;
  shaping-rate percent percent;
}
forwarding-classes {
  queue queue-number class-name priority (high | low);
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (percent percentage | rate | remainder) <exact>;
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      no-fragmentation;
    }
  }
}
```

Associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI by including the following statements at the **[edit class-of-service]** hierarchy level:

```
interfaces {
  lsq-fpc/pic/port {
    unit logical-unit-number { # Multilink PPP
    fragmentation-map map-name;
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

```
lsq-fpc/pic/port:channel { # MLFR FRF.16
unit logical-unit-number {
  fragmentation-map map-name;
}
```

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP

To configure an *N*xT1 bundle using MLPPP, you aggregate *N* different T1 links into a bundle. The *N*xT1 bundle is called a logical interface, because it can represent, for example, a routing adjacency. To aggregate T1 links into a an MLPPP bundle, include the **bundle** statement at the **[edit interfaces t1-*fpc/pic/port* unit *logical-unit-number* family mlppp]** hierarchy level:

```
[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
  address address;
}
```

The logical link services IQ interface represents the MLPPP bundle. For the MLPPP bundle, there are four associated queues on M Series routers and eight associated queues on M320 and T Series routers. A scheduler removes packets from the queues according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

For MLPPP, assign a single scheduler map to the link services IQ interface (**lsq**) and to each constituent link. The default schedulers for M Series and T Series routers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for MLPPP, you should configure a single scheduler with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ interface (**lsq**) and to each constituent link, as shown in "Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP" on page 383.

Copyright © 2013, Juniper Networks, Inc.



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

If the bundle has more than one link, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-*fpc/pic/port*]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port]
per-unit-scheduler;
```

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
  t1-fpc/pic/port unit logical-unit-number {
    scheduler-map map-name;
  }
}
forwarding-classes {
  queue queue-number class-name;
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (rate | percent percentage | remainder) <exact>;
  }
}
```

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the Junos OS Class of Service Configuration Guide.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
```

Copyright © 2013, Juniper Networks, Inc.

```
            multilink-class number;
            no-fragmentation;
        }
    }
}
```

For *N*xT1 bundles using MLPPP, the byte-wise load balancing used in multilink-encapsulated queues is superior to the flow-wise load balancing used in nonencapsulated queues. All other considerations are equal. Therefore, we recommend that you configure all queues to be multilink encapsulated. You do this by including the **fragment-threshold** statement in the configuration. If you choose to set traffic on a queue to be nonencapsulated rather than multilink encapsulated, include the **no-fragmentation** statement in the fragmentation map. You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For more information about MCML, see "Configuring Multiclass MLPPP on LSQ Interfaces" on page 367. For more information about fragmentation maps, see "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an MLPPP header. The MLPPP header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on one of the *N* different T1 links. The link is chosen on a packet-by-packet basis to balance the load across the various T1 links.

If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers. The outgoing link for each fragment is selected independently of all other fragments.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. Because there is no MLPPP header, there is no sequence number information. Therefore, the software must take special measures to avoid packet reordering. To avoid packet reordering, the software places the packet on one of the *N* different T1 links. The link is determined by hashing the values in the header. For IP, the software computes the hash based on source address, destination address, and IP protocol. For MPLS, the

 Copyright © 2013, Juniper Networks, Inc.

software computes the hash based on up to five MPLS labels, or four MPLS labels and the IP header.

For UDP and TCP the software computes the hash based on the source and destination ports, as well as source and destination IP addresses. This guarantees that all packets belonging to the same TCP/UDP flow always pass through the same T1 link, and therefore cannot be reordered. However, it does not guarantee that the load on the various T1 links is balanced. If there are many flows, the load is usually balanced.

The *N* different T1 interfaces link to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from all the T1 links. If a packet has an MLPPP header, the sequence number field is used to put the packet back into sequence number order. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives and makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP

```
[edit chassis]
fpc 1 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
t1-0/0/0 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1; # This adds t1-0/0/0 to the specified bundle.
    }
  }
}
t1-0/0/1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0 {
  unit 1 { # This is the virtual link that concatenates multiple T1s.
  encapsulation multilink-ppp;
  drop-timeout 1000;
  fragment-threshold 128;
  link-layer-overhead 0.5;
  minimum-links 2;
  mrru 4500;
  short-sequence;
  family inet {
    address 10.2.3.4/24;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
}
[edit interfaces]
lsq-1/3/0 {
    per-unit-scheduler;
}
[edit class-of-service]
interfaces {
    lsq-1/3/0 { # multilink PPP constituent link
    unit 0 {
        scheduler-map sched-map1;
    }
}
t1-0/0/0 { # multilink PPP constituent link
unit 0 {
    scheduler-map sched-map1;
}
t1-0/0/1 { # multilink PPP constituent link
unit 0 {
    scheduler-map sched-map1;
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
scheduler-maps {
    sched-map1 {
        forwarding-class af scheduler af-scheduler;
        forwarding-class be scheduler be-scheduler;
        forwarding-class ef scheduler ef-scheduler;
        forwarding-class nc scheduler nc-scheduler;
    }
}
schedulers {
    af-scheduler {
        transmit-rate percent 30;
        buffer-size percent 30;
        priority low;
    }
    be-scheduler {
        transmit-rate percent 25;
        buffer-size percent 25;
        priority low;
    }
    ef-scheduler {
        transmit-rate percent 40;
        buffer-size percent 40;
        priority strict-high; # voice queue
    }
    nc-scheduler {
        transmit-rate percent 5;
        buffer-size percent 5;
        priority high;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
}
fragmentation-maps {
  fragmap-1 {
    forwarding-class be {
      fragment-threshold 180;
    }
    forwarding-class ef {
      fragment-threshold 100;
    }
  }
}
[edit interfaces]
lsq-1/3/0 {
  unit 0 {
    fragmentation-map fragmap-1;
  }
}
```

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16

To configure an *N*xT1 bundle using FRF.16, you aggregate *N* different T1 links into a bundle. The *N*xT1 bundle carries a potentially large number of Frame Relay PVCs, identified by their DLCIs. Each DLCI is called a logical interface, because it can represent, for example, a routing adjacency.

To aggregate T1 links into an FRF.16 bundle, include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc *slot-number* pic *slot-number*]** hierarchy level and include the **bundle** statement at the **[edit interfaces t1-*fpc/pic/port* unit *logical-unit-number* family mlfr-uni-nni]** hierarchy level:

```
[edit chassis fpc slot-number pic slot-number]
mlfr-uni-nni-bundles number;

[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlfr-uni-nni]
bundle lsq-fpc/pic/port:channel;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq- *fpc/pic/port:channel*]** hierarchy level:

```
[edit interfaces lsq- fpc/pic/port:channel]
encapsulation multilink-frame-relay-uni-nni;
dce;
mlfr-uni-nni-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
```

Copyright © 2013, Juniper Networks, Inc.

```
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 number;
    t392 number;
    yellow-differential-delay milliseconds;
  }
  unit logical-unit-number {
    dlci dlci-identifier;
    family inet {
      address address;
    }
  }
```

The link services IQ channel represents the FRF.16 bundle. Four queues are associated with each DLCI. A scheduler removes packets from the queues according to a scheduling policy. On the link services IQ interface, you typically designate one queue to have strict priority. The remaining queues are serviced in proportion to weights you configure.

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high-priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the Junos OS Class of Service Configuration Guide.

If the bundle has more than one link, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-*fpc/pic/port*:*channel*]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port:channel]
per-unit-scheduler;
```

For FRF.16, you can assign a single scheduler map to the link services IQ interface (**lsq**) and to each link services IQ DLCI, or you can assign different scheduler maps to the various DLCIs of the bundle, as shown in "Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16" on page 388.

For the constituent links of an FRF.16 bundle, you do not need to configure a custom scheduler. Because LFI and multiclass are not supported for FRF.16, the traffic from each constituent link is transmitted from queue 0. This means you should allow most of the bandwidth to be used by queue 0. For M Series and T Series routers, the default schedulers' transmission rate and buffer size percentages for queues 0 through 3 are 95, 0, 0, and 5 percent. These default schedulers send all user traffic to queue 0 and all network-control traffic to queue 3, and therefore are well suited to the behavior of FRF.16. If desired, you can configure a custom scheduler that explicitly replicates the 95, 0, 0, and 5 percent queuing behavior, and apply it to the constituent links.

Copyright © 2013, Juniper Networks, Inc.



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
  lsq-fpc/pic/port:channel {
    unit logical-unit-number {
      scheduler-map map-name;
    }
  }
}
forwarding-classes {
  queue queue-number class-name;
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (rate | percent percentage | remainder) <exact>;
  }
}
```

To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
    }
  }
}
```

For FRF.16 traffic, only multilink encapsulated (fragmented and sequenced) queues are supported. This is the default queuing behavior for all forwarding classes. FRF.16 does not allow for nonencapsulated traffic because the protocol requires that all packets carry the fragmentation header. If a large packet is split into multiple fragments, the fragments must have consecutive sequential numbers. Therefore, you cannot include the **no-fragmentation** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level for FRF.16 traffic. For FRF.16, if you want to carry voice or any other latency-sensitive traffic, you should not use slow links. At T1 speeds and above, the serialization delay is small enough so that you do not need to use explicit LFI.

Copyright © 2013, Juniper Networks, Inc.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an FRF.16 header. The FRF.16 header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on one of the *N* different T1 links. The link is chosen on a packet-by-packet basis to balance the load across the various T1 links.

If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers. The outgoing link for each fragment is selected independently of all other fragments.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level. The MRRU is similar to the MTU but is specific to link services interfaces. By default, the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

The *N* different T1 interfaces link to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from all the T1 links. Because each packet has an FRF.16 header, the sequence number field is used to put the packet back into sequence number order.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16

Configure an *N*xT1 bundle using FRF.16 with multiple CoS scheduler maps:

```
[edit chassis fpc 1 pic 3]
adaptive-services {
   service-package layer-2;
}
mlfr-uni-nni-bundles 2; # Creates channelized LSQ interfaces/FRF.16 bundles.
[edit interfaces]
t1-0/0/0 {
   encapsulation multilink-frame-relay-uni-nni;
   unit 0 {
      family mlfr-uni-nni {
         bundle lsq-1/3/0:1;
      }
   }
}
t1-0/0/1 {
   encapsulation multilink-frame-relay-uni-nni;
```


Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family mlfr-uni-nni {
        bundle lsq-1/3/0:1;
      }
    }
}
lsq-1/3/0:1 { # Bundle link consisting of t1-0/0/0 and t1-0/0/1
per-unit-scheduler;
encapsulation multilink-frame-relay-uni-nni;
dce; # One end needs to be configured as DCE.
mlfr-uni-nni-bundle-options {
   drop-timeout 180;
   fragment-threshold 64;
   hello-timer 180;
   minimum-links 2;
   mrru 3000;
   link-layer-overhead 0.5;
}
unit 0 {
   dlci 26; # Each logical unit maps a single DLCI.
   family inet {
      address 10.2.3.4/24;
   }
}
unit 1 {
   dlci 42;
   family inet {
      address 10.20.30.40/24;
   }
}
unit 2 {
   dlci 69;
   family inet {
      address 10.20.30.40/24;
   }
}
[edit class-of-service]
scheduler-maps {
   sched-map-lsq0 {
      forwarding-class af scheduler af-scheduler-lsq0;
      forwarding-class be scheduler be-scheduler-lsq0;
      forwarding-class ef scheduler ef-scheduler-lsq0;
      forwarding-class nc scheduler nc-scheduler-lsq0;
   }
   sched-map-lsq1 {
      forwarding-class af scheduler af-scheduler-lsq1;
      forwarding-class be scheduler be-scheduler-lsq1;
      forwarding-class ef scheduler ef-scheduler-lsq1;
      forwarding-class nc scheduler nc-scheduler-lsq1;
   }
}
schedulers {
   af-scheduler-lsq0 {
      transmit-rate percent 60;
      buffer-size percent 60;
      priority low;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    be-scheduler-lsq0 {
      transmit-rate percent 30;
      buffer-size percent 30;
      priority low;
    }
    ef-scheduler-lsq0 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority strict-high;
    }
    nc-scheduler-lsq0 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority high;
    }
    af-scheduler-lsq1 {
      transmit-rate percent 50;
      buffer-size percent 50;
      priority low;
    }
    be-scheduler-lsq1 {
      transmit-rate percent 30;
      buffer-size percent 30;
      priority low;
    }
    ef-scheduler-lsq1 {
      transmit-rate percent 15;
      buffer-size percent 15;
      priority strict-high;
    }
    nc-scheduler-lsq1 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority high;
    }
  }
  interfaces {
    lsq-1/3/0:1 { # MLFR FRF.16
    unit 0 {
      scheduler-map sched-map-lsq0;
    }
    unit 1 {
      scheduler-map sched-map-lsq1;
    }
  }
```

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI

When you configure a single fractional T1 interface, it is called a logical interface, because it can represent, for example, a routing adjacency.

The logical link services IQ interface represents the MLPPP bundle. Four queues are associated with the logical interface. A scheduler removes packets from the queues

Copyright © 2013, Juniper Networks, Inc.

according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

To configure a single fractional T1 interface using MLPPP and LFI, you associate one DS0 (fractional T1) interface with a link services IQ interface. To associate a fractional T1 interface with a link services IQ interface, include the **bundle** statement at the **[edit interfaces ds-*fpc/pic/port*:*channel* unit *logical-unit-number* family mlppp]** hierarchy level:

```
[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
   address address;
}
```

For MLPPP, assign a single scheduler map to the link services IQ (**lsq**) interface and to each constituent link. The default schedulers for M Series and T Series routers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for MLPPP, you should configure a single scheduler with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ (**lsq**) interface and to each constituent link and to each constituent link, as shown in "Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI" on page 393.



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
   ds-fpc/pic/port.channel {
```

Copyright © 2013, Juniper Networks, Inc.

```
      scheduler-map map-name;
    }
  }
  forwarding-classes {
    queue queue-number class-name;
  }
  scheduler-maps {
    map-name {
      forwarding-class class-name scheduler scheduler-name;
    }
  }
  schedulers {
    scheduler-name {
      buffer-size (percent percentage | remainder | temporal microseconds);
      priority priority-level;
      transmit-rate (rate | percent percentage | remainder) <exact>;
    }
  }
```

For link services IQ interfaces, a strict-high-priority queue might starve all the other queues because traffic in a strict-high priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue receives infinite credits and does round-robin with high-priority queues, as described in the Junos OS Class of Service Configuration Guide.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      no-fragmentation;
    }
  }
}
```

If you require the queue to transmit small packets with low latency, configure the queue to be nonencapsulated by including the **no-fragmentation** statement. If you require the queue to transmit large packets with normal latency, configure the queue to be multilink encapsulated by including the **fragment-threshold** statement. If you require the queue to transmit large packets with low latency, we recommend using a faster link and configuring the queue to be nonencapsulated. For more information about fragmentation maps, see "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

When a packet is removed from a multilink-encapsulated queue, it is fragmented. If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured

 Copyright © 2013, Juniper Networks, Inc.

at the **[edit class-of-service fragmentation-maps** *map-name* **forwarding-class** *class-name***]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces** *interface-name* **unit** *logical-unit-number***]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces lsq-***fpc/pic/port* **unit** *logical-unit-number***]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an MLPPP header. The MLPPP header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on the fractional T1 link. Traffic from another queue might be interleaved between two fragments of the packet.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. The packet is then placed on the fractional T1 link as soon as possible. If necessary, the packet is placed between the fragments of a packet from another queue.

The fractional T1 interface links to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from the fractional T1 link. If a packet has an MLPPP header, the software assumes the packet is a fragment of a larger packet, and the fragment number field is used to reassemble the larger packet. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives, and the software makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI

Configure a single fractional T1 logical interface:

```
[edit interfaces]
lsq-0/2/0 {
   per-unit-scheduler;
   unit 0 {
      encapsulation multilink-ppp;
      link-layer-overhead 0.5;
      family inet {
         address 10.40.1.1/30;
      }
   }
}
ct3-1/0/0 {
   partition 1 interface-type ct1;
}
```

Copyright © 2013, Juniper Networks, Inc.

```
ct1-1/0/0:1 {
   partition 1 timeslots 1-2 interface-type ds;
}
ds-1/0/0:1:1 {
   encapsulation ppp;
   unit 0 {
      family mlppp {
         bundle lsq-0/2/0.0;
      }
   }
}
[edit class-of-service]
interfaces {
   ds-1/0/0:1:1 { # multilink PPP constituent link
   unit 0 {
      scheduler-map sched-map1;
   }
}
forwarding-classes {
   queue 0 be;
   queue 1 ef;
   queue 2 af;
   queue 3 nc;
}
scheduler-maps {
   sched-map1 {
      forwarding-class af scheduler af-scheduler;
      forwarding-class be scheduler be-scheduler;
      forwarding-class ef scheduler ef-scheduler;
      forwarding-class nc scheduler nc-scheduler;
   }
}
schedulers {
   af-scheduler {
      transmit-rate percent 20;
      buffer-size percent 20;
      priority low;
   }
   be-scheduler {
      transmit-rate percent 20;
      buffer-size percent 20;
      priority low;
   }
   ef-scheduler {
      transmit-rate percent 50;
      buffer-size percent 50;
      priority strict-high; # voice queue
   }
   nc-scheduler {
      transmit-rate percent 10;
      buffer-size percent 10;
      priority high;
   }
}
fragmentation-maps {
   fragmap-1 {
```


Copyright © 2013, Juniper Networks, Inc.

```
      forwarding-class be {
        fragment-threshold 180;
      }
      forwarding-class ef {
        fragment-threshold 100;
      }
    }
  }
  [edit interfaces]
  lsq-0/2/0 {
    unit 0 {
      fragmentation-map fragmap-1;
    }
  }
```

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12

To configure a single fractional T1 interface using FRF.16, you associate a DS0 interface with a link services IQ (**lsq**) interface. When you configure a single fractional T1, the fractional T1 carries a potentially large number of Frame Relay PVCs identified by their DLCIs. Each DLCI is called a logical interface, because it can represent, for example, a routing adjacency. To associate the DS0 interface with a link services IQ interface, include the **bundle** statement at the **[edit interfaces ds-*fpc/pic/port*:*channel* unit *logical-unit-number* family mlfr-end-to-end]** hierarchy level:

```
[edit interfaces ds-fpc/pic/port:channel unit logical-unit-number family mlfr-end-to-end]
bundle lsq-fpc/pic/port.logical-unit-number;
```



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the link services IQ interface properties, include the following statements at the **[edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-frame-relay-end-to-end;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
  address address;
}
```

The logical link services IQ interface represents the FRF.12 bundle. Four queues are associated with each logical interface. A scheduler removes packets from the queues according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

Copyright © 2013, Juniper Networks, Inc.

For FRF.12, assign a single scheduler map to the link services IQ interface (**lsq**) and to each constituent link. For M Series and T Series routers, the default schedulers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for FRF.12, you should configure schedulers with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign them to the link services IQ interface (**lsq**) and to each constituent link, as shown in "Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12" on page 398.



NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
  ds-fpc/pic/port.channel {
    scheduler-map map-name;
  }
}
forwarding-classes {
  queue queue-number class-name;
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (rate | percent percentage | remainder) <exact>;
  }
}
```

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high-priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the Junos OS Class of Service Configuration Guide.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
[edit class-of-service]
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      fragment-threshold bytes;
      no-fragmentation;
    }
  }
}
```

If you require the queue to transmit small packets with low latency, configure the queue to be nonencapsulated by including the **no-fragmentation** statement. If you require the queue to transmit large packets with normal latency, configure the queue to be multilink encapsulated by including the **fragment-threshold** statement. If you require the queue to transmit large packets with low latency, we recommend using a faster link and configuring the queue to be nonencapsulated. For more information about fragmentation maps, see "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

When a packet is removed from a multilink-encapsulated queue, it is fragmented. If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps** *map-name* **forwarding-class** *class-name***]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces** *interface-name* **unit** *logical-unit-number***]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces lsq-***fpc/pic/port* **unit** *logical-unit-number***]** hierarchy level. The MRRU is similar to the MTU but is specific to link services interfaces. By default, the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an FRF.12 header. The FRF.12 header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on the fractional T1 link. Traffic from another queue might be interleaved between two fragments of the packet.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain Frame Relay header. The packet is then placed on the fractional T1 link as soon as possible. If necessary, the packet is placed between the fragments of a packet from another queue.

The fractional T1 interface links to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from the fractional T1 link. If a packet has an FRF.12 header, the software assumes the packet is a fragment of a

Copyright © 2013, Juniper Networks, Inc.

larger packet, and the fragment number field is used to reassemble the larger packet. If the packet has a plain Frame Relay header, the software accepts the packet in the order in which it arrives, and the software makes no attempt to reassemble or reorder the packet.

A whole packet from a nonencapsulated queue can be placed between fragments of a multilink-encapsulated queue. However, fragments from one multilink-encapsulated queue cannot be interleaved with fragments from another multilink-encapsulated queue. This is the intent of the specification FRF.12, *Frame Relay Fragmentation Implementation Agreement*. If fragments from two different queues were interleaved, the header fields might not have enough information to separate the fragments.

Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12

FRF.12 with Fragmentation and Without LFI

This example shows a 128 KB DS0 interface. There is one traffic stream on **ge-0/0/0**, which is classified into queue 0 (**be**). Packets are fragmented in the link services IQ (**lsq-**) interface according to the threshold configured in the fragmentation map.

```
[edit chassis]
fpc 0 {
  pic 3 {
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 20.1.1.1/24 {
        arp 20.1.1.2 mac 00.90.1b.12.34.56;
      }
    }
  }
}
ce1-0/2/0 {
  partition 1 timeslots 1-2 interface-type ds;
}
ds-0/2/0:1 {
  no-keepalives;
  dce;
  encapsulation frame-relay;
  unit 0 {
    dlci 100;
    family mlfr-end-to-end {
      bundle lsq-0/3/0.0;
    }
  }
}
lsq-0/3/0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
    per-unit-scheduler;
    unit 0 {
      encapsulation multilink-frame-relay-end-to-end;
      family inet {
        address 10.200.0.78/30;
      }
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 172.16.1.162/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
    }
  }
  [edit class-of-service]
  forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
  }
  interfaces {
    lsq-0/3/0 {
      unit 0 {
        fragmentation-map map1;
      }
    }
  }
  fragmentation-maps {
    map1 {
      forwarding-class {
        be {
          fragment-threshold 160;
        }
      }
    }
  }
```

FRF.12 with Fragmentation and LFI

This example shows a 512 KB DS0 bundle and four traffic streams on **ge-0/0/0** that are classified into four queues. The fragment size is 160 for queue 0, queue 1, and queue 2. The voice stream on queue 3 has LFI configured.

```
  [edit chassis]
  fpc 0 {
    pic 3 {
```

Copyright © 2013, Juniper Networks, Inc.

```
    adaptive-services {
      service-package layer-2;
    }
  }
}
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 20.1.1.1/24 {
        arp 20.1.1.2 mac 00.90.1b.12.34.56;
      }
    }
  }
  ce1-0/2/0 {
    partition 1 timeslots 1-8 interface-type ds;
  }
  ds-0/2/0:1 {
    no-keepalives;
    dce;
    encapsulation frame-relay;
    unit 0 {
      dlci 100;
      family mlfr-end-to-end {
        bundle lsq-0/3/0.0;
      }
    }
  }
  lsq-0/3/0 {
    per-unit-scheduler;
    unit 0 {
      encapsulation multilink-frame-relay-end-to-end;
      family inet {
        address 10.200.0.78/30;
      }
    }
  }
  [edit class-of-service]
  classifiers {
    inet-precedence ge-interface-classifier {
      forwarding-class be {
        loss-priority low code-points 000;
      }
      forwarding-class ef {
        loss-priority low code-points 010;
      }
      forwarding-class af {
        loss-priority low code-points 100;
      }
      forwarding-class nc {
        loss-priority low code-points 110;
      }
    }
  }
  forwarding-classes {
    queue 0 be;
```


Copyright © 2013, Juniper Networks, Inc.

```
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
interfaces {
    lsq-0/3/0 {
        unit 0 {
            scheduler-map sched2;
            fragmentation-map map2;
        }
    }
    ds-0/2/0:1 {
        scheduler-map link-map2;
    }
    ge-0/0/0 {
        unit 0 {
            classifiers {
                inet-precedence ge-interface-classifier;
            }
        }
    }
}
scheduler-maps {
    sched2 {
        forwarding-class be scheduler economy;
        forwarding-class ef scheduler business;
        forwarding-class af scheduler stream;
        forwarding-class nc scheduler voice;
    }
    link-map2 {
        forwarding-class be scheduler link-economy;
        forwarding-class ef scheduler link-business;
        forwarding-class af scheduler link-stream;
        forwarding-class nc scheduler link-voice;
    }
}
fragmentation-maps {
    map2 {
        forwarding-class {
            be {
                fragment-threshold 160;
            }
            ef {
                fragment-threshold 160;
            }
            af {
                fragment-threshold 160;
            }
            nc {
                no-fragmentation;
            }
        }
    }
    schedulers {
        economy {
            transmit-rate percent 26;
```

Copyright © 2013, Juniper Networks, Inc.

```
            buffer-size percent 26;
        }
        business {
            transmit-rate percent 26;
            buffer-size percent 26;
        }
        stream {
            transmit-rate percent 35;
            buffer-size percent 35;
        }
        voice {
            transmit-rate percent 13;
            buffer-size percent 13;
        }
        link-economy {
            transmit-rate percent 26;
            buffer-size percent 26;
        }
        link-business {
            transmit-rate percent 26;
            buffer-size percent 26;
        }
        link-stream {
            transmit-rate percent 35;
            buffer-size percent 35;
        }
        link-voice {
            transmit-rate percent 13;
            buffer-size percent 13;
        }
    }
  }
}
```

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15

This example configures an ***N***xT1 bundle using FRF.15 on a link services IQ interface. FRF.15 is similar to FRF.12, as described in "Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12" on page 395. The difference is that FRF.15 supports multiple physical links in a bundle, whereas FRF.12 supports only one physical link per bundle. For the Junos OS implementation of FRF.15, you can configure one DLCI per physical link.



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. This example refers to T1 interfaces, but the configuration for E1 interfaces is similar.

```
[edit interfaces]
lsq-1/3/0 {
    per-unit-scheduler;
    unit 0 {
        encapsulation multilink-frame-relay-end-to-end;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
}
unit 1 {
  encapsulation multilink-frame-relay-end-to-end;
}
# First physical link
t1-1/1/0:1 {
  encapsulation frame-relay;
  unit 0 {
    dlci 69;
    family mlfr-end-to-end {
      bundle lsq-1/3/0.0;
    }
  }
}
# Second physical link
t1-1/1/0:2 {
  encapsulation frame-relay;
  unit 0 {
    dlci 13;
    family mlfr-end-to-end {
      bundle lsq-1/3/0.0;
    }
  }
}
```

Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP

This example bundles a single T3 interface on a link services IQ interface with MLPPP encapsulation. Binding a single T3 interface to a multilink bundle allows you to configure compressed RTP (CRTP) on the T3 interface.

This scenario applies to MLPPP bundles only. The Junos OS does not currently support CRTP over Frame Relay. For more information, see "Configuring Services Interfaces for Voice Services" on page 422.

There is no need to configure LFI at DS3 speeds, because the packet serialization delay is negligible.

```
[edit interfaces]
t3-0/0/0 {
  unit 0 {
    family mlppp {
      bundle lsq-1/3/0.1;
    }
  }
}
lsq-1/3/0.1 {
  encapsulation multilink-ppp;
}
compression {
  rtp {
    # cRTP parameters go here
    #
    port minimum 2000 maximum 64009;
  }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
```

This configuration uses a default fragmentation map, which results in all forwarding classes (queues) being sent out with a multilink header.

To eliminate multilink headers, you can configure a fragmentation map in which all queues have the **no-fragmentation** statement at the **[edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]** hierarchy level, and attach the fragmentation map to the **lsq-1/3/0.1** interface, as shown here:

```
[edit class-of-service]
fragmentation-maps {
  fragmap {
    forwarding-class {
      be {
        no-fragmentation;
      }
      af {
        no-fragmentation;
      }
      ef {
        no-fragmentation;
      }
      nc {
        no-fragmentation;
      }
    }
  }
}
interfaces {
  lsq-1/3/0.1 {
    fragmentation-map fragmap;
  }
}
```

Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12

This example configures a clear-channel T3 or OC3 interface with multiple logical interfaces (DLCIs) on the link. In this scenario, each DLCI represents a customer. DLCIs are shaped at the egress PIC to a particular speed (*N*xDS0). This allows you to configure LFI using FRF.12 End-to-End Protocol on Frame Relay DLCIs.

To do this, first configure logical interfaces (DLCIs) on the physical interface. Then bundle the DLCIs, so that there is only one DLCI per bundle.

The physical interface must be capable of per-DLCI scheduling, which allows you to attach shaping rates to each DLCI. For more information, see the Junos® OS Network Interfaces.

To prevent fragment drops at the egress PIC, you must assign a shaping rate to the link services IQ logical interfaces and to the egress DLCIs. Shaping rates on DLCIs specify how much bandwidth is available for each DLCI. The shaping rate on link services IQ interfaces should match the shaping rate assigned to the DLCI that is associated with the bundle.

Copyright © 2013, Juniper Networks, Inc.

Egress interfaces also must have a scheduler map attached. The queue that carries voice should be strict-high-priority, while all other queues should be low-priority. This makes LFI possible.

This example shows voice traffic in the **ef** queue. The voice traffic is interleaved with bulk data. Alternatively, you can use multiclass MLPPP to carry multiple classes of traffic in different multilink classes, as described in "Configuring Multiclass MLPPP on LSQ Interfaces" on page 367.

```
[edit interfaces]
t3-0/0/0 {
   per-unit-scheduler;
   encapsulation frame-relay;
   unit 0 {
      dlci 69;
      family mlfr-end-to-end {
         bundle lsq-1/3/0.0;
      }
   }
   unit 1 {
      dlci 42;
      family mlfr-end-to-end {
         bundle lsq-1/3/0.1;
      }
   }
}
lsq-1/3/0 {
   unit 0 {
      encapsulation multilink-frame-relay-end-to-end;
   }
   fragment-threshold 320; # Multilink packets must be fragmented
}
unit 1 {
   encapsulation multilink-frame-relay-end-to-end;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
   sched { # Scheduling parameters that apply to bundles on AS or Multiservices PICs.
   ...
}
pic-sched {
   # Scheduling parameters for egress DLCIs.
   # The voice queue should be strict-high priority.
   # All other queues should be low priority.
   ...
}
fragmentation-maps {
   fragmap {
      forwarding-class {
         ef {
            no-fragmentation;
         }
         # Voice is carried in the ef queue.
         # It is interleaved with bulk data.
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
    interfaces {
      t3-0/0/0 {
        unit 0 {
          shaping-rate 512k;
          scheduler-map pic-sched;
        }
        unit 1 {
          shaping-rate 128k;
          scheduler-map pic-sched;
        }
      }
      lsq-1/3/0 { # Assign fragmentation and scheduling to LSQ interfaces.
      unit 0 {
        shaping-rate 512k;
        scheduler-map sched;
        fragmentation-map fragmap;
      }
      unit 1 {
        shaping-rate 128k;
        scheduler-map sched;
        fragmentation-map fragmap;
      }
    }
```

For more information about how FRF.12 works with links services IQ interfaces, see "Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12" on page 395.

Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP

This example configures an ATM2 IQ interface with MLPPP bundled with link services IQ interfaces. This allows you to configure LFI on ATM virtual circuits.

For this type of configuration, the ATM2 IQ interface must have LLC encapsulation.

The following ATM PICs are supported in this scenario:

- 2-port OC-3/STM1 ATM2 IQ
- 4-port DS3 ATM2 IQ

Virtual circuit multiplexed PPP over AAL5 is not supported. Frame Relay is not supported. Bundling of multiple ATM VCs into a single logical interface is not supported.

Unlike DS3 and OC3 interfaces, there is no need to create a separate scheduler map for the ATM PIC. For ATM, you define CoS components at the **[edit interfaces at-*fpc/pic/port* atm-options]** hierarchy level, as described in the Junos® OS Network Interfaces.



NOTE: Do not configure RED profiles on ATM logical interfaces that are bundled. Drops do not occur at the ATM interface.

Copyright © 2013, Juniper Networks, Inc.

In this example, two ATM VCs are configured and bundled into two link services IQ bundles. A fragmentation map is used to interleave voice traffic with other multilink traffic. Because MLPPP is used, each link services IQ bundle can be configured for CRTP.

```
[edit interfaces]
at-1/2/0 {
   atm-options {
      vpi 0;
      pic-type atm2;
   }
   unit 0 {
      vci 0.69;
      encapsulation atm-mlppp-llc;
      family mlppp {
         bundle lsq-1/3/0.10;
      }
   }
   unit 1 {
      vci 0.42;
      encapsulation atm-mlppp-llc;
      family mlppp {
         bundle lsq-1/3/0.11;
      }
   }
}
lsq-1/3/0 {
   unit 10 {
      encapsulation multilink-ppp;
   }
   # Large packets must be fragmented.
   # You can specify fragmentation for each forwarding class.
   fragment-threshold 320;
   compression {
      rtp {
         port minimum 2000 maximum 64009;
      }
   }
}
unit 11 {
   encapsulation multilink-ppp;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
   sched { # Scheduling parameters that apply to LSQ bundles on AS or Multiservices PICs.
   ...
}
fragmentation-maps {
   fragmap {
      forwarding-class {
         ef {
            no-fragmentation;
         }
      }
   }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
interfaces { # Assign fragmentation and scheduling parameters to LSQ interfaces.
lsq-1/3/0 {
   unit 0 {
      shaping-rate 512k;
      scheduler-map sched;
      fragmentation-map fragmap;
   }
   unit 1 {
      shaping-rate 128k;
      scheduler-map sched;
      fragmentation-map fragmap;
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 19

Summary of Link Services IQ Configuration Statements

The following sections explain each of the Link Services Intelligent Queuing (IQ) statements. The statements are organized alphabetically.

cisco-interoperability

Syntax cisco-interoperability send-lip-remove-link-for-link-reject;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description FRF.16 interoperability settings.

Options **send-lip-remove-link-for-link-reject**—Send Link Integrity Protocol remove link when an add-link rejection message is received.

Usage Guidelines See "Configuring SONET APS Interoperability with Cisco Systems FRF.16" on page 351.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

forwarding-class

Syntax

```
forwarding-class class-name {
  (fragment-threshold bytes | no-fragmentation);
  multilink-class number;
}
```

Hierarchy Level [edit class-of-service fragmentation-maps]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, define a forwarding class name and associated fragmentation properties within a fragmentation map.

The **fragment-threshold** and **no-fragmentation** statements are mutually exclusive.

Default If you do not include this statement, the traffic in forwarding class ***class-name*** is fragmented.

Options ***class-name***—Name of the forwarding class.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

fragment-threshold

Syntax fragment-threshold *bytes*;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, set the fragmentation threshold for an individual forwarding class.

Default If you do not include this statement, the fragmentation threshold you set at the **[edit interfaces *interface-name* unit *logical-unit-number*]** or **[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest maximum transmission unit (MTU) of all the links in the bundle.

Options *bytes*—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.

Range: 128 through 16,320 bytes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 365

fragmentation-map

Syntax fragmentation-map *map-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI.

Default If you do not include this statement, traffic in all forwarding classes is fragmented.

Options *map-name*—Name of the fragmentation map.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

fragmentation-maps

Syntax

```
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      (fragment-threshold bytes | no-fragmentation);
      multilink-class number;
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, define fragmentation properties for individual forwarding classes.

Default If you do not include this statement, traffic in all forwarding classes is fragmented.

Options *map-name*—Name of the fragmentation map.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

hot-standby

Syntax hot-standby;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options],
[edit interfaces rlsq*number:number* redundancy-options]

Release Information Statement introduced in Junos OS Release 7.6.

Description For one-to-one AS or Multiservices PIC redundancy configurations, specify that the failure detection and recovery must take place in less than 5 seconds. For FRF.15 (MLFR) and FRF.16 (MFR) configuration, specify the switch over time of 5 seconds and less for FRF.15 and a maximum of 10 seconds for FRF.16.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 352.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

link-layer-overhead

Syntax link-layer-overhead *percent*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, configure the percentage of total bundle bandwidth to be set aside for link-layer overhead. Link-layer overhead accounts for the bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information. Overhead resulting from link-layer encapsulation and framing is computed automatically.

Options *percent*—Percentage of total bundle bandwidth to be set aside for link-layer overhead.

Range: 0 through 50 percent

Default: 0 percent

Usage Guidelines See "Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 361.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

lsq-failure-options

Syntax
```
lsq-failure-options {
  no-termination-request;
  trigger-link-failure interface-name;
}
```

Hierarchy Level [edit interfaces lsq-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, define the failure recovery option settings.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring the Association between LSQ and SONET Interfaces" on page 350.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

multilink-class

Syntax multilink-class *number*;

Hierarchy Level [edit class-of-service fragmentation-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, map a forwarding class into a multiclass MLPPP (MCML).

The **multilink-class** statement and **no-fragmentation** statements are mutually exclusive.

Options *number*—The multilink class assigned to this forwarding class.
Range: 0 through 7
Default: None

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 365
- Configuring Multiclass MLPPP on LSQ Interfaces on page 367
- Configuring Fragmentation by Forwarding Class
- Junos Services Interfaces Configuration Release 12.3
- multilink-max-classes on page 306

multilink-max-classes

Syntax multilink-max-classes *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, configure the number of multilink classes to be negotiated when a link joins the bundle.

Options *number*—The number of multilink classes to be negotiated when a link joins the bundle.
Range: 1 through 8
Default: None

Usage Guidelines See "Configuring Multiclass MLPPP on LSQ Interfaces" on page 367.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

no-fragmentation

Syntax no-fragmentation;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services IQ (**lsq**) interfaces only, set traffic on a particular forwarding class to be interleaved, rather than fragmented. This statement specifies that no extra fragmentation header is prepended to the packets received on this queue and that static-link load balancing is used to ensure in-order packet delivery.

Static-link load balancing is done based on packet payload. For IP version 4 (IPv4) and IP version 6 (IPv6) traffic, the link is chosen based on a hash computed from the source address, destination address, and protocol. If the IP payload is Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) traffic, the hash also includes source port and destination port. For MPLS traffic, the hash includes all MPLS labels and fields in the payload, whether the MPLS payload is IPv4 or IPv6.

Default If you do not include this statement, the traffic in forwarding class ***class-name*** is fragmented.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 365

no-per-unit-scheduler

Syntax no-per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 11.4.

Description To enable traffic control profiles to be applied at FRF.16 bundle (physical) interface level, disable the per-unit scheduler, which is enabled by default. This statement and the **shared-scheduler** statement are mutually exclusive.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Oversubscribing Interface Bandwidth

Copyright © 2013, Juniper Networks, Inc.

no-termination-request

Syntax no-termination-request;

Hierarchy Level [edit interfaces *interface-name* ppp-options],
[edit interfaces lsq-*fpc/pic/port* lsq-failure-options]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit interfaces *interface-name* ppp-options]** hierarchy level added in Junos OS Release 8.3.

Description Inhibit PPP termination-request messages to the remote host if the primary circuit fails.

Usage Guidelines See "Configuring the Association between LSQ and SONET Interfaces" on page 350.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

per-unit-scheduler

Syntax per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For channelized OC3 IQ, channelized OC12 IQ, channelized STM1 IQ, channelized T3 IQ, channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), link services (ls-) on J Series routers, Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-Gigabit Ethernet interfaces only, enable association of scheduler map names with logical interfaces.



NOTE: Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.



NOTE: On Gigabit Ethernet IQ2 and IQ2-E PICs without the **per-unit-scheduler** statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the **per-unit-scheduler** statement, the entire PIC supports 1024 – 2 * number of ports (1024 minus two times the number of ports), because each port is allocated two default schedulers.

When including the **per-unit-scheduler** statement for interfaces on the IQ2 and IQ2-E PIC, you must also include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- vlan-tagging

Copyright © 2013, Juniper Networks, Inc.

preserve-interface

Syntax preserve-interface;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced in Junos OS Release 7.6.

Description Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:

- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC

Link PIC replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without triggering link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation.

Usage Guidelines See "Configuring Link State Replication for Redundant Link PICs" on page 356.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Junos® OS Network Interfaces

primary

Syntax primary *interface-name*;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the primary Link Services IQ PIC interface.

Options ***interface-name***—The identifier for the Link Services IQ PIC interface, which must be of the form **lsq-*fpc/pic/port***.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 352.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

redundancy-options

Syntax

```
redundancy-options {
  (hot-standby | warm-standby);
  primary lsq-fpc/pic/port;
  secondary lsq-fpc/pic/port;
}
```

Hierarchy Level [edit interfaces rlsq*number*]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the primary and secondary (backup) Link Services IQ PIC interfaces.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 352.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

secondary

Syntax secondary *interface-name*;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the secondary (backup) Link Services IQ PIC interface.

Options ***interface-name***—The identifier for the Link Services IQ PIC interface, which must be of the form **lsq-*fpc/pic/port***.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 352.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

trigger-link-failure

Syntax trigger-link-failure *interface-name*;

Hierarchy Level [edit interfaces lsq-*fpc/pic/port* lsq-failure-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the Link Services IQ PIC fails.

Options ***interface-name***—Name of SONET interface.

Usage Guidelines See "Configuring the Association between LSQ and SONET Interfaces" on page 350.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

warm-standby

Syntax warm-standby;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in Junos OS Release 8.0.

Description For AS or Multiservices PIC redundancy configurations, specify that the failure detection and recovery involves one backup PIC supporting multiple working PICs. Recovery time is not guaranteed.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 352.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 20

Voice Services Configuration Guidelines

The Adaptive Services (AS) and Multiservices PICs support the compressed Real-Time Transport Protocol (CRTP) on the **lsq-*fpc/pic/port*** interface type. This enables voice over IP (VoIP) traffic to use low-speed links more effectively, by compressing the 40-byte IP/User Datagram Protocol (UDP)/RTP header down to from 2 to 4 bytes in most cases.



NOTE: J Series routers also support VoIP routing through the Avaya TGM550 media gateway module. This is a separate product from the adaptive services suite and is not supported on M Series and T Series routers. For more information, see the *Junos OS Feature Support Reference for SRX Series and J Series Devices*.

For link services IQ interfaces (**lsq**) only, you can configure CRTP with multiclass MLPPP (MCML). MCML greatly simplifies packet ordering issues that occur when multiple links are used. Without MCML, all voice traffic belonging to a single flow is hashed to a single link in order to avoid packet ordering issues. With MCML, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about MCML support on link services IQ interfaces, see "Configuring Link Services and CoS on Services PICs" on page 377.

Link services IQ interfaces use a bundle configuration. For more information, see "Layer 2 Service Package Capabilities and Interfaces" on page 348 and "Multilink and Link Services Logical Interface Configuration Overview" on page 1136.



NOTE: On LSQ interfaces, all multilink traffic for a single bundle is sent to a single processor. If CRTP is enabled on the bundle, it adds overhead to the CPU. Because T3 network interfaces support only one link per bundle, make sure you configure a fragmentation map for compressed traffic on these interfaces and specify the **no-fragmentation** option. For more information, see "Configuring Delay-Sensitive Packet Interleaving" on page 424 and "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

Copyright © 2013, Juniper Networks, Inc.

Voice services do not require a separate service rules configuration, but you need to configure both services interfaces and network interfaces, as described in the following topics:

- Configuring Services Interfaces for Voice Services on page 422
- Configuring Encapsulation for Voice Services on page 425
- Configuring Network Interfaces for Voice Services on page 425
- Examples: Configuring Voice Services on page 426

Configuring Services Interfaces for Voice Services

You define voice service properties such as compression by configuring statements and values for a voice services interface, specified by the interface type **lsq-**. You can include the following statements:

```
encapsulation mlppp;
family inet {
  address address;
}
compression {
  rtp {
    f-max-period number;
    maximum-contexts number <force>;
    port {
      minimum port-number;
      maximum port-number;
    }
    queues [ queue-numbers ];
  }
}
fragment-threshold bytes;
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number*]**

The following sections provide detailed instructions for configuring for voice services on services interfaces:

- Configuring the Logical Interface Address for the MLPPP Bundle on page 422
- Configuring Compression of Voice Traffic on page 423
- Configuring Delay-Sensitive Packet Interleaving on page 424
- Example: Configuring Compression of Voice Traffic on page 424

Configuring the Logical Interface Address for the MLPPP Bundle

To configure the logical address for the MLPPP bundle, include the **address** statement:

```
address address {
```


Copyright © 2013, Juniper Networks, Inc.

```
    ...
}
```

You can configure this statement at the following hierarchy levels:

- **[edit interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number* family inet]**
- **[edit logical-systems *logical-system-name* interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number* family inet]**

address specifies an IP address for the interface. AS and Multiservices PICs support only IP version 4 (IPv4) addresses, which are therefore configured under the **family inet** statement.

For information on other addressing properties you can configure that are not specific to service interfaces, see the Junos® OS Network Interfaces.

Configuring Compression of Voice Traffic

You can specify how a services interface handles voice traffic compression by including the **compression** statement:

```
compression {
  rtp {
    f-max-period number;
    maximum-contexts number <force>;
    port {
      minimum port-number;
      maximum port-number;
    }
    queues [ queue-numbers ];
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces (lsq | ls)-*fpc/pic/port* unit *logical-unit-number*]**

The following statements configure the indicated compression properties:

- **f-max-period *number***—Sets the maximum number of compressed packets to insert between the transmission of full headers. If you do not include the statement, the default is 255 packets.
- **maximum-contexts *number* <force>**—Specifies the maximum number of RTP contexts to accept during negotiation. The optional **force** statement requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option enables interoperation with Junos OS Releases that base the RTP context value on link speed.
- **port**, **minimum *port-number***, and **maximum *port-number***—Specify the lower and upper boundaries for a range of UDP destination port values on which RTP compression takes

Copyright © 2013, Juniper Networks, Inc.

effect. Values for ***port-number*** can range from 0 through 65,535. RTP compression is applied to traffic transiting the ports within the specified range.

- **queues [** ***queue-numbers*** **]**—Specifies one or more of queues **q0**, **q1**, **q2**, and **q3** . RTP compression is applied to the traffic in the specified queues.



NOTE: If you specify both a port range and one or more queues, compression takes place if either condition is met.

Configuring Delay-Sensitive Packet Interleaving

When you configure CRTP, the software automatically enables link fragmentation and interleaving (LFI). LFI reduces excessive delays by fragmenting long packets into smaller packets and interleaving them with real-time frames. This allows real-time and non-real-time data frames to be carried together on lower-speed links without causing excessive delays to the real-time traffic. When the peer interface receives the smaller fragments, it reassembles the fragments into their original packet. For example, short delay-sensitive packets, such as packetized voice, can race ahead of larger delay-insensitive packets, such as common data packets.

By default, LFI is always active when you include the **compression rtp** statement at the **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number*****]** hierarchy level. You control the operation of LFI indirectly by setting the **fragment-threshold** statement on the same logical interface. For example, if you include the **fragment-threshold 256** statement at the **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number*****]** hierarchy level, all IP packets larger than 256 bytes are fragmented.

Example: Configuring Compression of Voice Traffic

Configure compression on a T1 interface with MLPPP encapsulation. Configure fragmentation for all IP packets larger than 128 bytes.

```
[edit interfaces]
t1-1/0/0 {
   unit 0 {
      family mlppp {
         bundle lsq-1/1/0.1;
      }
   }
}
lsq-1/1/0 {
   encapsulation mlppp;
   unit 1 {
      compression {
         rtp {
            port minimum 2000 maximum 64009;
         }
      }
      family inet {
         address 30.1.1.2/24;
      }
```


Copyright © 2013, Juniper Networks, Inc.

```
        fragment-threshold 128;
    }
}
```

Configuring Encapsulation for Voice Services

Voice services interfaces support the following logical interface encapsulation types:

- Multilink Point-to-Point Protocol (MLPPP), which is the default encapsulation
- ATM2 IQ MLPPP over AAL5 LLC
- Frame Relay PPP

For general information on encapsulation, see the Junos® OS Network Interfaces. You can also configure physical interface encapsulation on voice services interfaces.

To configure voice services encapsulation, include the **encapsulation** statement:

```
encapsulation type;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

For voice services interfaces, the valid values for the ***type*** variable are **atm-mlppp-llc**, **frame-relay-ppp** or **multilink-ppp**.

You must also configure the physical interface with the corresponding encapsulation type, either Frame Relay or PPP. LSQ interfaces are supported by the following physical interface types: ATM2 IQ, DS3, E1, E3, OC3, OC12, STM1, and T1, including the channelized versions of these interfaces. For examples, see "Examples: Configuring Voice Services" on page 426.



NOTE: **The only protocol type supported with frame-relay-ppp encapsulation is family mlppp.**

Configuring Network Interfaces for Voice Services

To complete a voice services interface configuration, you need to configure the physical network interface with either MLPPP encapsulation and a voice services bundle or PPP encapsulation and a compression interface, as described in the following sections:

- Configuring Voice Services Bundles with MLPPP Encapsulation on page 426
- Configuring the Compression Interface with PPP Encapsulation on page 426

Copyright © 2013, Juniper Networks, Inc.

Configuring Voice Services Bundles with MLPPP Encapsulation

For voice services interfaces, you configure the link bundle as a channel. The physical interface is usually connected to networks capable of supporting MLPPP; the interface types supported for voice traffic are T1, E1, T3, E3, OC3, OC12, and STM1, including channelized versions of these interfaces.



NOTE:

For M Series routers and T Series routers, the following caveats apply:

- Maximum supported throughput on the bundle interfaces is 45 Mbps.
- Bundling of the logical interfaces under a T3 physical interface into the same or different bundles is not supported.

To configure a physical interface link for MLPPP, include the following statement:

```
bundle interface-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family mlppp]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family mlppp]**

When you configure **family mlppp**, no other protocol configuration is allowed. For more information on link bundles, see "Configuring the Links in a Multilink or Link Services Bundle" on page 1135.

Configuring the Compression Interface with PPP Encapsulation

To configure the physical interface for PPP encapsulation, you also need to specify the services interface to be used for voice compression: a Link Services IQ (**lsq-**) interface.

To configure the compression interface, include the **compression-device** statement:

```
compression-device interface-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit interfaces (lsq | ls)-***fpc/pic/port* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces (lsq | ls)-***fpc/pic/port* **unit** *logical-unit-number***]**

Examples: Configuring Voice Services

Configure voice services using a T1 physical interface and MLPPP bundle encapsulation:

```
[edit interfaces]
t1-0/2/0:1 {
   encapsulation ppp;
```


Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family mlppp {
        bundle lsq-1/3/0.1;
      }
    }
  }
  lsq-1/3/0 {
    unit 1 {
      encapsulation mlppp;
      family inet {
        address 10.5.5.2/30;
      }
      compression {
        rtp {
          f-max-period 100;
          queues [ q1 q2 ];
          port {
            minimum 16384;
            maximum 32767;
          }
        }
      }
      fragment-threshold 128;
    }
  }
```

Configure voice services using Frame Relay encapsulation without bundling:

```
  [edit interfaces]
  t1-1/0/0 {
    encapsulation frame-relay;
    unit 0 {
      dlci 100;
      encapsulation frame-relay-ppp;
      compression-device lsq-2/0/0.0;
    }
  }
  lsq-2/0/0 {
    unit 0 {
      compression {
        rtp {
          f-max-period 100;
          queues [ q1 q2 ];
          port {
            minimum 16000;
            maximum 32000;
          }
        }
      }
      family inet {
        address 10.1.1.1/32;
      }
    }
  }
```

Copyright © 2013, Juniper Networks, Inc.

Configure voice services using an ATM2 physical interface (the corresponding class-of-service configuration is provided for illustration):

```
[edit interfaces]
at-1/2/0 {
  atm-options {
    vpi 0;
    pic-type atm2; # only ATM2 PICs are supported
  }
  unit 0 {
    vci 0.69;
    encapsulation atm-mlppp-llc;
    family mlppp {
      bundle lsq-1/3/0.10;
    }
  }
  unit 1 {
    vci 0.42;
    encapsulation atm-mlppp-llc;
    family mlppp {
      bundle lsq-1/3/0.11;
    }
  }
}
lsq-1/3/0 {
  unit 10 {
    encapsulation multilink-ppp;
  }
  # Large packets need to be fragmented.
  # Fragmentation can also be specified per forwarding class.
  fragment-threshold 320;
  compression {
    rtp {
      port minimum 2000 maximum 64009;
    }
  }
}
unit 11 {
  encapsulation multilink-ppp;
}
fragment-threshold 160;
[edit class-of-service]
scheduler-maps {
  sched {
    # Scheduling parameters apply to bundles on the AS or Multiservices PIC.
    # Unlike DS3/SONET interfaces, there is no need to create
    # a separate scheduler map for the ATM PIC. ATM defines
    # CoS constructs under the [edit interfaces at-fpc/pic/port] hierarchy.
    ...
  }
}
fragmentation-maps {
  fragmap {
    forwarding-class {
      ef {
        # In this example, voice is carried in the ef queue.
```


Copyright © 2013, Juniper Networks, Inc.

```
            # It is interleaved with bulk data.
            # Alternatively, you could use multiclass MLPPP to
            # carry multiple classes of traffic in different
            # multilink classes.
            no-fragmentation;
        }
      }
    }
}
interfaces {
    # Assign fragmentation and scheduling parameters to LSQ interfaces.
    lsq-1/3/0 {
        unit 0 {
            shaping-rate 512k;
            scheduler-map sched;
            fragmentation-map fragmap;
        }
        unit 1 {
            shaping-rate 128k;
            scheduler-map sched;
            fragmentation-map fragmap;
        }
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 21

Summary of Voice Services Configuration Statements

The following sections explain each of the voice services statements. The statements are organized alphabetically.

address (Interfaces)

Syntax

```
address address {
    ...
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the interface address.

Options

address—Address of the interface.

Usage Guidelines

See "Configuring the Logical Interface Address for the MLPPP Bundle" on page 422; for a general discussion of **address** statement options, see the Junos® OS Network Interfaces.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

bundle

Syntax bundle (lsq-*fpc/pic/port* | ...);

Hierarchy Level [edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number* family mlppp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Associate the voice services interface with the logical interface it is joining.

Options **lsq-*fpc/pic/port***—Name of the voice services interface you are linking.

Usage Guidelines See "Configuring Voice Services Bundles with MLPPP Encapsulation" on page 426.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

compression

Syntax

```
compression {
  rtp {
    f-max-period number;
    maximum-contexts number <force>;
    port {
      minimum port-number;
      maximum port-number;
    }
    queues [ queue-numbers ];
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the compression properties for voice services traffic.

The remaining statements are described separately.

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 423.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

compression-device (Interfaces)

Syntax compression-device *interface-name*;

Hierarchy Level [edit interfaces *interface-name* unit (Interfaces) *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 7.5.

Description Specify the compression interface for voice services traffic.

Usage Guidelines See "Configuring the Compression Interface with PPP Encapsulation" on page 426.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

encapsulation

Syntax encapsulation *type*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the logical link-layer encapsulation type.

Options

atm-mlppp-llc—For ATM2 IQ physical interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC encapsulation.

frame-relay-ppp—For Frame Relay circuits, use Frame Relay PPP encapsulation.

multilink-ppp—By default, voice services logical interfaces use MLPPP encapsulation.

Usage Guidelines See "Configuring Encapsulation for Voice Services" on page 425; for information about **encapsulation** statement options used with other interface types, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

f-max-period

Syntax f-max-period *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* compression rtp],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* compression rtp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-time Transport Protocol (RTP) traffic stream.

Options *number*—Maximum number of packets.

Default: 256

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 423.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

family (Interfaces)

Syntax

```
family (inet | mlppp | ...) {
  address address {
    ...
  }
  bundle interface-name;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure protocol family information for the logical interface.

Options ***family***—Protocol family:

- **inet**—IP version 4
- **mlppp**—MLPPP

The remaining statements are explained separately.

Usage Guidelines See "Configuring Network Interfaces for Voice Services" on page 425; for a general discussion of **family** statement options, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

fragment-threshold (Interfaces LSQ)

Syntax fragment-threshold *bytes*;

Hierarchy Level [edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces lsq-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For voice services interfaces, set the fragmentation threshold, in bytes.

Options ***bytes***—Maximum size, in bytes, for multilink packet fragments. The value must be a multiple of 64 bytes, because zero is also a multiple of 64 bytes.

Range: 128 through 16,320 bytes

Default: 0 bytes (no fragmentation)

Usage Guidelines See "Configuring Delay-Sensitive Packet Interleaving" on page 424.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines See the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

maximum-contexts

Syntax maximum-contexts *number* <force>;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* compression rtp],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* compression rtp]

Release Information Statement introduced in Junos OS Release 7.5.

Description Specify the maximum number of RTP contexts to accept during negotiation.

Options *number*—Maximum number of contexts.

force—(Optional) Requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option allows the software to interoperate with Junos OS Releases that base the RTP context value on link speed.

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 423.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

port

Syntax
```
port {
  minimum port-number;
  maximum port-number;
}
```

Hierarchy Level [edit interfaces lsq-*fpc/pic/port* unit *logical-unit-number* compression rtp],
[edit logical-systems *logical-system-name* interfaces lsq-*fpc/pic/port* unit *logical-unit-number* compression rtp]

Release Information Statement introduced before Junos OS Release 7.4.

Description For voice services interfaces only, specify a range of User Datagram Protocol (UDP) destination port numbers in which RTP compression takes place.

Options **minimum *port-number***—Specify the minimum port number.

Range: 0 through 65,535

maximum *port-number*—Specify the maximum port number.

Range: 0 through 65,535

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 423.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

queues

Syntax queues [*queue-numbers*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* compression rtp],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* compression rtp]

Release Information Statement introduced before Junos OS Release 7.4.

Description For voice services interfaces only, assign queue numbers on which RTP compression takes place.

Options **queues *queue-numbers***—Assign one or more of the following queues: **q0**, **q1**, **q2**, and **q3**.

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 423.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

rtp

Syntax

```
rtp {
  f-max-period number;
  maximum-contexts number <force>;
  port {
    minimum port-number;
    maximum port-number;
  }
  queues [ queue-numbers ];
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* compression],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* compression]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the RTP properties for voice services traffic.

The remaining statements are described separately.

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 423.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

unit (Interfaces)

Syntax

```
unit logical-unit-number {
  compression {
    rtp {
      f-max-period number;
      maximum-contexts number <force>;
      port {
        minimum port-number;
        maximum port-number;
      }
      queues [ queue-numbers ];
    }
  }
  compression-device interface-name;
  encapsulation type;
  family family {
    address address {
      ...
    }
    bundle (lsq-fpc/pic/port | ...);
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Usage Guidelines See "Configuring Services Interfaces for Voice Services" on page 422; for a general discussion of logical interface properties, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 22

Class-of-Service Configuration Guidelines

To configure class of service (CoS) features on Adaptive Services (AS) and Multiservices PICs, include the **cos** statement at the **[edit services]** hierarchy level:

```
cos {
  application-profile profile-name {
    ftp {
      data {
        dscp (alias | bits);
        forwarding-class class-name;
      }
    }
    sip {
      video {
        dscp (alias | bits);
        forwarding-class class-name;
      }
      voice {
        dscp (alias | bits);
        forwarding-class class-name;
      }
    }
  }
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        application-sets set-name;
        applications [ application-names ];
        destination-address (CoS) address;
        destination-prefix-list list-name <except>;
        source-address address;
        source-prefix-list list-name <except>;
      }
      then {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        syslog;
        (reflexive | reverse) {
          application-profile profile-name;
          dscp (alias | bits);
          forwarding-class class-name;
```

Copyright © 2013, Juniper Networks, Inc.

```
            syslog;
          }
        }
      }
    }
    rule-set rule-set-name {
      [ rule rule-names ];
    }
  }
```



NOTE: CoS behavior aggregate (BA) classification is not supported on services interfaces.

This chapter contains the following sections:

- Restrictions and Cautions for CoS Configuration on Services Interfaces on page 442
- Configuring CoS Rules on page 443
- Configuring CoS Rule Sets on page 448
- Examples: Configuring CoS on Services Interfaces on page 448

Restrictions and Cautions for CoS Configuration on Services Interfaces

The following restrictions and cautions apply to CoS configuration on services interfaces:

- The adaptive services interface does not support scheduling, only DiffServ marking and queue assignment. You must configure scheduling at the **[edit class-of-service]** hierarchy level on the output interface or fabric.
- In the default configuration, queues 1 and 2 receive 0 percent bandwidth. If packets will be assigned to these queues, you must configure a scheduling map.
- You must issue a **commit full** command before using custom forwarding-class names in the configuration.
- Only the Junos standard DiffServ names can be used in the configuration. Custom names are not recognized.
- On M Series routers, you can configure rewrite rules that change packet headers and attach the rules to output interfaces. These rules might overwrite the DSCP marking configured on an AS or MultiServices PIC. It is important to keep this adverse effect in mind and use care when creating system-wide configurations.

  For example, knowing that the AS or MultiServices PIC can mark packets with any ToS or DSCP value and the output interface is restricted to only eight DSCP values, rewrite rules on the output interface condense the mapping from 64 to 8 values with overall loss of granularity. In this case, you have the following options:

  - Remove the rewrite rules from the output interface.
  - Configure the output interface to include the most important mappings.

Copyright © 2013, Juniper Networks, Inc.

Configuring CoS Rules

To configure a CoS rule, include the **rule *rule-name*** statement at the **[edit services cos]** hierarchy level:

```
[edit services cos]
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (CoS) address;
      destination-prefix-list list-name <except>;
      source-address address;
      source-prefix-list list-name <except>;
    }
    then {
      application-profile profile-name;
      dscp (alias | bits);
      forwarding-class class-name;
      syslog;
      (reflexive | reverse) {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        syslog;
      }
    }
  }
}
```

Each CoS rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how to configure the components of CoS rules:

- Configuring Match Direction for CoS Rules on page 443
- Configuring Match Conditions In CoS Rules on page 444
- Configuring Actions in CoS Rules on page 445
- Example: Configuring CoS Rules on page 447

Configuring Match Direction for CoS Rules

Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services cos rule *rule-name*]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, bidirectional rule creation is allowed.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the AS or Multiservices PIC, the packet direction is output. For more information on inside and outside interfaces, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

On the AS or Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

Configuring Match Conditions In CoS Rules

To configure CoS match conditions, include the **from** statement at the **[edit services cos rule *rule-name* term *term-name*]** hierarchy level:

```
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (CoS) address;
  destination-prefix-list list-name <except>;
  source-address address;
  source-prefix-list list-name <except>;
}
```

The source address and destination address can be either IPv4 or IPv6. You can use either the source address or the destination address as a match condition, in the same way that you would configure a firewall filter; for more information, see the Routing Policy Configuration Guide.

Alternatively, you can specify a list of source or destination prefixes by configuring the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or **source-prefix-list** statement in the CoS rule. For an example, see "Examples: Configuring Stateful Firewall Rules" on page 136.

If you omit the **from** term, the router accepts all traffic and the default protocol handlers take effect:

- User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP) create a bidirectional flow with a predicted reverse flow.
- IP creates a unidirectional flow.

Copyright © 2013, Juniper Networks, Inc.

You can also include application protocol definitions you have configured at the **[edit applications]** hierarchy level; for more information, see "Configuring Application Protocol Properties" on page 74.

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services cos rule *rule-name* term *term-name* from]** hierarchy level.
- To apply one or more sets of application protocol definitions you have defined, include the **application-sets** statement at the **[edit services cos rule *rule-name* term *term-name* from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit applications]** hierarchy level; you cannot specify these properties as match conditions.

Configuring Actions in CoS Rules

To configure CoS actions, include the **then** statement at the **[edit services cos rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services cos rule rule-name term term-name]
then {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  syslog;
  (reflexive | reverse) {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    syslog;
  }
}
```

The principal CoS actions are as follows:

- **dscp**—Causes the packet to be marked with the specified DiffServ code point (DSCP) value or alias.
- **forwarding-class**—Causes the packet to be assigned to the specified forwarding class.

For detailed information about DSCP values and forwarding classes, see "Examples: Configuring CoS on Services Interfaces" on page 448 or the Junos OS Class of Service Configuration Guide.

You can optionally set the configuration to record information in the system logging facility by including the **syslog** statement at the **[edit services cos rule *rule-name* term *term-name* then]** hierarchy level. This statement overrides any **syslog** setting included in the service set or interface default configuration.

Copyright © 2013, Juniper Networks, Inc.

For information about some additional CoS actions, see the following sections:

- Configuring Application Profiles for Use as CoS Rule Actions on page 446
- Configuring Reflexive and Reverse CoS Rule Actions on page 446

Configuring Application Profiles for Use as CoS Rule Actions

You can optionally define one or more application profiles for inclusion in CoS actions. To configure application profiles, include the **application-profile** statement at the **[edit services cos]** hierarchy level:

```
[edit services cos]
application-profile profile-name {
  ftp {
    data {
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
  sip {
    video {
      dscp (alias | bits);
      forwarding-class class-name;
    }
    voice {
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
}
```

The **application-profile** statement includes two main components and three traffic types: **ftp** with the **data** traffic type and **sip** with the **video** and **voice** traffic types. You can set the appropriate **dscp** and **forwarding-class** values for each component within the application profile.



NOTE: The **ftp** and **sip** statements are not supported on Juniper Network MX Series 3D Universal Edge Routers.

You can apply the application profile to a CoS configuration by including it at the **[edit services cos rule *rule-name* term *term-name* then]** hierarchy level.

Configuring Reflexive and Reverse CoS Rule Actions

CoS services are unidirectional. It might be necessary to specify different treatments for flows in opposite directions.

Regardless of whether a packet matches the input, output or input-output direction, flows in both directions are created. A forward, reverse, or forward-and-reverse CoS action is associated with each flow. Bear in mind that the flow in the opposite direction might end up having a CoS action associated with it that you have not specifically configured.

Copyright © 2013, Juniper Networks, Inc.

To control the direction in which service is applied, as distinct from the direction in which the rule match is applied, you can configure the (**reflexive | reverse**) statement at the **[edit services cos rule *rule-name* term *term-name* then]** hierarchy level:

```
[edit services cos rule rule-name term term-name then]
(reflexive | reverse) {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  syslog;
}
```

The two actions are mutually exclusive:

- **reflexive** causes the equivalent opposing CoS action to be applied to flows in the opposite direction.
- **reverse** allows you to define the CoS behavior for flows in the reverse direction.

If you omit the statement, data flows inherit the CoS behavior of the forward control flow.

Example: Configuring CoS Rules

The following example show a CoS configuration containing two rules, one for input matching on a specified application set and the other for output matching on a specified source address:

```
[edit services]
cos {
  rule my-cos-rule {
    match-direction input-output;
    term term1 {
      from {
        source-address 10.1.3.2/32;
        applications sip;
      }
      then {
        dscp ef;
        syslog;
      }
    }
    term term2 {
      from {
        destination-address 10.2.3.2;
        applications http;
      }
      then {
        dscp af21;
      }
    }
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

Configuring CoS Rule Sets

The **rule-set** statement defines a collection of CoS rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then you specify the order of the rules by including the **rule-set** statement at the **[edit services cos]** hierarchy level with a **rule** statement for each rule:

```
rule-set rule-set-name {
  rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Examples: Configuring CoS on Services Interfaces

To make settings consistent across Juniper Networks routers, you configure many CoS settings at the **[edit class-of-service]** hierarchy level to be used on services interfaces. When you commit this configuration along with what you configure at the **[edit services cos]** hierarchy level, these properties are applied to the AS or MultiServices PIC.

The following configuration examples at the **[edit class-of-service]** hierarchy level can be applied on services interfaces. For more information, see the Junos OS Class of Service Configuration Guide.



NOTE: The first two configurations, mapping forwarding-class name to forwarding-class ID and mapping forwarding-class name to queue number, are mutually exclusive.

Mapping Forwarding-Class Name to Forwarding-Class ID

Map forwarding-class names to forwarding-class IDs:

```
[edit class-of-service]
forwarding-classes {
   forwarding-class fc0 0;
   forwarding-class fc1 0;
   forwarding-class fc2 1;
   forwarding-class fc3 1;
   forwarding-class fc4 2;
   forwarding-class fc5 2;
   forwarding-class fc6 3;
   forwarding-class fc7 3;
   forwarding-class fc8 4;
   forwarding-class fc9 4;
   forwarding-class fc10 5;
   forwarding-class fc11 5;
   forwarding-class fc12 6;
```


Copyright © 2013, Juniper Networks, Inc.

```
    forwarding-class fc13 6;
    forwarding-class fc14 7;
    forwarding-class fc15 7;
}
```

Mapping Forwarding-Class Name to Queue Number

Map forwarding-class names to queue numbers:

```
[edit class-of-service]
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
    queue 4 ef1;
    queue 5 ef2;
    queue 6 af1;
    queue 7 nc1;
}
```

Mapping Diffserv Code Point Aliases to DSCP Bits

Map alias names to DSCP bit values. The aliases then can be used instead of the DSCP bits in adaptive services configurations.

```
[edit class-of-service]
code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) {
        alias | bits;
    }
}
```

Here is an example:

```
code-point-aliases {
    dscp {
        my1 110001;
        my2 101110;
        be 000001;
        cs7 110000;
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 23

Summary of Class-of-Service Configuration Statements

The following sections explain each of the class-of-service (CoS) statements. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

application-profile

Syntax

```
application-profile profile-name {
  ftp {
    data {
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
  sip {
    video {
      dscp (alias | bits);
      forwarding-class class-name;
    }
    voice {
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
}
```

Hierarchy Level

[edit services cos],
[edit services cos rule *rule-name* term *term-name* then],
[edit services cos rule *rule-name* term *term-name* then (reflexive | reverse)]

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Define or apply a CoS application profile. When you apply a CoS application profile in a CoS rule, terminate the profile name with a semicolon (;).

Options

profile-name—Identifier for the application profile.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles for Use as CoS Rule Actions on page 446

Copyright © 2013, Juniper Networks, Inc.

application-sets (Services CoS)

Syntax applications-sets *set-name*;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define one or more target application sets.

Options ***set-name***—Name of the target application set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions In CoS Rules on page 444

applications (Services CoS)

Syntax applications [*application-name*];

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define one or more applications to which the CoS services apply.

Options ***application-name***—Name of the target application.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in a CoS Rule
- Configuring Match Conditions In CoS Rules on page 444

Copyright © 2013, Juniper Networks, Inc.

data (FTP)

Syntax

```
data {
  dscp (alias | bits);
  forwarding-class class-name;
}
```

Hierarchy Level [edit services cos application-profile *profile-name* ftp]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the appropriate **dscp** and **forwarding-class** value for FTP data.

Default By default, the system will not alter the DSCP or forwarding class for FTP data traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles
- video (Application Profile) on page 464
- voice (Application Profile) on page 465

destination-address (CoS)

Syntax destination-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.1.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Specify the destination address for rule matching.

Options ***address***—Destination IPv4 or IPv6 address or prefix value.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Match Conditions in a CoS Rule
- Configuring Match Conditions In CoS Rules on page 444

Copyright © 2013, Juniper Networks, Inc.

destination-prefix-list (Services CoS)

Syntax destination-prefix-list *list-name* <except>;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the destination prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Usage Guidelines See "Configuring Match Conditions In CoS Rules" on page 444.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Routing Policy Configuration Guide

dscp

Syntax dscp (*alias* | *bits*);

Hierarchy Level [edit services cos application-profile *profile-name* ftp data],
[edit services cos application-profile *profile-name* sip (video | voice)],
[edit services cos rule *rule-name* term *term-name* then],
[edit services cos rule *rule-name* term *term-name* then (reflexive | reverse)]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.

Options ***alias***—Name assigned to a set of CoS markers.

bits—Mapping value in the packet header.

Usage Guidelines See "Configuring Actions in CoS Rules" on page 445.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

forwarding-class

Syntax forwarding-class *class-name*;

Hierarchy Level [edit services cos application-profile *profile-name* ftp data],
[edit services cos application-profile *profile-name* sip (video | voice)],
[edit services cos rule *rule-name* term *term-name* then],
[edit services cos rule *rule-name* term *term-name* then (reflexive | reverse)]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the forwarding class to which packets are assigned.

Options *class-name*—Name of the target application.

Usage Guidelines See "Configuring Actions in CoS Rules" on page 445.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

from (Services CoS)

Syntax

```
from {
  application-sets set-name;
  applications [ application-names ];
  destination-address (CoS) address;
  destination-prefix-list list-name <except>;
  source-address address;
  source-prefix-list list-name <except>;
}
```

Hierarchy Level [edit services cos rule *rule-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Specify input conditions for a CoS term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Rules" on page 443.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ftp (Class-Of-Service)

Syntax

```
ftp {
  data {
    dscp (alias | bits);
    forwarding-class class-name;
  }
}
```

Hierarchy Level [edit services cosapplication-profile *profile-name* ftp]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the appropriate **dscp** and **forwarding-class** value for FTP.

Default By default, the system does not alter the DSCP or forwarding class for FTP traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles
- sip (Application Profile) on page 461

match-direction

Syntax match-direction (input | output | input-output);

Hierarchy Level [edit services cos rule *rule-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Specify the direction in which the rule match is applied.

Options

input—Apply the rule match on the input side of the interface.

output—Apply the rule match on the output side of the interface.

input-output—Apply the rule match bidirectionally.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring CoS Rules

Copyright © 2013, Juniper Networks, Inc.

(reflexive | reverse)

Syntax

```
(reflexive | reverse) {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  syslog;
}
```

Hierarchy Level [edit services cos rule *rule-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 8.1.

Description **reflexive**—Applies the equivalent opposing CoS action to flows in the opposite direction.

reverse—Allows you to define CoS behavior for flows in the reverse direction.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring CoS Rules
- Configuring Reflexive and Reverse CoS Rule Actions on page 446

Copyright © 2013, Juniper Networks, Inc.

rule

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-names ];
      destination-address (CoS) address;
      destination-prefix-list list-name <except>;
      source-address address;
      source-prefix-list list-name <except>;
    }
    then {
      application-profile profile-name;
      dscp (alias | bits);
      forwarding-class class-name;
      syslog;
      (reflexive | reverse) {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
        syslog;
      }
    }
  }
}
```

Hierarchy Level

[edit services cos],
[edit services cos rule-set *rule-set-name*]

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring CoS Rules" on page 443.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rule-set (Services CoS)

Syntax

```
rule-set rule-set-name {
    [ rule rule-name ];
}
```

Hierarchy Level [edit services cos]

Release Information Statement introduced in Junos OS Release 8.1.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring CoS Rule Sets

services (COS)

Syntax services cos { ... }

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the service rules to be applied to traffic.

Options **cos**—Identifier for the class-of-service set of rules statements.

Usage Guidelines See Class-of-Service Properties.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

sip (Application Profile)

Syntax

```
sip {
  video {
    dscp (alias | bits);
    forwarding-class class-name;
  }
  voice {
    dscp (alias | bits);
    forwarding-class class-name;
  }
}
```

Hierarchy Level [edit services cos application-profile*profile-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the appropriate **dscp** and **forwarding-class** value for SIP traffic.

Default By default, the system will not alter the DSCP or forwarding class for SIP traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles
- ftp (Class-Of-Service) on page 457

source-address (Services CoS)

Syntax source-address *address*;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.1.
address option enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description Source address for rule matching.

Options ***address***—Source IPv4 or IPv6 address or prefix value.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Match Conditions in a CoS Rule
- Configuring Match Conditions In CoS Rules on page 444

Copyright © 2013, Juniper Networks, Inc.

source-prefix-list

Syntax source-prefix-list *list-name* <except>;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the source prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Usage Guidelines See "Configuring CoS Rules" on page 443.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Routing Policy Configuration Guide

syslog (Services CoS)

Syntax syslog;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* then],
[edit services cos rule *rule-name* term *term-name* then (reflexive | reverse)]

Release Information Statement introduced in Junos OS Release 8.1.

Description Enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the **/var/log** directory. This setting overrides any **syslog** statement setting included in the service set or interface default configuration.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Actions in a CoS Rule
- Configuring Actions in CoS Rules on page 445

Copyright © 2013, Juniper Networks, Inc.

term (Services CoS)

Syntax

```
term term-name {
  from {
    application-sets set-name;
    applications [ application-names ];
    destination-address (CoS) address;
    destination-prefix-list list-name <except>;
    source-address address;
    source-prefix-list list-name <except>;
  }
  then {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    syslog;
    (reflexive | reverse) {
      application-profile profile-name;
      dscp (alias | bits);
      forwarding-class class-name;
      syslog;
    }
  }
}
```

Hierarchy Level [edit services cos rule *rule-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the CoS term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Rules" on page 443.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

then

Syntax

```
then {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  syslog;
  (reflexive | reverse) {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    syslog;
  }
}
```

Hierarchy Level [edit services cos rule *rule-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the CoS term actions.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Actions in a CoS Rule
- Configuring Actions in CoS Rules on page 445

video (Application Profile)

Syntax

```
video {
  dscp (alias | bits);
  forwarding-class class-name;
}
```

Hierarchy Level [edit services cos application-profile*profile-name* sip]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the appropriate **dscp** and **forwarding-class** values for SIP video traffic.

Default By default, the system will not alter the DSCP or forwarding class for SIP video traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles
- voice (Application Profile) on page 465

Copyright © 2013, Juniper Networks, Inc.

voice (Application Profile)

Syntax

```
voice {
  dscp (alias | bits);
  forwarding-class class-name;
}
```

Hierarchy Level [edit services cos application-profile*profile-name* sip]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the appropriate **dscp** and **forwarding-class** values for SIP voice traffic.

Default By default, the system will not alter the DSCP or forwarding class for SIP voice traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles
- video (Application Profile) on page 464

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 24

Service Set Configuration Guidelines

A *service set* is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. To configure service sets, include the following statements at the **[edit services]** hierarchy level:

```
[edit services]
service-set service-set-name {
   (ids-rules rule-names | ids-rule-sets rule-set-name);
   (ipsec-vpn-rules rule-names | ipsec-vpn-rule-sets rule-set-name);
   (nat-rules rule-names | nat-rule-sets rule-set-name);
   (pgcp-rules rule-names | pgcp-rule-sets rule-set-name);
   (ptsp-rules rule-names | ptsp-rule-sets rule-set-name);
   (stateful-firewall-rules rule-names | stateful-firewall-rule-sets rule-set-name);
   allow-multicast;
   extension-service service-name {
      provider-specific rules;
   }
   interface-service {
      service-interface interface-name;
   }
   ipsec-vpn-options {
      anti-replay-window-size bits;
      clear-dont-fragment-bit;
      ike-access-profile profile-name;
      local-gateway address;
      no-anti-replay;
      passive-mode-tunneling;
      trusted-ca [ ca-profile-names ];
      tunnel-mtu bytes;
   }
   max-flows number;
   next-hop-service {
      inside-service-interface interface-name.unit-number;
      outside-service-interface interface-name.unit-number;
      service-interface-pool name;
   }
   syslog {
      host hostname {
         services severity-level;
         facility-override facility-name;
         log-prefix prefix-value;
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  adaptive-services-pics {
    traceoptions {
      file filename <files number> <match regex> <size size> <(world-readable |
        no-world-readable)>;
      flag flag;
    }
  }
  logging {
    traceoptions {
      file filename <files number> <match regex> <size size> <(world-readable |
        no-world-readable)>;
      flag flag;
    }
  }
```

This chapter contains the following sections:

- Configuring Service Sets to be Applied to Services Interfaces on page 468
- Configuring Service Rules on page 472
- Configuring IPsec Service Sets on page 473
- Configuring Service Set Limitations on page 478
- Configuring System Logging for Service Sets on page 479
- Enabling Services PICs to Accept Multicast Traffic on page 480
- Tracing Services PIC Operations on page 480
- Example: Configuring Service Sets on page 483

Configuring Service Sets to be Applied to Services Interfaces

You configure a services interface to specify the adaptive services interface on which the service is to be performed. Services interfaces are used with either of the service set types described in the following sections.

- Configuring Interface Service Sets on page 468
- Configuring Next-Hop Service Sets on page 470
- Determining Traffic Direction on page 471

Configuring Interface Service Sets

An interface service set is used as an action modifier across an entire interface. To configure the services interface, include the **interface-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level:

```
[edit services service-set service-set-name]
interface-service {
  service-interface interface-name;
}
```

Only the device name is needed, because the router software manages logical unit numbers automatically. The services interface must be an adaptive services interface

Copyright © 2013, Juniper Networks, Inc.

for which you have configured **unit 0 family inet** at the **[edit interfaces *interface-name*** hierarchy level.

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces installed on the router. When you apply the service set to an interface, it automatically ensures that packets are directed to the PIC.

To associate a defined service set with an interface, include a **service-set** statement with the **input** or **output** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet service]
input {
   service-set service-set-name <service-filter filter-name>;
   post-service-filter filter-name;
}
output {
   service-set service-set-name <service-filter filter-name>;
}
```

If a packet is entering the interface, the match direction is **input**. If a packet is leaving the interface, the match direction is **output**. The service set retains the input interface information even after services are applied, so that functions such as filter-class forwarding and destination class usage (DCU) that depend on input interface information continue to work.

You configure the same service set on the input and output sides of the interface. You can optionally include filters associated with each service set to refine the target and additionally process the traffic. If you include the **service-set** statement without a **service-filter** definition, the router software assumes the match condition is true and selects the service set for processing automatically.



NOTE: If you configure service sets with filters, they must be configured on the input and output sides of the interface.

You can include more than one service set definition on each side of the interface. If you include multiple service sets, the router software evaluates them in the order in which they appear in the configuration. The system executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions. A maximum of six service sets can be applied to an interface. When you apply multiple service sets to an interface, you must also configure and apply a service filter to the interface.

An additional statement allows you to specify a filter for processing the traffic after the input service set is executed. To configure this type of filter, include the **post-service-filter** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service input]** hierarchy level:

```
post-service-filter filter-name;
```

For an example, see "Example: Configuring Service Sets" on page 483.

Copyright © 2013, Juniper Networks, Inc.

NOTE: When the MultiServices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the **bypass-traffic-on-pic-failure** statement at the **[edit services service-set *service-set-name* service-set-options]** hierarchy level. When this statement is configured, the affected packets are forwarded in the event of a MultiServices PIC failure or offlining, as though interface-style services were not configured. This issue applies only Dynamic Application Awareness for Junos OS configurations using IDP service sets. This forwarding feature worked only with the Packet Forwarding Engine (PFE) initially. Starting with Junos OS Release 11.3, the packet-forwarding feature is extended to packets generated by the Routing Engine for bypass service sets as well.

Configuring Next-Hop Service Sets

A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes. This configuration is useful when services need to be applied to an entire virtual private network (VPN) routing and forwarding (VRF) table, or when routing decisions determine that services need to be performed.

When a next-hop service is configured, the AS or Multiservices PIC is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).

NOTE: You can create IFL indexes greater than 8000 only if the interface service set is not configured.

To configure the domain, include the **service-domain** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
service-domain (inside | outside);
```

The **service-domain** setting must match the configuration for the next-hop service inside and outside interfaces. To configure the inside and outside interfaces, include the **next-hop-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level. The interfaces you specify must be logical interfaces on the same AS PIC. You cannot configure **unit 0** for this purpose, and the logical interface you choose must not be used by another service set.

```
next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
}
```

Traffic on which the service is applied is forced to the inside interface using a static route. For example:

Copyright © 2013, Juniper Networks, Inc.

```
routing-options {
   static {
      route 10.1.2.3 next-hop sp-1/1/0.1;
   }
}
```

After the service is applied, traffic exits by way of the outside interface. A lookup is then performed in the Packet Forwarding Engine (PFE) to send the packet out of the AS or Multiservices PIC.

The reverse traffic enters the outside interface, is serviced, and sent to the inside interface. The inside interface forwards the traffic out of the AS or Multiservices PIC.

Determining Traffic Direction

When you configure next-hop service sets, the AS PIC functions as a two-part interface, in which one part is the *inside* interface and the other part is the *outside* interface. The following sequence of actions takes place:

1. To associate the two parts with logical interfaces, you configure two logical interfaces with the **service-domain** statement, one with the **inside** value and one with the **outside** value, to mark them as either an inside or outside service interface.
2. The router forwards the traffic to be serviced to the inside interface, using the next-hop lookup table.
3. After the service is applied, the traffic exits from the outside interface. A route lookup is then performed on the packets to be sent out of the router.
4. When the reverse traffic returns on the outside interface, the applied service is undone; for example, IPsec traffic is decrypted or NAT addresses are unmasked. The serviced packets then emerge on the inside interface, the router performs a route lookup, and the traffic exits the router.

A service rule's match direction, whether input, output, or input/output, is applied with respect to the traffic flow through the AS PIC, not through a specific inside or outside interface.

When a packet is sent to an AS PIC, packet direction information is carried along with it. This is true for both interface style and next-hop style service sets.

Interface Style Service Sets

Packet direction is determined by whether a packet is entering or leaving any Packet Forwarding Engine interface (with respect to the forwarding plane) on which the **interface-service** statement is applied. This is similar to the input and output direction for stateless firewall filters.

The match direction can also depend on the network topology. For example, you might route all the external traffic through one interface that is used to protect the other interfaces on the router, and configure various services on this interface specifically. Alternatively, you might use one interface for priority traffic and configure special services on it, but not care about protecting traffic on the other interfaces.

Copyright © 2013, Juniper Networks, Inc.

Next-Hop Style Service Sets

Packet direction is determined by the AS PIC interface used to route packets to the AS PIC. If you use the **inside-interface** statement to route traffic, then the packet direction is **input**. If you use the **outside-interface** statement to direct packets to the AS PIC, then the packet direction is **output**.

The interface to which you apply the service sets affects the match direction. For example, apply the following configuration:

```
sp-1/1/0 unit 1 service-domain inside;
sp-1/1/0 unit 2 service-domain outside;
```

If you configure **match-direction input**, you include the following statements:

```
[edit]
services service-set test1 next-hop-service inside-service-interface sp-1/0/0.1;
services service-set test1 next-hop-service outside-service-interface sp-1/0/0.2;
services ipsec-vpn rule test-ipsec-rule match-direction input;
routing-options static route 10.0.0.0/24 next-hop sp-1/1/0.1;
```

If you configure **match-direction output**, you include the following statements:

```
[edit]
services service-set test2 next-hop-service inside-service-interface sp-1/0/0.1;
services service-set test2 next-hop-service outside-service-interface sp-1/0/0.2;
services ipsec-vpn rule test-ipsec-rule match-direction output;
routing-options static route 10.0.0.0/24 next-hop sp-1/1/0.2;
```

The essential difference between the two configurations is the change in the match direction and the static routes' next hop, pointing to either the AS PIC's inside or outside interface.

Configuring Service Rules

You specify the collection of rules and rule sets that constitute the service set. The router performs rule sets in the order in which they appear in the configuration. You can include only one rule set for each service type. You configure the rule names and content for each service type at the **[edit services *name*]** hierarchy level for each type:

- You configure intrusion detection service (IDS) rules at the **[edit services ids]** hierarchy level; for more information, see "Configuring IDS Rules" on page 187.
- You configure IP Security (IPsec) rules at the **[edit services ipsec-vpn]** hierarchy level; for more information, see IPsec Properties.
- You configure Network Address Translation (NAT) rules at the **[edit services nat]** hierarchy level; for more information, see Network Address Translation.
- You configure packet-triggered subscribers and policy control (PTSP) rules at the **[edit services ptsp]** hierarchy level; for more information, see PTSP for Subscriber Access.

Copyright © 2013, Juniper Networks, Inc.

- You configure softwire rules for DS-Lite or 6rd softwires at the **[edit services softwire]** hierarchy level; for more information, see Softwire Services.
- You configure stateful firewall rules at the **[edit services stateful-firewall]** hierarchy level; for more information, see Stateful Firewall for Juniper Service Framework (JSF).

To configure the rules and rule sets that constitute a service set, include the following statements at the **[edit services service-set *service-set-name*]** hierarchy level:

```
([ ids-rules rule-names ] | ids-rule-sets rule-set-name);
([ ipsec-vpn-rules rule-names ] | ipsec-vpn-rule-sets rule-set-name);
([ nat-rules rule-names ] | nat-rule-sets rule-set-name);
([ pgcp-rules rule-names] | pgcp-rule-sets rule-set-name);
([softwire-rules rule-names] | softwire-rule-sets rule-set-name);
([ stateful-firewall-rules rule-names ] | stateful-firewall-rule-sets rule-set-name);
```

For each service type, you can include one or more individual rules, or one rule set.

If you configure a service set with IPsec rules, it must not contain rules for any other services. You can, however, configure another service set containing rules for the other services and apply both service sets to the same interface.

NOTE: You can also include Dynamic Application Awareness for Junos OS functionality within service sets. To do this, you must include an **idp-profile** statement at the **[edit services service-set]** hierarchy level, along with application identification (APPID) rules, and, as appropriate, application-aware access list (AACL) rules and a **policy-decision-statistics-profile**. Only one service sets can be applied to a single interface when Dynamic Application Awareness functionality is used. For more information, see Intrusion Detection and Prevention, Application Identification, and Application-Aware Access List.

Configuring IPsec Service Sets

IPsec service sets require additional specifications that you configure at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level:

```
[edit services service-set service-set-name ipsec-vpn-options]
anti-replay-window-size bits;
clear-dont-fragment-bit;
ike-access-profile profile-name;
local-gateway address;
no-anti-replay;
passive-mode-tunneling;
trusted-ca [ ca-profile-names ];
tunnel-mtu bytes;
```

Configuration of these statements is described in the following sections:

- Configuring the Local Gateway Address for IPsec Service Sets on page 474
- Configuring IKE Access Profiles for IPsec Service Sets on page 475

Copyright © 2013, Juniper Networks, Inc.

- Configuring Certification Authorities for IPsec Service Sets on page 475
- Configuring or Disabling Antireplay Service on page 476
- Clearing the Don't-Fragment Bit on page 476
- Configuring Passive-Mode Tunneling on page 477
- Configuring the Tunnel MTU Value on page 477

Configuring the Local Gateway Address for IPsec Service Sets

If you configure an IPsec service set, you must also configure a local IPv4 or IPv6 address by including the **local-gateway** statement:

- If the Internet Key Exchange (IKE) gateway IP address is in **inet.0** (the default situation), you configure the following statement:

  local-gateway *address*;

- If the IKE gateway IP address is in a VPN routing and forwarding (VRF) instance, you configure the following statement:

  local-gateway *address* routing-instance *instance-name*;

You can configure all the link-type tunnels that share the same local gateway address in a single next-hop-style service set. The value you specify for the **inside-service-interface** statement at the **[edit services service-set *service-set-name*]** hierarchy level should match the **ipsec-inside-interface** value, which you configure at the **[edit services ipsec-vpn rule *rule-name* term *term-name* from]** hierarchy level. For more information about IPsec configuration, see "Configuring IPsec Rules" on page 243.

IKE Addresses in VRF Instances

You can configure Internet Key Exchange (IKE) gateway IP addresses that are present in a VPN routing and forwarding (VRF) instance as long as the peer is reachable through the VRF instance.

For next-hop service sets, the key management process (kmd) places the IKE packets in the routing instance that contains the **outside-service-interface** value you specify, as in this example:

```
routing-instances vrf-nxthop {
    instance-type vrf;
    interface sp-1/1/0.2;
    ...
}
services service-set service-set-1 {
    next-hop-service {
        inside-service-interface sp-1/1/0.1;
        outside-service-interface sp-1/1/0.2;
    }
    ...
}
```

For interface service sets, the **service-interface** statement determines the VRF, as in this example:

Copyright © 2013, Juniper Networks, Inc.

```
routing-instances vrf-intf {
   instance-type vrf;
   interface sp-1/1/0.3;
   interface ge-1/2/0.1; # interface on which service set is applied
   ...
}
services service-set service-set-2 {
   interface-service {
      service-interface sp-1/1/0.3;
   }
   ...
}
```

Configuring IKE Access Profiles for IPsec Service Sets

For dynamic endpoint tunneling only, you need to reference the IKE access profile configured at the **[edit access]** hierarchy level. To do this, include the **ike-access-profile** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level:

```
[edit services service-set service-set-name ipsec-vpn-options]
ike-access-profile profile-name;
```

The **ike-access-profile** statement must reference the same name as the **profile** statement you configured for IKE access at the **[edit access]** hierarchy level. You can reference only one access profile in each service set. This profile is used to negotiate IKE and IPsec security associations with dynamic peers only.



NOTE: If you configure an IKE access profile in a service set, no other service set can share the same **local-gateway** address.

Also, you must configure a separate service set for each VRF. All interfaces referenced by the **ipsec-inside-interface** statement within a service set must belong to the same VRF.

Configuring Certification Authorities for IPsec Service Sets

You can specify one or more trusted certification authorities by including the **trusted-ca** statement:

```
trusted-ca [ ca-profile-names ];
```

When you configure public key infrastructure (PKI) digital certificates in the IPsec configuration, each service set can have its own set of trusted certification authorities. The names you specify for the **trusted-ca** statement must match profiles configured at the **[edit security pki]** hierarchy level; for more information, see the Junos OS System Basics Configuration Guide. For more information about IPsec digital certificate configuration, see "Configuring IPsec Rules" on page 243.

Copyright © 2013, Juniper Networks, Inc.

Configuring or Disabling Antireplay Service

You can include the **anti-replay-window-size** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level to specify the size of the antireplay window.

anti-replay-window-size *bits*;

This statement is useful for dynamic endpoint tunnels for which you cannot configure the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement sets the antireplay window size for all the static tunnels within this service set. If a particular tunnel needs a specific value for antireplay window size, set the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level. If antireplay check has to be disabled for a particular tunnel in this service set, set the **no-anti-replay** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: The **anti-replay-window-size** and **no-anti-replay** settings at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level override the settings specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

You can also include the **no-anti-replay** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level to disable IPsec antireplay service. It occasionally causes interoperability issues for security associations.

no-anti-replay;

This statement is useful for dynamic endpoint tunnels for which you cannot configure the **no-anti-reply** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement disables the antireplay check for all the tunnels within this service set. If antireplay check has to be enabled for a particular tunnel, then set the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: Setting the **anti-replay-window-size** and **no-anti-replay** statements at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level overrides the settings specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

Clearing the Don't-Fragment Bit

You can include the **clear-dont-fragment-bit** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level to clear the Don't Fragment (DF) bit

Copyright © 2013, Juniper Networks, Inc.

on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation.

clear-dont-fragment-bit;

This statement is useful for dynamic endpoint tunnels, for which you cannot configure the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, setting this statement clears the DF bit on packets entering all the static tunnels within this service set. If you want to clear the DF bit on packets entering a specific tunnel, set the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

Configuring Passive-Mode Tunneling

You can include the **passive-mode-tunneling** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level to enable the service set to tunnel malformed packets.

[edit services service-set *service-set-name* ipsec-vpn-options]
passive-mode-tunneling;

This functionality bypasses the active IP checks, such as version, TTL, protocol, options, address and other land attack checks, and tunnels the packets as is. If this statement is not configured, packets failing the IP checks are dropped in the PIC. In passive mode, the inner packet is not touched; hence, an ICMP error is not generated, if the packet size exceeds the tunnel MTU value.

The IPsec tunnel is not treated as a next hop and TTL is not decremented. Because an ICMP error is not generated if the packet size exceeds the tunnel MTU value, the packet will be tunnelled even if it crosses the tunnel MTU threshold.



NOTE: This functionality is similar to that provided by the **no-ipsec-tunnel-in-traceroute** statement, described in "Disabling IPsec Tunnel Endpoint in Traceroute" on page 256.

Configuring the Tunnel MTU Value

You can include the **tunnel-mtu** statement at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level to set the maximum transmission unit (MTU) value for IPsec tunnels.

tunnel-mtu *bytes*;

This statement is useful for dynamic endpoint tunnels for which you cannot configure the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement sets the tunnel MTU value for all the tunnels within this service set. If you need a specific value for a particular tunnel, then set the

Copyright © 2013, Juniper Networks, Inc.

tunnel-mtu statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: The **tunnel-mtu** setting at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level overrides the value specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

Configuring Service Set Limitations

You can set the following limitations on service set capacity:

- You can limit the maximum number of flows allowed per service set. To configure the maximum value, include the **max-flows** statement at the **[edit services service-set *service-set-name*]** hierarchy level:

  max-flows *number*;

  The **max-flows** statement permits you to assign a single flow limit value. For IDS service sets only, you can specify various types of flow limits with a finer degree of control. For more information, see the description of the **session-limit** statement in "Configuring IDS Rule Sets" on page 193.

- You can limit the maximum segment size (MSS) allowed by the Transmission Control Protocol (TCP). To configure the maximum value, include the **tcp-mss** statement at the **[edit services service-set *service-set-name*]** hierarchy level:

  tcp-mss *number*;

  The TCP protocol negotiates an MSS value during session connection establishment between two peers. The MSS value negotiated is primarily based on the MTU of the interfaces to which the communicating peers are directly connected to. However in the network, due to variation in link MTU on the path taken by the TCP packets, some packets which are still well within the MSS value may be fragmented when the concerned packet's size exceeds the link's MTU.

  If the router receives a TCP packet with the SYN bit and MSS option set and the MSS option specified in the packet is larger than the MSS value specified by the **tcp-mss** statement, the router replaces the MSS value in the packet with the lower value specified by the **tcp-mss** statement. The range for the **tcp-mss *mss-value*** parameter is from **536** through **65535**.

  To view statistics of SYN packets received and SYN packets whose MSS value, is modified, issue the **show services service-sets statistics tcp-mss** operational mode command. For more information on this topic, see the Junos OS System Basics Configuration Guide.

Related Documentation

- Configuring SNMP Traps for Flow Limits

Copyright © 2013, Juniper Networks, Inc.

Configuring System Logging for Service Sets

You specify properties that control how system log messages are generated for the service set. These values override the values configured at the **[edit interfaces** *interface-name* **services-options]** hierarchy level.

To configure service-set-specific system logging values, include the **syslog** statement at the **[edit services service-set** *service-set-name***]** hierarchy level:

```
syslog {
  host hostname {
    class class-name
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number
    services severity-level;
    source-address source-address
  }
}
```

Configure the **host** statement with a hostname or an IP address that specifies the system log target server. The hostname **local** directs system log messages to the Routing Engine. For external system log servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify only one system logging hostname. The **source-address** parameter is supported on the ms, rms, and mams interfaces.

Table 17 on page 479 lists the severity levels that you can specify in configuration statements at the **[edit services service-set** *service-set-name* **syslog host** *hostname***]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 17: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling

Copyright © 2013, Juniper Networks, Inc.

Table 17: System Log Message Severity Levels *(continued)*

Severity Level	Description
info	Events or non-rror conditions of interest

We recommend setting the system logging severity level to **error** during normal operation. To monitor PIC resource usage, set the level to **warning**. To gather information about an intrusion attack when an intrusion detection system error is detected, set the level to **notice** for a specific service set. To debug a configuration or log NAT functionality, set the level to **info**.

For more information about system log messages, see the *Junos OS System Log Messages Reference*.

To select the class of messages to be logged to the specified system log host, include the **class** statement at the **[edit services service-set *service-set-name* syslog host *hostname*]** hierarchy level:

class *class-name*;

To use one particular facility code for all logging to the specified system log host, include the **facility-override** statement at the **[edit services service-set *service-set-name* syslog host *hostname*]** hierarchy level:

facility-override *facility-name*;

The supported facilities are: **authorization**, **daemon**, **ftp**, **kernel**, **user**, and **local0** through **local7**.

To specify a text prefix for all logging to this system log host, include the **log-prefix** statement at the **[edit services service-set *service-set-name* syslog host *hostname*]** hierarchy level:

log-prefix *prefix-value*;

Enabling Services PICs to Accept Multicast Traffic

To allow multicast traffic to be sent to the Adaptive Services or Multiservices PIC, include the **allow-multicast** statement at the **[edit services service-set *service-set-name*]** hierarchy level. If this statement is not included, multicast traffic is dropped by default. This statement applies only to multicast traffic using a next-hop service set; interface service set configuration is not supported. Only unidirectional flows are created for multicast packets. For a configuration example, see Example: Configuring NAT for Multicast Traffic.

Tracing Services PIC Operations

Tracing operations track all adaptive services operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

Copyright © 2013, Juniper Networks, Inc.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit services adaptive-services-pics]** or **[edit services logging]** hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **serviced** located in the **/var/log** directory.
- When the file **serviced** reaches 128 kilobytes (KB), it is renamed **serviced.0**, then **serviced.1**, and so on, until there are three trace files. Then the oldest trace file (**serviced.2**) is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configures the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements:

```
file filename <files number> <match regular-expression> <size size> <world-readable |
  no-world-readable>;
flag {
  all;
  command-queued;
  config;
  handshake;
  init;
  interfaces;
  mib;
  removed-client;
  show;
}
```

You include these statements at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level.

These statements are described in the following sections:

- Configuring the Adaptive Services Log Filename on page 481
- Configuring the Number and Size of Adaptive Services Log Files on page 482
- Configuring Access to the Log File on page 482
- Configuring a Regular Expression for Lines to Be Logged on page 482
- Configuring the Trace Operations on page 482

Configuring the Adaptive Services Log Filename

By default, the name of the file that records trace output is **serviced**. You can specify a different name by including the **file** statement at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file filename;
```

Copyright © 2013, Juniper Networks, Inc.

Configuring the Number and Size of Adaptive Services Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed ***filename*.0**, then ***filename*.1**, and so on, until there are three trace files. Then the oldest trace file (***filename*.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file <filename> files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (***filename***) reaches 2 MB, ***filename*** is renamed ***filename*.0**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0** is renamed ***filename*.1** and ***filename*** is renamed ***filename*.0**. This process repeats until there are 20 trace files. Then the oldest file (***filename*.19**) is overwritten by the newest file (***filename*.0**).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file <filename> world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

```
file <filename> no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the **match** statement at the **[edit services adaptive-services-pics traceoptions file *filename*]** or **[edit services logging traceoptions]** hierarchy level and specifying a regular expression (regex) to be matched:

```
file <filename> match regular-expression;
```

Configuring the Trace Operations

By default, if the **traceoptions** configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit services adaptive-services-pics traceoptions]** or **[edit services logging traceoptions]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
flag {
  all;
  configuration;
  routing-protocol;
  routing-socket;
  snmp;
}
```

Table 18 on page 483 describes the meaning of the adaptive services tracing flags.

Table 18: Adaptive Services Tracing Flags

Flag	Description	Default Setting
all	Trace all operations.	Off
command-queued	Trace command enqueue events.	Off
config	Log reading of the configuration at the **[edit services]** hierarchy level.	Off
handshake	Trace handshake events.	Off
init	Trace initialization events.	Off
interfaces	Trace interface events.	Off
mib	Trace GGSN SNMP MIB events.	Off
removed-client	Trace client cleanup events.	Off
show	Trace CLI command servicing.	Off

To display the end of the log, issue the **show log serviced | last** operational mode command:

```
[edit]
user@host# run show log serviced | last
```

Example: Configuring Service Sets

Apply two service sets, **my-input-service-set** and **my-output-service-set**, on an interface-wide basis. All traffic has **my-input-service-set** applied to it. After the service set is applied, additional filtering is done using **my_post_service_input_filter**.

```
[edit interfaces fe-0/1/0]
unit 0 {
  family inet {
    service {
      input {
        service-set my-input-service-set;
        post-service-filter my_post_service_input_filter;
      }
      output {
```

Copyright © 2013, Juniper Networks, Inc.

```
            service-set my-output-service-set;
        }
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 25

Summary of Service Set Configuration Statements

The following sections explain each of the service set configuration statements. The statements are organized alphabetically.

adaptive-services-pics

Syntax

```
adaptive-services-pics {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4. The **file** option was added in Release 8.0.

Description Define global services properties.

Options The remaining statement is explained separately.

Usage Guidelines See "Tracing Services PIC Operations" on page 480.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

allow-multicast

Syntax allow-multicast;

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced in Junos OS Release 8.0.

Description Allow multicast traffic to be sent to the Adaptive Services or Multiservices PIC.

Usage Guidelines See "Enabling Services PICs to Accept Multicast Traffic" on page 480.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

anti-replay-window-size (Services Service Set)

Syntax anti-replay-window-size *bits*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the size of the IPsec antireplay window. This statement is useful for dynamic endpoint tunnels for which you cannot configure the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement sets the antireplay window size for all the static tunnels within this service set. If a particular tunnel needs a specific value for antireplay window size, set the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level. If antireplay check has to be disabled for a particular tunnel in this service set, set the **no-anti-replay** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: The **anti-replay-window-size** and **no-anti-replay** settings at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level override the settings specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

Options ***bits***—Size of the antireplay window, in bits.

Default: 64 bits (AS PICs), 128 bits (Multiservices PICs and DPCs)

Range: 64 through 4096 bits

Usage Guidelines See "Configuring IPsec Service Sets" on page 473 or "Configuring or Disabling IPsec Anti-Replay" on page 248.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

bypass-traffic-on-exceeding-flow-limits

Syntax bypass-traffic-on-exceeding-flow-limits;

Hierarchy Level [editservices service-set *service-set-name* service-set-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable packets to bypass without creating a new session when the flow in the service set exceeds the limit that is set by the **max-flows** statement at the **[edit services service-set *service-set-name*]** hierarchy level.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

bypass-traffic-on-pic-failure

Syntax bypass-traffic-on-pic-failure;

Hierarchy Level [edit services service-set *service-set-name* service-set-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description When the MultiServices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the **bypass-traffic-on-pic-failure** statement. When this statement is configured, the affected packets are forwarded in the event of a MultiServices PIC failure or offlining, as though interface-style services were not configured.

This issue applies only to Dynamic Application Awareness for Junos OS configurations with IDP service sets.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

clear-dont-fragment-bit (Services Service Set)

Syntax clear-dont-fragment-bit;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the IPsec tunnel. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. This statement is useful for dynamic endpoint tunnels, for which you cannot configure the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, setting this statement clears the DF bit on packets entering all the static tunnels within this service set. If you want to clear the DF bit on packets entering a specific tunnel, set the **clear-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

Usage Guidelines See "Configuring IPsec Service Sets" on page 473 or "Configuring Actions in IPsec Rules" on page 246.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

facility-override

Syntax facility-override *facility-name*;

Hierarchy Level [edit services service-set *service-set-name* syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Override the default facility for system log reporting.

Options ***facility-name***—Name of the facility that overrides the default assignment. Valid entries are:

- **authorization**
- **daemon**
- **ftp**
- **kernel**
- **local0** through **local7**
- **user**

Usage Guidelines See "Configuring System Logging for Service Sets" on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

host (service-set)

Syntax

```
host hostname {
    classclass-name
    facility-override facility-name;
    interface-service prefix-value;
    log-prefixprefix-value
    port port-number
    services severity-level;
    source-addresssource-address
}
```

Hierarchy Level [edit services service-set *service-set-name* syslog]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the hostname for the system logging utility.

Options ***hostname***—Name of the system logging utility host machine.

The remaining statements are explained separately.

Usage Guidelines See "Configuring System Logging for Service Sets" on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

ids-rules

Syntax (ids-rules *rule-name* | ids-rule-sets *rule-set-name*);

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the intrusion detection service (IDS) rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Usage Guidelines See "Configuring Service Rules" on page 472.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ike-access-profile

Syntax ike-access-profile *profile-name*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description Define the access profile for the IPsec traffic on dynamic tunnels.

Options ***profile-name***—Identifier for access profile, which must match the name configured at the **[edit access profile *name* client * ike]** hierarchy level.

Usage Guidelines See "Configuring Dynamic Endpoints for IPsec Tunnels" on page 250 or "Configuring IPsec Service Sets" on page 473.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

interface-service

Syntax
```
interface-service {
  service-interface name;
}
```

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the device name for the interface service Physical Interface Card (PIC).

Options **service-interface *name***—Name of the service device associated with the interface-wide service set.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ip-reassembly-rules (Service Set)

Syntax ip-reassembly-rules *rule-name*;

Hierarchy Level [edit services service-set*service-set-name*]

Release Information Statement introduced in Junos OS Release 13.1

Description Specify one or more previously configured IP reassembly rules to associate with the service set.



NOTE: **The IP reassembly rule must be defined at the [edit services ip-reassembly rule] hierarchy level.**

Options ***rule-name***—Name of an IP reassembly rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 321
- IP Packet Fragment Reassembly for L2TP Overview on page 320

Copyright © 2013, Juniper Networks, Inc.

ipsec-vpn-options

Syntax

```
ipsec-vpn-options {
    anti-replay-window-size bits;
    clear-dont-fragment-bit;
    ike-access-profile profile-name;
    local-gateway address;
    no-anti-replay;
    passive-mode-tunneling;
    trusted-ca [ ca-profile-names ];
    tunnel-mtu bytes;
}
```

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify IP Security (IPsec) service options.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Service Rules" on page 472.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

ipsec-vpn-rules

Syntax (ipsec-vpn-rules *rule-name* | ipsec-vpn-rule-sets *rule-set-name*);

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the IPsec rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Usage Guidelines See "Configuring Service Rules" on page 472.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

local-gateway

Syntax local-gateway *address*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the local IPv4 or IPv6 address for the IPsec traffic.

Options ***address***—Local address.

Usage Guidelines See "Configuring Service Rules" on page 472.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

log-prefix (Services)

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit services service-set *service-set-name* syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the system logging prefix value.

Options ***prefix-value***—System logging prefix value.

Usage Guidelines See "Configuring System Logging for Service Sets" on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

logging (Services)

Syntax

```
logging {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 8.0.

Description Define global services properties.

Options The remaining statement is explained separately.

Usage Guidelines See "Tracing Services PIC Operations" on page 480.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

max-flows

Syntax max-flows *number*;

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Maximum number of flows allowed for the service set.

Options ***number***—Maximum number of flows.

Usage Guidelines See "Configuring Service Set Limitations" on page 478.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

message-rate-limit

Syntax message-rate-limit *messages-per-second*

Hierarchy Level

```
interfaces interface-name {
  services-options {
    cgn-pic;
    disable-global-timeout-override;
    ignore-errors <alg> <tcp>;
    inactivity-non-tcp-timeout seconds;
    inactivity-tcp-timeout seconds;
    inactivity-timeout seconds;
    open-timeout seconds;
    session-limit {
      maximum number;
      rate new-sessions-per-second;
    }
    session-timeout seconds;
    syslog {
    }
  }
}
```

Release Information Statement introduced Junos OS Release 11.1.

Description Maximum system log messages per second allowed from this interface.



NOTE: The message-rate-limit command can be configured only for physical service interfaces (sp-*x/x/x*) and not for redundancy services PIC interfaces (rsp*x*).

Options *messages-per-second*——This option configures the maximum number of system log messages per second that can be formatted and sent from the PIC to either the Routing Engine (local) or to an external server (remote). The default rates are 10,000 for the Routing Engine and 200,000 for an external server.

Usage Guidelines See "Configuring System Logging for Service Sets" on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

nat-options

Syntax

```
nat-options {
  stateful-nat64 {
    clear-dont-fragment-bit
  }
}
```

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced with Junos OS Release 12.1.

Description Specity that the DF (don't fragment) bit in an translated IPv4 packet is cleared if its packet size is less than 1280 bytes.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

nat-rules

Syntax (nat-rules *rule-name* | nat-rule-sets *rule-set-name*);

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the Network Address Translation (NAT) rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Usage Guidelines See "Configuring Service Rules" on page 472.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

next-hop-service

Syntax

```
next-hop-service {
  inside-service-interface interface-name.unit-number;
  outside-service-interface interface-name.unit-number;
  outside-service-interface-type local;
  service-interface-pool name;
}
```

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.
service-interface-pool option added in Junos OS Release 9.3.

Description Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface.

Options

inside-service-interface ***interface-name.unit-number***—Name and logical unit number of the service interface associated with the service set applied inside the network.

outside-service-interface ***interface-name.unit-number***—Name and logical unit number of the service interface associated with the service set applied outside the network.

outside-service-interface-type ***interface-type***—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.

service-interface-pool ***name***—Name of the pool of logical interfaces configured at the **[edit services** service-interface-pools **pool** ***pool-name*****]** hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.



NOTE: **service-interface-pool** is not applicable for IP reassembly configuration on L2TP.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

no-anti-replay (Services Service Set)

Syntax no-anti-replay;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Disable IPsec antireplay service for this service set, which occasionally causes interoperability issues for security associations. This statement is useful for dynamic endpoint tunnels for which you cannot configure the **no-anti-reply** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement disables the antireplay check for all the tunnels within this service set. If antireplay check has to be enabled for a particular tunnel, then set the **anti-replay-window-size** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: Setting the **anti-replay-window-size** and **no-anti-replay** statements at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level overrides the settings specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

Usage Guidelines See "Configuring IPsec Service Sets" on page 473 or "Configuring or Disabling IPsec Anti-Replay" on page 248.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

passive-mode-tunneling

Syntax passive-mode-tunneling;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Allows tunneling of malformed packets. When this statement is enabled, traffic bypasses the usual active IP checks. The IPsec tunnel is not treated as a next hop and TTL is not decremented. If the packet size exceeds the tunnel MTU value, an ICMP error is not generated.

Usage Guidelines See "Configuring IPsec Service Sets" on page 473.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

pgcp-rules

Syntax (pgcp-rules *rule-name* | pgcp-rules-sets *rule-set-name*);

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Specify the Packet Gateway Control Protocol (PGCP) rules or rule set included in this service set. You can configure multiple rules but only one rule set for each service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface–control—To add this statement to the configuration.

port (syslog)

Syntax port *port-number*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced in Junos OS Release 11.1.

Description UDP port for system log messages on the host. The default port is 514.

Options ***port-number***—Port number for system log messages.

Usage Guidelines See "Configuring System Logging for Services Interfaces" on page 518.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ptsp-rules

Syntax (ptsp-rules *rule-name* | ptsp-rules-sets *rule-set-name*);

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the PTSP rules or rule set included in this service set. You can configure multiple rules but only one rule set for each service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

rule-set-name—Identifier for the set of rules to be included.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface–control—To add this statement to the configuration.

service-interface

Syntax service-interface *interface-name*;

Hierarchy Level [edit services service-set *service-set-name* interface-service]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the name for the adaptive services interface associated with an interface-wide service set.

Options ***interface-name***—Identifier of the service interface.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

service-set (Services)

Syntax

```
service-set service-set-name {
  allow-multicast;
  extension-service service-name {
    provider-specific-rules-configuration;
  }
  (ids-rules rule-name | ids-rule-sets rule-set-name);
  interface-service {
    service-interface interface-name;
  }
  ipsec-vpn-options {
    anti-replay-window-size bits;
    clear-dont-fragment-bit;
    ike-access-profile profile-name;
    local-gateway address;
    no-anti-replay;
    passive-mode-tunneling;
    trusted-ca [ ca-profile-names ];
    tunnel-mtu bytes;
  }
  ip-reassembly-rules rule-name};
  (ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
  max-flows number;
  max-drop-flows {
    ingress ingress-flows;
    egress egress-flows;
  }
  nat-options {
    stateful-nat64 {
      clear-dont-fragment-bit
    }
  }
  (nat-rules rule-name | nat-rule-sets rule-set-name);
  next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
    outside-service-interface-type local;
    service-interface-pool name;
  }
  (pgcp-rules rule-name | pgcp-rule-sets rule-set-name);
  (ptsp-rules rule-name | ptsp-rule-sets rule-set-name);
  snmp-trap-thresholds{
    flows high high-threshold | low low-threshold;
    nat-address-port high-threshold | low low-threshold;
    }
  }
  (softwire-rules rule-name | softwire-rule-sets rule-set-name);
  (stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);
  syslog {
    host hostname {
      class class-name;
      facility-override facility-name;
      log-prefix prefix-value;
```

Copyright © 2013, Juniper Networks, Inc.

```
        port port-number;
        services severity-level;
        source-address source-address
      }
    }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4. The **pgcp-rules** and **pgcp-rule-sets** options were added in Release 8.4. The **ptsp-rules** and **ptsp-rule-sets** options were added in Release 10.2. The **softwire-rules** and **softwire-rule-sets** options were added in Release 10.4.

Description Define the service set.

Options ***service-set-name***—Identifies the service set.

The remaining statements are explained separately.

Usage Guidelines See Service Set Properties.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

services

See the following sections:

- services (Hierarchy) on page 505
- services (System Logging) on page 506

services (Hierarchy)

Syntax services { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service rules to be applied to traffic.

Usage Guidelines See Service Set Properties.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

services (System Logging)

Syntax services *severity-level*;

Hierarchy Level [edit services service-set *service-set-name* syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the severity level for system logging messages.

Options ***severity-level***—Assigns a severity level to the facility. Valid entries are:

- **alert**—Conditions that should be corrected immediately.
- **any**—Matches any level.
- **critical**—Critical conditions.
- **emergency**—Panic conditions.
- **error**—Error conditions.
- **info**—Informational messages.
- **notice**—Conditions that require special handling.
- **warning**—Warning messages.

Usage Guidelines See "Configuring System Logging for Service Sets" on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

stateful-firewall-rules

Syntax (stateful-firewall-rules *rule-names* | stateful-firewall-rule-sets *rule-set-name*);

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the stateful firewall rules or rule set included in this service set. You can configure multiple rules, but only one rule set for each service.

Options ***rule-name***—Identifier for the collection of terms that make up this rule.

rule-set-name—Identifier for the set of rules to be included.

Usage Guidelines See "Configuring Service Rules" on page 472.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

syslog (Services Service Set)

Syntax

```
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    interface-service prefix-value;
  }
}
```

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure generation of system log messages for the service set. The system log information is passed to the kernel for logging in the **/var/log** directory. These settings override the values defined at the **[edit interfaces *interface-name* services-options]** hierarchy level; for more information on configuring those values, see "Configuring System Logging for Services Interfaces" on page 518.

Options The remaining statements are described separately.

Usage Guidelines See "Configuring System Logging for Service Sets" on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

tcp-mss

Syntax tcp-mss *number*;

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the TCP Maximum Segment Size (MSS) allowed for the service set.

Options ***number***—MSS value.

Usage Guidelines See "Configuring Service Set Limitations" on page 478.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

traceoptions (Services Logging)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit services adaptive-services-pics],
[edit services logging]

Release Information Statement introduced before Junos OS Release 7.4.
file option added in Release 8.0.

Description Configure Adaptive Services or Multiservices PIC tracing operations. The messages are output to **/var/log/serviced**.

Options

file *filename*—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 3 files

flag *flag*—Tracing operation to perform:

- **all**—Trace everything.
- **command-queued**—Trace command enqueue events.
- **config**—Trace configuration events.
- **handshake**—Trace handshake events.
- **init**—Trace initialization events.
- **interfaces**—Trace interface events.
- **mib**—Trace GGSN SNMP MIB events.
- **removed-client**—Trace client cleanup events.
- **show**—Trace CLI command servicing.

match *regex*—(Optional) Match output to a defined regular expression (regex).

Copyright © 2013, Juniper Networks, Inc.

Default: If you do not include this option, the trace operation output includes all lines relevant to the logged events.

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Usage Guidelines See "Tracing Services PIC Operations" on page 480.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

trusted-ca

Syntax trusted-ca *ca-profile-name*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 7.5.

Description Identify one or more trusted IPsec certification authorities.

Options ***ca-profile-name***—Name of certification authority profile, which is configured at the **[edit security pki]** hierarchy level.

Usage Guidelines See "Configuring IPsec Service Sets" on page 473.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

tunnel-mtu (Services Service Set)

Syntax tunnel-mtu *bytes*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Maximum transmission unit (MTU) size for IPsec tunnels. This statement is useful for dynamic endpoint tunnels for which you cannot configure the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.

For static IPsec tunnels, this statement sets the tunnel MTU value for all the tunnels within this service set. If you need a specific value for a particular tunnel, then set the **tunnel-mtu** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level.



NOTE: The **tunnel-mtu** setting at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level overrides the value specified at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

Options ***bytes***—MTU size.

Default: 1500 bytes

Range: 256 through 9192 bytes

Usage Guidelines See "Configuring IPsec Service Sets" on page 473 or "Specifying the MTU for IPsec Tunnels" on page 249.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- mtu on page 1205

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 26

Service Interface Configuration Guidelines

For the interfaces on a router to function, you must configure them, specifying properties such as the interface location (that is, which slot the Flexible PIC Concentrator [FPC] is installed in and which location on the FPC the Physical Interface Card [PIC] is installed in), the interface type (such as SONET/SDH or Asynchronous Transfer Mode [ATM]), encapsulation, and interface-specific properties. You can configure the interfaces that are currently present in the router, and you can also configure interfaces that are not currently present but that you might add in the future. When a configured interface appears, the Junos OS detects its presence and applies the appropriate configuration to it. For more information on the general configuration of interfaces, see the Junos® OS Network Interfaces.

You can configure two different sets of properties at the interface level:

- Properties that apply to an entire Adaptive Services (AS) or Multiservices PIC interface on a global level, including default values for system logging and timeout properties.
- Assignment of service sets and filters to a network interface.

To configure default properties for the adaptive services interface, include the **sp-*fpc/pic/port*** or **rsp*number*** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
(sp-fpc/pic/port | rspnumber) {
  services-options {
    cgn-pic;
    disable-global-timeout-override;
    ignore-errors <alg> <tcp>;
    inactivity-non-tcp-timeout seconds;
    inactivity-tcp-timeout seconds;
    inactivity-timeout seconds;
    open-timeout seconds;
    session-limit {
      maximum number;
      rate new-sessions-per-second;
    }
    session-timeout seconds;
    syslog {
      host hostname {
        facility-override facility-name;
        log-prefix prefix-value;
        port port-number;
```

Copyright © 2013, Juniper Networks, Inc.

```
            services severity-level;
        }
        message-rate-limit messages-per-second;
    }
  }
}
```

To apply services on network interfaces, include the **unit** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
unit logical-unit-number {
  clear-dont-fragment-bit;
  encapsulation type;
  family (Interfaces) inet {
    address (Interfaces) address {
      ...
    }
    mtu bytes;
    service {
      input (Interfaces) {
        [ service-set (Interfaces) service-set-name <service-filter filter-name> ];
        post-service-filter filter-name;
      }
      output {
        [ service-set (Interfaces) service-set-name <service-filter filter-name> ];
      }
    }
    service-domain (inside | outside);
  }
}
```

To configure AS or Multiservices PIC redundancy, include the **redundancy-options** statement at the **[edit interfaces rsp *number*]** hierarchy level:

```
rspnumber {
  redundancy-options {
    primary sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
  }
}
```

To configure an MX-DPC interface to be used exclusively for carrier-grade NAT (CGN) and related services (intrusion detection, stateful firewall, and softwire),include the **cgn-pic** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level.

This chapter contains the following sections:

- Services Interface Naming Overview on page 515
- Configuring the Address and Domain for Services Interfaces on page 516
- Configuring Default Timeout Settings for Services Interfaces on page 516
- Configuring System Logging for Services Interfaces on page 518
- Enabling Fragmentation on GRE Tunnels on page 519
- Applying Filters and Services to Interfaces on page 520

Copyright © 2013, Juniper Networks, Inc.

- Configuring AS or Multiservices PIC Redundancy on page 522
- Flow Offloading on page 525
- Examples: Configuring Services Interfaces on page 525

Services Interface Naming Overview

Each interface has an interface name, which specifies the media type, the slot the FPC is located in, the location on the FPC that the PIC is installed in, and the PIC port. The interface name uniquely identifies an individual network connector in the system. You use the interface name when configuring interfaces and when enabling various functions and properties, such as routing protocols, on individual interfaces. The system uses the interface name when displaying information about the interface, for example, in the **show interfaces** command.

The interface name is represented by a physical part, a logical part, and a channel part in the following format:

```
physical<:channel>.logical
```

The channel part of the name is optional for all interfaces except channelized DS3, E1, OC12, and STM1 interfaces.

The physical part of an interface name identifies the physical device, which corresponds to a single physical network connector. This part of the interface name has the following format:

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

type is the media type, which identifies the network device. For service interfaces, it can be one of the following:

- **cp**—Flow collector interface.
- **es**—Encryption interface.
- **gr**—Generic routing encapsulation tunnel interface.
- **gre**—This interface is internally generated and not configurable.
- **ip**—IP-over-IP encapsulation tunnel interface.
- **ipip**—This interface is internally generated and not configurable.
- **ls**—Link services interface.
- **lsq**—Link services intelligent queuing (IQ) interface; also used for voice services.
- **ml**—Multilink interface.
- **mo**—Monitoring services interface. The logical interface **mo-*fpc/pic/port***.16383 is an internally generated, nonconfigurable interface for router control traffic.
- **mt**—Multicast tunnel interface. This interface is automatically generated, but you can configure properties on it if needed.
- **mtun**—This interface is internally generated and not configurable.

Copyright © 2013, Juniper Networks, Inc.

- **rlsq**—Redundancy LSQ interface.
- **rsp**—Redundancy adaptive services interface.
- **si**—Services inline interface, configured on MX3D Series routers only.
- **sp**—Adaptive services interface. The logical interface **sp-*fpc/pic/port***.16383 is an internally generated, nonconfigurable interface for router control traffic.
- **tap**—This interface is internally generated and not configurable.
- **vp**—Voice over IP (VoIP) interface, configured on J Series Services Routers only.
- **vt**—Virtual loopback tunnel interface.

Configuring the Address and Domain for Services Interfaces

On the AS or Multiservices PIC, you configure a source address for system log messages by including the **address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
address address {
    ...
}
```

Assign an IP address to the interface by configuring the ***address*** value. The AS or Multiservices PIC generally supports only IP version 4 (IPv4) addresses configured using the **family inet** statement, but IPsec services support IP version 6 (IPv6) addresses as well, configured using the **family inet6** statement.

For information on other addressing properties you can configure that are not specific to service interfaces, see the Junos® OS Network Interfaces.

The **service-domain** statement specifies whether the interface is used within the network or to communicate with remote devices. The software uses this setting to determine which default stateful firewall rules to apply, and to determine the default direction for service rules. To configure the domain, include the **service-domain** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
service-domain (inside | outside);
```

If you are configuring the interface in a next-hop service-set definition, the **service-domain** setting must match the configuration for the **inside-service-interface** and **outside-service-interface** statements; for more information, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

Configuring Default Timeout Settings for Services Interfaces

You can specify global default settings for certain timers that apply for the entire interface. There are two statements of this type:

- **inactivity-timeout**—Sets the inactivity timeout period for established flows, after which they are no longer valid.

Copyright © 2013, Juniper Networks, Inc.

- **open-timeout**—Sets the timeout period for Transmission Control Protocol (TCP) session establishment, for use with SYN-cookie defenses against network intrusion.

To configure a setting for the inactivity timeout period, include the **inactivity-timeout** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
inactivity-timeout seconds;
```

The default value is 30 seconds. The range of possible values is from 4 through 86,400 seconds. Any value you configure in the application protocol definition overrides the value specified here; for more information, see "Configuring Application Protocol Properties" on page 74.

To configure a setting for the TCP session establishment timeout period, include the **open-timeout** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
open-timeout seconds;
```

The default value is 30 seconds. The range of possible values is from 4 through 224 seconds. Any value you configure in the intrusion detection service (IDS) definition overrides the value specified here; for more information, see Intrusion Detection Properties.

Use of Keep-Alive Messages for Greater Control of TCP Inactivity Timeouts

Keep-alive messages are generated automatically to prevent TCP inactivity timeouts. The default number of keep-alive messages is 4. However, you can configure the number of keep-alive messages by entering the **tcp-tickles** statement at the **[edit interaces *interface-name* service-options]** hierarchy level.

When timeout is generated for a bidirectional TCP flow, keep-alive packets are sent to reset the timer. If number of consecutive keep-alive packets sent in a flow reaches the default or configured limit, the conversation is deleted. There are several possible scenarios, depending on the setting of the **inactivity-timer** and the default or configured maximum number of keep-alive messages.

- If the configured value of keep-alive messages is zero and **inactivity-timeout** is NOT configured (in which case the default timeout value of 30 is used), no keep-alive packets are sent. The conversation is deleted when any flow in the conversation is idle for more than 30 seconds.
- If the configured value of keep-alive messages is zero and the **inactivity-timeout** is configured, no keep-alive packets are sent, and the conversation is deleted when any flow in the conversation is idle for more than the configured timeout value.
- If the default or configured maximum number of keep-alive messages is some positive integer, and any of the flows in a conversation is idle for more than the default or configured value for **inactivity-timeout** keep-alive packets are sent. If hosts do not respond to the configured number of consecutive keep-alive packets, the conversation is deleted. The interval between keep-alive packets will be 1 second. However, if the

Copyright © 2013, Juniper Networks, Inc.

host sends back an ACK packet, the corresponding flow becomes active, and keep-alive packets are not sent until the flow becomes idle again.

Configuring System Logging for Services Interfaces

You specify properties that control how system log messages are generated for the interface as a whole. If you configure different values for the same properties at the **[edit services service-set *service-set-name*]** hierarchy level, the service-set values override the values configured for the interface. For more information on configuring service-set properties, see "Configuring System Logging for Service Sets" on page 479.

To configure interface-wide default system logging values, include the **syslog** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level:

```
[edit interfaces interface-name services-options]
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number;
  }
}
```

Configure the **host** statement with a hostname or an IP address that specifies the system log target server. The hostname **local** directs system log messages to the Routing Engine. For external system log servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify only one system logging hostname.

Table 19 on page 518 lists the severity levels that you can specify in configuration statements at the **[edit interfaces *interface-name* services-options syslog host *hostname*]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 19: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the router to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring

Copyright © 2013, Juniper Networks, Inc.

Table 19: System Log Message Severity Levels *(continued)*

Severity Level	Description
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

We recommend setting the system logging severity level to **error** during normal operation. To monitor PIC resource usage, set the level to **warning**. To gather information about an intrusion attack when an intrusion detection system error is detected, set the level to **notice** for a specific interface. To debug a configuration or log Network Address Translation (NAT) functionality, set the level to **info**.

For more information about system log messages, see the *Junos OS System Log Messages Reference*.

To use one particular facility code for all logging to the specified system log host, include the **facility-override** statement at the **[edit interfaces *interface-name* services-options syslog host *hostname*]** hierarchy level:

```
[edit interfaces interface-name services-options]
facility-override facility-name;
```

The supported facilities include **authorization**, **daemon**, **ftp**, **kernel**, **user**, and **local0** through **local7**.

To specify a text prefix for all logging to this system log host, include the **log-prefix** statement at the **[edit interfaces *interface-name* services-options syslog host *hostname*]** hierarchy level:

```
[edit interfaces interface-name services-options]
log-prefix prefix-value;
```

Enabling Fragmentation on GRE Tunnels

To enable fragmentation of IPv4 packets in generic routing encapsulation (GRE) tunnels, include the **clear-dont-fragment-bit** statement and a maximum transmission unit (MTU) setting for the tunnel as part of an existing GRE configuration at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
gr-fpc/pic/port {
    unit logical-unit-number {
        clear-dont-fragment-bit;
        ...
        family inet {
            mtu 1000;
            ...
        }
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

This statement clears the Don't Fragment (DF) bit in the packet header, regardless of the packet size. If the packet size exceeds the tunnel MTU value, the packet is fragmented before encapsulation. The maximum MTU size configurable on the AS or Multiservices PIC is 9192 bytes.



NOTE: The **clear-dont-fragment-bit** statement is supported only on MX Series routers and all M Series routers except the M320 router.

Fragmentation is enabled only on IPv4 packets being encapsulated in IPv4-based GRE tunnels.



NOTE: This configuration is supported only on GRE tunnels on AS or Multiservices interfaces. If you commit **gre-fragmentation** as the encapsulation type on a standard Tunnel PIC interface, the following console log message appears when the PIC comes online:

gr-*fpc/pic/port*: does not support this encapsulation

The Packet Forwarding Engine updates the IP identification field in the outer IP header of GRE-encapsulated packets, so that reassembly of the packets is possible after fragmentation. The previous CLI constraint check that required you to configure either the **clear-dont-fragment-bit** statement or a tunnel key with the **allow-fragmentation** statement is no longer enforced.

Applying Filters and Services to Interfaces

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces on the router. To associate a defined service set with an interface, include the **service-set** statement with the **input** or **output** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet service]
input {
   service-set service-set-name <service-filter filter-name>;
   post-service-filter filter-name;
}
output {
   service-set service-set-name <service-filter filter-name>;
}
```



NOTE: When you enable services on an interface, reverse-path forwarding is not supported. You cannot configure services on the management interface (**fxp0**) or the loopback interface (**lo0**).

You can configure different service sets on the input and output sides of the interface. However, for service sets with bidirectional service rules, you must include the same

Copyright © 2013, Juniper Networks, Inc.

service set definition in both the **input** and **output** statements. Any service set you include in the **service** statement must be configured with the **interface-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level; for more information, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.



NOTE: If you configure an interface with an input firewall filter that includes a reject action and with a service set that includes stateful firewall rules, the router executes the input firewall filter before the stateful firewall rules are run on the packet. As a result, when the Packet Forwarding Engine sends an Internet Control Message Protocol (ICMP) error message out through the interface, the stateful firewall rules might drop the packet because it was not seen in the input direction.

Possible workarounds are to include a forwarding-table filter to perform the reject action, because this type of filter is executed after the stateful firewall in the input direction, or to include an output service filter to prevent the locally generated ICMP packets from going to the stateful firewall service.

Configuring Service Filters

You can optionally include filters associated with each service set to refine the target and additionally process the traffic. If you include the **service-set** statement without a **service-filter** definition, the router software assumes that the match condition is true and selects the service set for processing automatically.

To configure service filters, include the **firewall** statement at the **[edit]** hierarchy level:

```
firewall {
    family inet {
        service-filter filter-name {
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
        }
    }
}
```



NOTE: You must specify inet as the address family to configure a service filter.

You configure service filters in a similar way to firewall filters. Service filters have the same match conditions as firewall filters, but the following specific actions:

Copyright © 2013, Juniper Networks, Inc.

- **count**—Add the packet to a counter total.
- **log**—Log the packet.
- **port-mirror**—Port-mirror the packet.
- **sample**—Sample the packet.
- **service**—Forward the packet for service processing.
- **skip**—Omit the packet from service processing.

For more information about configuring firewall filters, see the Routing Policy Configuration Guide.

You can also include more than one service set definition on each side of the interface. If you include multiple service sets, the router software evaluates them in the order specified in the configuration. It executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions.

An additional statement allows you to specify a filter for processing the traffic after the input service set is executed. To configure this type of filter, include the **post-service-filter** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service input]** hierarchy level:

```
post-service-filter filter-name;
```



NOTE: The software performs postservice filtering only when it has selected and executed a service set. If the traffic does not meet the match criteria for any of the configured service sets, the postservice filter is ignored.

For an example of applying a service set to an interface, see "Examples: Configuring Services Interfaces" on page 525.

For more information on applying filters to interfaces, see the Junos® OS Network Interfaces. For general information on filters, see the Routing Policy Configuration Guide.



NOTE: After NAT processing is applied to packets, they are not subject to output service filters. The service filters affect only untranslated traffic.

Configuring AS or Multiservices PIC Redundancy

You can configure AS or Multiservices PIC redundancy on M Series and T Series routers, except TX Matrix routers, that have multiple AS or Multiservices PICs. To configure redundancy, you specify a redundancy services PIC (**rsp**) interface in which the primary PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary AS or Multiservices PIC is restored, it remains on standby and does not preempt the secondary PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the **show interfaces redundancy** command.

Copyright © 2013, Juniper Networks, Inc.

Failover to the secondary PIC occurs under the following conditions:

- The primary PIC, FPC, or Packet Forwarding Engine goes down, resets, or is physically removed from the router.
- The PIC or FPC is taken offline using the **request chassis pic fpc-slot** *slot-number* **pic-slot** *slot-number* **offline** or **request chassis fpc slot** *slot-number* **offline** command. For more information, see the Junos OS Operational Mode Commands.
- The driver watchdog timer expires.
- The **request interface switchover** command is issued. For more information, see the Junos OS Operational Mode Commands.



NOTE: Adaptive Services and Multiservices PICs in Layer-2 mode (running Layer 2 services) are not rebooted when a MAC flow-control situation is detected.

The physical interface type **rsp** specifies the pairings between primary and secondary **sp** interfaces to enable redundancy. To configure an AS or Multiservices PIC as the backup, include the **redundancy-options** statement at the **[edit interfaces rsp***number***]** hierarchy level:

```
[edit interfaces rspnumber]
redundancy-options {
    primary sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
}
```



NOTE: You can include a similar redundancy configuration for Link Services IQ (LSQ) PICs at the **[edit interfaces rlsq***number***]** hierarchy level. For more information, see "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 352.

Copyright © 2013, Juniper Networks, Inc.

The following constraints apply to redundant AS or Multiservices PIC configurations:

- The services supported in redundancy configurations include stateful firewall, NAT, IDS, and IPsec. Services mounted on the AS or Multiservices PIC that use interface types other than **sp-** interfaces, such as tunneling and voice services, are not supported. For information on flow monitoring redundancy, see "Configuring Services Interface Redundancy with Flow Monitoring" on page 975.

  

  NOTE: For IPsec functionality, the router no longer needs to renegotiate security associations (SAs) during warm standby PIC switchover. Instead, the warm standby feature has been made stateful by periodically setting a checkpoint between the working state of the PIC and the Routing Engine, which should lessen the downtime during switchover. If you prefer to retain the earlier behavior, you can include the **clear-ipsec-sas-on-pic-restart** statement at the **[edit services ipsec-vpn]** hierarchy level. If you enable this capability, the router renegotiates the IPsec SAs on warm standby PIC switchover. For more information, see "Clearing Security Associations" on page 228.

- We recommend that you pair the same model type in RSP configurations, such as two ASMs or two AS2 PICs. If you pair unlike models, the two PICs may perform differently.
- You can specify an AS or Multiservices PIC (**sp** interface) as the primary for only one **rsp** interface.
- An **sp** interface can be a secondary for multiple **rsp** interfaces. However, the same **sp** interface cannot be configured as a primary interface in one **rsp** configuration and as a secondary in another configuration.
- When the secondary PIC is active, if another primary PIC that is paired with it in an **rsp** configuration fails, no failover takes place.
- When you configure an AS or Multiservices PIC within a redundant configuration, the **sp** interface cannot have any configured services. Apply the configurations at the **[edit interfaces rsp*number*]** hierarchy level, using, for example, the **unit** and **services-options** statements. Exceptions include the **multiservice-options** statement used in flow monitoring configurations, which can be configured separately for the primary and secondary **sp** interfaces, and the **traceoptions** statement.
- All the operational mode commands that apply to **sp** interfaces also apply to **rsp** interfaces. You can issue **show** commands for the **rsp** interface or the primary and secondary **sp** interfaces.
- If a secondary PIC fails while it is in use, the **rsp** interface returns to the " not present" state. If the primary PIC comes up later, service is restored to it.

For a sample configuration, see "Examples: Configuring Services Interfaces" on page 525.

Copyright © 2013, Juniper Networks, Inc.

Flow Offloading

The Junos OS enables you to configure flow offloading for PICS on MX Series routers using Modular Port Concentrator (MPCs) with Modular Interface Cards (MICs). Flows are offloaded to Fast Update Filters (FUFs) on the Packet Forwarding Engine. Offloading produces the greatest benefits when applied to long-lasting or high-bandwidth flows

The maximum number of active offloads is 200,000 per PIC. When offloaded flows are deleted, more flows can be offloaded.

To configure flow offloading:

- At the **[edit interfaces *interface-name* services-options]** hierarchy level, enter the **trio-flow-offload minimum-bytes *minimum-bytes*** statement.

  ```
  user@host# edit services interface-name
  [edit services interface-name services-options]
  user@host# set trio-flow-offload minimum-bytes minimum- bytes
  ```

  In the following example, flows are offloaded when they consist of no less than 1024 bytes:

  ```
  user@host# edit services ms-0/1/0
  [edit services ms-0/1/0 services-options]
  user@host# set trio-flow-offload minimum-bytes 1024
  ```

Related Documentation

- trio-flow-offload on page 545

Examples: Configuring Services Interfaces

Apply the **my-service-set** service set on an interface-wide basis. All traffic that is accepted by **my_input_filter** has **my-input-service-set** applied to it. After the service set is applied, additional filtering is done using the **my_post_service_input_filter** filter.

```
[edit interfaces fe-0/1/0]
unit 0 {
  family inet {
    filter {
      input my_input_filter;
      output my_output_filter;
    }
    service {
      input {
        service-set my-input-service-set;
        post-service-filter my_post_service_input_filter;
      }
      output {
        service-set my-output-service-set;
      }
    }
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

Configure two redundancy interfaces, **rsp0** and **rsp1**, and associated services.

```
[edit interfaces]
rsp0 {
   redundancy-options {
      primary sp-0/0/0;
      secondary sp-1/3/0;
   }
   unit 0 {
      family inet;
   }
   unit 30 {
      family inet;
      service-domain inside;
   }
   unit 31 {
      family inet;
      service-domain outside;
   }
}
rsp1 {
   redundancy-options {
      primary sp-0/1/0;
      secondary sp-1/3/0;
   }
   unit 0 {
      family inet;
   }
   unit 20 {
      family inet;
      service-domain inside;
   }
   unit 21 {
      family inet;
      service-domain outside;
   }
}
[edit services]
service-set null-sfw-with-nat {
   stateful-firewall-rules allow-all;
   nat-rules rule1;
   next-hop-service {
      inside-service-interface rsp0.30;
      outside-service-interface rsp0.31;
   }
}
[edit routing-instances]
vpna {
   interface rsp0.0;
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 27

Summary of Service Interface Configuration Statements

The following sections explain each of the service interface configuration statements. The statements are organized alphabetically.

address (Interfaces)

Syntax

address *address* {
  ...
}

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family (Interfaces) *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (Interfaces) *family*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the interface address.

Options

address—Address of the interface.

Usage Guidelines

See "Configuring the Address and Domain for Services Interfaces" on page 516; for a general discussion of **address** statement options, see the Junos® OS Network Interfaces.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

cgn-pic

Syntax cgn-pic;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 11.2.

Description Restrict usage of the service PIC to carrier-grade NAT (CGN) or associated services (intrusion detection, stateful firewall, and softwire). All memory is available for CGN or related services and can be used for CGN scaling.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring NAT Service Sets

clear-dont-fragment-bit

Syntax clear-dont-fragment-bit;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Clear the Don't Fragment (DF) bit on all IP version 4 (IPv4) packets entering the generic routing encapsulation (GRE) tunnel on Adaptive Services (AS) or Multiservices interfaces. If the encapsulated packet size exceeds the tunnel maximum transmission unit (MTU), the packet is fragmented before encapsulation. The statement is supported only on MX Series routers and all M Series routers except the M320 router.

Usage Guidelines See "Enabling Fragmentation on GRE Tunnels" on page 519.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

dial-options

Syntax

```
dial-options {
  ipsec-interface-id name;
  l2tp-interface-id name;
  (shared | dedicated);
}
```

Hierarchy Level

[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces si-*fpc/pic/port* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
The **[edit ...si-...]** hierarchy levels introduced in Junos OS Release 11.4.

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name*—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level.

l2tp-interface-id *name*—Interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 318
- Configuring Dynamic Endpoints for IPsec Tunnels on page 250
- Configuring Options for the LNS Inline Services Logical Interface

Copyright © 2013, Juniper Networks, Inc.

facility-override

Syntax facility-override *facility-name*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Override the default facility for system log reporting.

Options ***facility-name***—Name of the facility that overrides the default assignment. Valid entries include:

- **authorization**
- **daemon**
- **ftp**
- **kernel**
- **local0** through **local7**
- **user**

Usage Guidelines See "Configuring System Logging for Services Interfaces" on page 518.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

family (Interfaces)

Syntax

```
family inet {
  address address {
    ...
  }
  service {
    input {
      [ service-set service-set-name <service-filterfilter-name> ];
      post-service-filter filter-name;
    }
    output {
      [ service-set service-set-name <service-filter filter-name> ];
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure protocol family information for the logical interface.

Options *family*—Protocol family. Valid settings for service interfaces include **inet** (IPv4) and **mpls**.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the Address and Domain for Services Interfaces" on page 516 or; for a general discussion of **family** statement options, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

host (Interfaces)

Syntax

```
host hostname {
  services severity-level;
  facility-override facility-name;
  log-prefix prefix-value;
  port port-number;
}
```

Hierarchy Level [edit interfaces *interface-name* services-options syslog]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the hostname for the system logging utility.

Options ***hostname***—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.

The remaining statements are explained separately.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 520.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

inactivity-timeout

Syntax inactivity-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the inactivity timeout period for established flows. The timeout value configured in the application protocol definition overrides this value.

Options ***seconds***—Timeout period.

Default: 30 seconds

Range: 4 through 86,400 seconds

Usage Guidelines See "Configuring Default Timeout Settings for Services Interfaces" on page 516.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

input (Interfaces)

Syntax

```
input {
  service-set service-set-name <service-filter filter-name>;
  post-service-filter filter-name;
}
```

Hierarchy Level

[edit interface *interface-name* unit *logical-unit-number* family inet service],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the input service sets and filters to be applied to traffic.

Options The remaining statements are explained separately.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 520.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines For a complete description, see the Junos® OS Network Interfaces.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

log-prefix (Interfaces)

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the system logging prefix value.

Options ***prefix-value***—System logging prefix value.

Usage Guidelines See "Configuring System Logging for Services Interfaces" on page 518.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos Services Interfaces Configuration Release 12.3

maximum

Syntax maximum *number*;

Hierarchy Level [edit interfaces *interface-name* services-options session-limit]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the maximum number of sessions allowed simultaneously.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

open-timeout

Syntax open-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a timeout period for Transmission Control Protocol (TCP) session establishment.

Options ***seconds***—Timeout period.
Default: 30 seconds
Range: 4 through 224 seconds

Usage Guidelines See "Configuring Default Timeout Settings for Services Interfaces" on page 516.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

output

Syntax
```
output {
  [ service-set service-set-name <service-filter filter-name> ];
}
```

Hierarchy Level [edit interface *interface-name* unit *logical-unit-number* family inet service],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the output service sets and filters to be applied to traffic.

Options The remaining statements are explained separately.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 520.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

post-service-filter

Syntax post-service-filter *filter-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet service input],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service input]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a postservice filter on the input side of the interface only.

Options ***filter-name***—Identifier for the post-service filter.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 520.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

primary (Interfaces)

Syntax primary *interface-name*;

Hierarchy Level [edit interfaces (rsp0 | rsp1) redundancy-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary adaptive services interface.

Options ***interface-name***—The identifier for the AS or Multiservices PIC interface, which must be of the form **sp-*fpc/pic/port***.

Usage Guidelines See "Configuring AS or Multiservices PIC Redundancy" on page 522.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rate

Syntax rate *new-sessions-per-second*;

Hierarchy Level [edit interfaces *interface-name* services-options session-limit]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the maximum number of new sessions allowed per second.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

redundancy-options

Syntax
```
redundancy-options {
    primary sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
}
```

Hierarchy Level [edit interfaces (rsp0 | rsp1)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary (backup) adaptive services interfaces.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring AS or Multiservices PIC Redundancy" on page 522.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

secondary (Interfaces)

Syntax secondary *interface-name*;

Hierarchy Level [edit interfaces (rsp0 | rsp1) redundancy-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the secondary (backup) adaptive services interface.

Options ***interface-name***—The identifier for the adaptive services interface, which must be of the form **sp-*fpc/pic/port***.

Usage Guidelines See "Configuring AS or Multiservices PIC Redundancy" on page 522.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

service

Syntax

```
service {
  input {
    [ service-set service-set-name <service-filter filter-name> ];
    post-service-filter filter-name;
  }
  output {
    [ service-set service-set-name <service-filter filter-name> ];
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service sets and filters to be applied to an interface.

Options The remaining statements are explained separately.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 520.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

service-domain

Syntax service-domain (inside | outside);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the service interface domain. If you specify this interface using the **next-hop-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level, the interface domain must match that specified with the **inside-service-interface** and **outside-service-interface** statements.

Options **inside**—Interface used within the network.

outside—Interface used outside the network.

Usage Guidelines See "Configuring the Address and Domain for Services Interfaces" on page 516.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

service-filter (Interfaces)

Syntax service-filter *filter-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output) service-set *service-set-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service (input | output) service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the **service-set** statement without a **service-filter** definition, Junos OS assumes the match condition is true and selects the service set for processing automatically.

Options ***filter-name***—Identifies the filter to be applied in service processing.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Applying Filters and Services to Interfaces on page 520
- Junos Services Interfaces Configuration Release 12.3

Copyright © 2013, Juniper Networks, Inc.

service-set (Interfaces)

Syntax service-set *service-set-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output)],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service (input | output)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define one or more service sets to be applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration.

Options ***service-set-name***—Identifies the service set.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 520.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

services (Interfaces)

Syntax services *severity-level*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the system logging severity level.

Options ***severity-level***—Assigns a severity level to the facility. Valid entries include:

- **alert**—Conditions that should be corrected immediately.
- **any**—Matches any level.
- **critical**—Critical conditions.
- **emergency**—Panic conditions.
- **error**—Error conditions.
- **info**—Informational messages.
- **notice**—Conditions that require special handling.
- **warning**—Warning messages.

Usage Guidelines See "Configuring System Logging for Services Interfaces" on page 518.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

services-options

Syntax

```
services-options {
  cgn-pic;
  close-timeout
  fragment-limit
  disable-global-timeout-override;
  ignore-errors <alg> <tcp>;
  inactivity-non-tcp-timeout seconds;
  inactivity-tcp-timeout seconds;
  inactivity-timeout seconds;
  open-timeout seconds;
  reassembly-timeout
  session-limit {
    maximum number;
    rate new-sessions-per-second;
  }
  session-timeout seconds;
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-value;
      port port-number;
      services severity-level;
    }
    message-rate-limit messages-per-second;
  }
  tcp-tickles tcp-tickles;
  trio-flow-offload minimum-bytes minimum-bytes;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service options to be applied on an interface.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Interface Properties.

Copyright © 2013, Juniper Networks, Inc.

session-limit

Syntax

```
session-limit {
  maximum number;
  rate new-sessions-per-second;
}
```

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 9.6.

Description Restrict the maximum number of sessions and the session rate on Multiservices PICs.

Options **session-limit**—Restricts the maximum number of sessions and the session rate for Multiservices PICs.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

source-address

Syntax source-address *source-address*

Hierarchy Level [edit services service-set *service-set-name* syslog host *hostname*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Specify a source address to record in system log messages that are directed to a remote machine specified in the ***hostname*** statement.



NOTE: The supported interfaces are ms, rms, and mams interfaces. If you do not specify the interface parameter, the command loops on all supported interfaces.

Options ***source-address***—A valid IP address, which is recorded as the message source in messages sent to the remote machines specified in the **host *hostname*** statement

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring System Logging for Service Sets on page 479
- host on page 491
- service-set on page 503

syslog (Interfaces)

Syntax

```
syslog {
  host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number;
  }
  message-rate-limit messages-per-second;
}
```

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure generation of system log messages for the service set. System log information is passed to the kernel for logging in the **/var/log** directory. Any values configured in the service set definition override these values.

Options The remaining statements are described separately.

Usage Guidelines See "Configuring System Logging for Services Interfaces" on page 518.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

tcp-tickles

Syntax tcp-tickles *tcp-tickles*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Define the maximum number of keep-alive messages sent before a TCP session is allowed to timeout.

Options ***tcp-tickles***—Number of keep-alive messages.

Range: 0 through 30

Default: 4

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Default Timeout Settings for Services Interfaces on page 516

Copyright © 2013, Juniper Networks, Inc.

trio-flow-offload

Syntax trio-flow-offload minimum-bytes *minimum-bytes*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 12.1.

Description Enable any plug-in or daemon on a PIC to generate a flow offload request to of-load flows to the Packet Forwarding Engine. This command is available on MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs).

Options ***minimum-bytes***—The minimum number of bytes that trigger offloading. When this option is omitted, offloading is triggered when both the forward and reverse flows of the session have begun, meaning that at least one packet has flowed in each direction.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Flow Offloading on page 525

Copyright © 2013, Juniper Networks, Inc.

unit

Syntax

```
unit logical-unit-number {
  family inet {
    address address {
    }
    service {
      input {
        [ service-set service-set-name <service-filter filter-name> ];
        post-service-filter filter-name;
      }
      output {
        [ service-set service-set-name <service-filter filter-name> ];
      }
    }
    service-domain (inside | outside);
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Usage Guidelines For a general discussion of logical interface properties, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 28

PGCP Configuration Guidelines for the BGF Feature

To configure the border gateway function (BGF), include the **pgcp** statement at the **[edit services]** hierarchy level:

```
[edit services]
pgcp {
  gateway gateway-name {
    cleanup-timeout seconds;
    gateway-address gateway-address;
    fast-update-filters {
      maximum-terms number-of-terms;
      maximum-fuf-percentage percentage;
    }
    gateway-controller gateway-controller-name {
      controller-address ip-address;
      controller-port port-number;
      interim-ah-scheme {
        algorithm algorithm;
      }
    }
    gateway-port gateway-port;
    graceful-restart {
      maximum-synchronization-mismatches number-of-mismatches;
      seconds;
    }
    data-inactivity-detection {
      inactivity-delay;
      latch-deadlock-delay seconds;
      no-rtcp-check;
      send-notification-on-delay;
      inactivity-duration seconds;
      stop-detection-on-drop;
      report-service-change {
        service-change-type (forced-906) | forced-910);
      }
    }
    h248-options {
      audit-observed-events-returns;
      encoding {
        no-dscp-bit-mirroring;
```

Copyright © 2013, Juniper Networks, Inc.

```
        use-lower-case
    }
    h248-profile {
        profile-name profile-name;
        profile-version version-number;
    }
    service-change {
        control-association-indications {
            disconnect {
                controller-failure (failover-909 | restart-902);
                reconnect (disconnected-900 | restart-902);
            }
            down {
                administrative (forced-905 | forced-908 | none);
                failure (forced-904 | forced-908 | none);
                graceful (graceful-905 | none);
            }
            up {
                cancel-graceful (none | restart-918);
                failover-cold (failover-920 | restart-901);
                failover-warm (failover-919 | restart-902);
            }
        }
        virtual-interface-indications {
            virtual-interface-down {
                administrative (forced-905 | forced-906 | none);
                graceful (graceful-905 | none);
            }
            virtual-interface-up {
                cancel-graceful (none | restart-918);
                warm (none | restart-900);
            }
        }
        context-indications {
            state-loss (forced-910 | forced-915 | none);
        }
        use-wildcard-response;
    }
}
h248-properties {
    application-data-inactivity-detection {
        ip-flow-stop-detection (regulated-notify | immediate-notify);
    }
    base-root {
        mg-provisional-response-timer-value {
            default milliseconds;
            maximum milliseconds;
            minimum milliseconds;
        }
        mgc-provisional-response-timer-value {
            default milliseconds;
            maximum milliseconds;
            minimum milliseconds;
        }
        mg-originated-pending-limit {
            default number-of-messages;
```


Copyright © 2013, Juniper Networks, Inc.

```
      maximum number-of-messages;
      minimum number-of-messages;
    }
    mgc-originated-pending-limit {
      default number-of-messages;
      maximum number-of-messages;
      minimum number-of-messages;
    }
    normal-mg-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    normal-mgc-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
  }
  diffserv {
    dscp {
      default (dscp-value | alias | do-not-change);
    }
  }
  event-timestamp-notification {
    request-timestamp (requested | suppressed | autonomous);
  {
  hanging-termination-detection {
    timerx seconds;
  }
  ipsec-transport-security-association security-association-name;
  notification-behavior {
    notification-regulation default (once | 0 - 100);
  }
  platform {
    device interface-name;
    routing-engine;
  }
  segmentation {
    mg-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mgc-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mg-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
    mgc-maximum-pdu-size {
      default bytes;
```

Copyright © 2013, Juniper Networks, Inc.

```
            maximum bytes;
            minimum bytes;
        }
    }
    traffic-management {
        max-burst-size {
            default bytes;
            maximum bytes;
            minimum bytes;
            rtcp {
                (fixed-value bytes-per-second | percentage percentage);
            }
        }
        peak-data-rate {
            default bytes-per-second;
            maximum bytes-per-second;
            minimum bytes-per-second;
            rtcp {
                (fixed-value bytes-per-second | percentage percentage);
            }
        }
        sustained-data-rate {
            default bytes-per-second;
            maximum bytes-per-second;
            minimum bytes-per-second;
            rtcp {
                (fixed-value bytes-per-second | percentage percentage);
            }
        }
    }
    inactivity-timer {
        inactivity-timeout {
            detect;
            maximum-inactivity-time {
                default 10-millisecond-units;
                maximum 10-millisecond-units;
                minimum 10-millisecond-units;
            }
        }
    }
}
h248-timers {
    initial-average-ack-delay milliseconds;
    maximum-net-propagation-delay milliseconds;
    maximum-waiting-delay milliseconds;
    tmax-retransmission-delay milliseconds;
}
max-concurrent-calls number-of-calls;
monitor {
    media {
        rtcp;
        rtp;
    }
}
service-state (in-service | out-of-service-forced | out-of-service-graceful);
session-mirroring {
```


Copyright © 2013, Juniper Networks, Inc.

```
    delivery-function delivery-function-name {
      destination-address destination-address;
      destination-port destination-port;
      network-operator-id network-operator-id;
      source-address source-address;
      source-port (Mirrored BGF Packets) source-port;
    }
    disable-session-mirroring;
  }
}
rule rule-name {
  gateway gateway-name;
  nat-pool nat-pool-name;
}
rule-set rule-set-name {
  rule rule-name;
}
session-mirroring {
  delivery-function delivery-function-name {
    destination-address destination-address;
    destination-port destination-port;
    network-operator-id network-operator-id;
    source-address source-address;
    source-port (Mirrored BGF Packets) source-port;
  }
  disable-session-mirroring;
}
traceoptions {
  file <filename filename> <files number> <match regex> <size size> <world-readable
    | no-world-readable>;
  flag {
    bgf-core {
      common trace-level;
      default trace-level;
      firewall trace-level;
      gate-logic trace-level;
      pic-broker trace-level;
      policy trace-level;
      statistics trace-level;
    }
    default trace-level;
    h248-stack {
      control-association trace-level;
      default trace-level;
      messages;
      media-gateway trace-level;
    }
    sbc-utils {
      common trace-level;
      configuration trace-level;
      default trace-level;
      device-monitor trace-level;
      ipc trace-level;
      memory-management trace-level;
      messaging trace-level;
      user-interface trace-level;
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
    virtual-interface number {
      nat-pool nat-pool-name;
        routing-instance instance-name {
        service-interface interface-name.unit-number;
      }
      service-state (in-service | out-of-service-forced | out-of-service-graceful);
    }
  }
```

For information about using the PGCP statements to configure the BGF feature, see the *Session Border Control Solutions Guide Using BGF and IMSG*.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 29

Summary of PGCP Configuration Statements

The following sections explain each of the PGCP statements, which are used to configure the BGF feature. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

administrative

See the following sections:

- administrative (Control Association) on page 554
- administrative (Virtual Interface) on page 555

administrative (Control Association)

Syntax administrative (forced-905 | forced-908 | none);

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name* h248-options service-change control-association-indications down]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Unregistration Messages in ServiceChange commands that it sends to the gateway controller when a control association transitions to Out-of-Service because of an administrative operation.

Default If you do not specify an option, the virtual BGF includes FO/905 (**forced-905**).

Options

forced-905—Termination is being taken out of service. The virtual BGF is transitioning to Out-of-Service because of an administrative operation.

forced-908—Termination is being taken out of service. The virtual BGF is transitioning to Out-of-Service because of an administrative operation or error.

none—The virtual BGF does not send a ServiceChange command to the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

administrative (Virtual Interface)

Syntax administrative (forced-905 | forced-906 | none);

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name* h248-options service-change virtual-interface-indications virtual-interface-down]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Service-Interruption ServiceChange commands that it sends to the gateway controller when a virtual interface changes to Out-of-Service because of an administrative operation.

Default If you do not specify an option, the virtual BGF includes FO/905 (**forced-905**).

Options

forced-905—Termination is being taken out of service. The virtual interface is transitioning to Out-of-Service because of an administrative operation.

forced-906—Loss of lower-layer connectivity. The virtual interface is transitioning to Out-of-Service because of a loss of Layer 2 connectivity caused by the logical or physical interface being administratively disabled.

none—Virtual BGF does not send a ServiceChange command.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

algorithm

Syntax algorithm *algorithm*;

Hierarchy Level [edit services pgcp gateway *gateway-name* gateway-controller *gateway-controller-name* interim-ah-scheme]

Release Information Statement introduced in Junos OS Release 8.4.

Description Specify the algorithm for the interim AH scheme. Once you set the algorithm for the interim AH scheme, to disable the interim AH scheme, you need to remove the algorithm and restart the PGCP service.

Options ***algorithm***—Algorithm used for the interim AH scheme. HMAC null is currently the only algorithm supported.

Values: hmac-null

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

application-data-inactivity-detection

Syntax
```
application-data-inactivity {
  ip-flow-stop-detection (regulated-notify | immediate-notify);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 9.3.

Description Activate or deactivate regulated notification of media inactivity events.

Options The statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring H.248 Notification Behavior to Prevent Excessive Media Inactivity Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

audit-observed-events-returns

Syntax audit-observed-events-returns;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options]

Release Information Statement introduced in Junos OS Release 9.3.

Description Enable a history of media inactivity events to be viewed by the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring H.248 Notification Behavior to Prevent Excessive Media Inactivity Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

base-root

Syntax

```
base-root {
  mg-provisional-response-timer-value {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
  }
  mgc-provisional-response-timer-value {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
  }
  mg-originated-pending-limit {
    default number-of-messages;
    maximum number-of-messages;
    minimum number-of-messages;
  }
  mgc-originated-pending-limit {
    default number-of-messages;
    maximum number-of-messages;
    minimum number-of-messages;
  }
  normal-mg-execution-time {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
  }
  normal-mgc-execution-time {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure default values for properties in the base root package defined in Annex E of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

bgf-core

Syntax

```
bgf-core {
  default trace-level;
  firewall trace-level;
  gate-logic trace-level;
  pic-broker trace-level;
  policy trace-level;
  statistics trace-level;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure trace-level options for the BGF core component of the virtual BGF.

Options

default *trace-level*—Default trace level for all **bgf-core** messages.

firewall *trace-level*—Trace level for the **firewall** subcomponent, which controls firewall filters, connections, and all relevant firewall activities.

gate-logic *trace-level*—Trace level for the **gate-logic** subcomponent, which controls gate common logic, gate lookup, and gate manager activities.

pic-broker *trace-level*—Trace level for the **pic-broker** subcomponent, which controls gates on the PIC.

policy *trace-level*—Trace level for the **policy** subcomponent, which controls media function and socket policy.

statistics *trace-level*—Trace level for the **statistics** subcomponent, which provides pgcpd statistics.

trace-level—Trace-level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

cancel-graceful

See the following sections:

- cancel-graceful (Control Association) on page 561
- cancel-graceful (Virtual Interface) on page 562

cancel-graceful (Control Association)

Syntax cancel-graceful (none | restart-918);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications up]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Notification ServiceChange commands that it sends to the gateway controller when the control association transitions from the Draining state to the Forwarding state.

Default If you do not specify an option, the virtual BGF does not send a ServiceChange command.

Options **none**—The virtual BGF does not send a ServiceChange command to the gateway controller.

restart-918—The control association has returned to the Forwarding state.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

cancel-graceful (Virtual Interface)

Syntax cancel-graceful (none | restart-918);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change virtual-interface-indications virtual-interface-up]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Notification ServiceChange commands that it sends to the gateway controller when the virtual interface transitions from In-Service to Out-of-Service-Graceful.

Default If you do not specify an option, the virtual BGF does not send a ServiceChange command.

Options **none**—Virtual BGF does not send a ServiceChange command.

restart-918—Cancel graceful. The virtual interface has entered the Draining state.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

cleanup-timeout

Syntax cleanup-timeout *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the number of seconds before the virtual BGF automatically deletes all gates following a disconnection from the gateway controller.

Options ***seconds***—Interval before inactivity detection starts.
Range: 0 through 65,535 seconds
Default: 3600 seconds

Required Privilege Level interface-level—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

context-indications

Syntax

```
context-indications {
  state-loss (forced-910 | forced-915 | none);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Service-Interruption ServiceChange commands that it sends to the gateway controller when the gates of a context no longer provide their configured services. When the virtual BGF sends a Service-Interruption message, both terminations in the context become Out-of-Service.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Contexts in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

control-association-indications

Syntax

```
control-association-indications {
  disconnect {
    controller-failure (failover-909 | restart-902);
    reconnect (disconnected-900 | restart-902);
  }
  down {
    administrative (forced-905 | forced-908 | none);
    failure (forced-904 | forced-908 | none);
    graceful (graceful-905 | none);
  }
  up {
    cancel-graceful (none | restart-918);
    failover-cold (failover-920 | restart-901);
    failover-warm (failover-919 | restart-902);
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in ServiceChange commands that it sends to the gateway controller when the state of the control association changes.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

controller-address

Syntax controller-address *ip-address*;

Hierarchy Level [edit services pgcp gateway *gateway-name* gateway-controller *gateway-controller-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure an IP address for the gateway controller.

Options ***ip-address***—IP address of the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Gateway Controller in *Session Border Control Solutions Guide Using BGF and IMSG*

controller-failure

Syntax controller-failure (failover-909 | restart-902);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications disconnect]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Registration Request ServiceChange commands when it attempts to reregister with the gateway controller or register with a new gateway controller after the control association is disconnected.

Default If you do not specify an option, the virtual BGF includes RS/902 (**restart-902**).

Options **failover-909**—Gateway controller impending failure. The virtual BGF is reregistering with a new gateway controller following a disconnection of the virtual BGF and gateway controller.

restart-902—Warm boot. The virtual BGF is attempting to reregister with existing states after a gateway controller failure.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

controller-port

Syntax controller-port *port-number*;

Hierarchy Level [edit services pgcp gateway *gateway-name* gateway-controller *gateway-controller-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the port number of the gateway controller listening port. The virtual BGF sends H.248 messages to this port.

Options ***port-number***—Port number of the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

data-inactivity-detection (Services PGCP)

Syntax
```
data-inactivity-detection {
  inactivity-delay seconds;
  inactivity-duration seconds;
  latch-deadlock-delay seconds;
  no-rtcp-check;
  send-notification-on-delay;
  stop-detection-on-drop;
  report-service-change {
    service-change-type (forced-906) | forced-910);
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure data inactivity detection to detect latch deadlocks or other media inactivity on a gate.

Options The statements are described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

default

Syntax default *trace-level*;

Hierarchy Level [edit services pgcp gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the minimum trace level for all selected PGCP trace options. This option overrides individual trace options that are set at a lower level.

Default warning

Options ***trace-level***—Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

delivery-function

Syntax

```
delivery-function delivery-function-name {
  destination-address destination-address;
  destination-port destination-port;
  network-operator-id network-operator-id;
  source-address source-address;
  source-port (Mirrored BGF Packets) source-port;
}
```

Hierarchy Level [edit services pgcp session-mirroring],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring]

Release Information Statement introduced in Junos OS Release 9.2

Description Configure the delivery function that receives the session mirroring information. You can configure only one delivery function.

Options ***delivery-function-name***—Name of the delivery function that receives the session mirroring information.

Required Privilege Level pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation

- Setting Up Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

destination-address

Syntax destination-address *destination-address*;

Hierarchy Level [edit services pgcp session-mirroring delivery-function *delivery-function-name*],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring delivery-function *delivery-function-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the address of the delivery function server to which the BGF sends session-mirroring information.

Options ***destination-address***—Address of the server to which the BGF sends session-mirroring information.

Required Privilege Level pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation

- Setting Up Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

destination-port

Syntax destination-port *destination-port*;

Hierarchy Level [edit services pgcp session-mirroring delivery-function *delivery-function-name*],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring delivery-function *delivery-function-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the port on the delivery function server that receives session-mirroring information.

Options ***destination-port***—Port on the delivery function server that receives session-mirroring information.

Range: 1 through 65,535

Required Privilege Level pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation
- Setting Up Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

detect

Syntax detect;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties inactivity-timer inactivity-timeout]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify whether the BGF detects inactivity timeout events received from the BGF by default.

Default The BGF does not detect inactivity timeout events by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Detecting Gateway Controller Failures in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

diffserv

Syntax

```
diffserv {
  dscp {
    default (dscp-value | alias | do-not-change);
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure default values for properties in the Differentiated Services (DiffServ) package defined in Annex A.2 of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options Statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface–control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

disable-session-mirroring

Syntax disable-session-mirroring;

Hierarchy Level [edit services pgcp session-mirroring],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring]

Release Information Statement introduced in Junos OS Release 9.2.

Description Disable or enable session mirroring on the BGF. To disable session mirroring, enter **set disable-session-mirroring**. To enable session mirroring, enter **delete disable-session-mirroring**.

Required Privilege Level pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation

- Disabling Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

disconnect

Syntax

```
disconnect {
  controller-failure (failover-909 | restart-902)
  reconnect (disconnected-900 | restart-902)
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Registration Request ServiceChange commands when it attempts to reregister with the gateway controller or register with a new gateway controller after the control association is disconnected.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in Service Change Commands in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

down

Syntax

```
down {
  administrative (forced-905 | forced-908 | none);
  failure (forced-904 | forced-908 | none);
  graceful (graceful-905 | none);
}
```

Hierarchy Level

[edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Specify the method and reason that the virtual BGF includes in Unregistration Messages in ServiceChange commands that it sends to the gateway controller when a control association transitions to Out-of-Service because of a failure. The failure can be the result of a services PIC or DPC, or because the services PIC or DPC was powered off or removed.

Options

The options are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfacesin *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

dscp (Services PGCP)

Syntax

```
dscp {
  default (dscp-value | alias | do-not-change);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties diffserv]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure default values for DSCP marking that the virtual BGF uses for outgoing traffic when the DSCP value is not already defined by the gateway controller.

Default The default DSCP value that the virtual BGF uses is zero (0x00).

Options

dscp-value—Specify a string of eight bits or a 1-byte hexadecimal value using the format: 0x*XX*. Currently, only six bits are used by the packet.

alias—Specify a standard DSCP name. The standard name is translated to an 8-bit string with the two least significant bits (LSBs) as zeros; for example, **EF=10111000**.

do-not-change—Specify that no DSCP action be performed on the PIC or DPC. The egress value on the gate is the same as the ingress DSCP value.

Required Privilege Level interface—To view this statement in the configuration.
interface–control—To add this statement to the configuration.

Related Documentation

- Quality of Service for VoIP Traffic Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

encoding

Syntax

```
encoding {
  no-dscp-bit-mirroring;
  use-lower-case;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options]

Release Information Statement introduced in Junos OS Release 9.3.
use-lower-case option introduced in Release 9.5.

Description Change encoding defaults.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

event-timestamp-notification

Syntax
```
event-timestamp-notification {
  request-timestamp (requested | suppressed | autonomous);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248–properties]

Release Information Statement introduced in Junos OS Release 9.3.

Description Enable or disable access by the gateway controller to timestamp information for media inactivity event notifications.

Options The statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Quality of Service for VoIP Traffic Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

failover-cold

Syntax failover-cold (failover-920 | restart-901);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications up]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Registration ServiceChange commands when it attempts to register with a new gateway controller following a cold failover.

Default If you do not specify an option, the virtual BGF includes RS/901 (**restart-901**).

Options

failover-920—Cold failover. The virtual BGF is registering following a graceful Routing Engine switchover. The installed state is reset.

restart-901—Cold boot. The virtual BGF is transitioning to In-Service. The previously installed state is not retained.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

failover-warm

Syntax failover-warm (failover-919 | restart-902);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications up]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Registration ServiceChange commands when it attempts to register with a new gateway controller following a warm failover.

Default If you do not specify an option, the virtual BGF includes RS/902 (**restart-902**).

Options

failover-919—Gateway controller impending failure. The virtual BGF is registering with a new gateway controller after the virtual BGF and the gateway controller were disconnected.

restart-902—Warm boot. The virtual BGF is transitioning to In-Service. The previously installed state is retained.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

failure

Syntax failure (forced-904 | forced-908 | none);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications down]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Unregistration or Notification Messages in ServiceChange commands when a control association transitions to Out-of-Service.

Default If you do not specify an option, the virtual BGF sends ServiceChange command forced-904 to the gateway controller.

Options

forced-904—Termination malfunctioning. The virtual BGF is transitioning to Out-of-Service because of a failure.

forced-908—The virtual BGF is transitioning to Out-of-Service due to administrator action or a failure.

none—The virtual BGF does not send a ServiceChange command to the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

fast-update-filters

Syntax

```
fast-update-filters {
  maximum-terms number-of-terms;
  maximum-fuf-percentage percentage-of-gates;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.1.

Description Limit the number of FUF terms installed on the Packet Forwarding Engine for a virtual BGF to improve performance when the software is collecting statistics on packets that are dropped because they exceed the rate limits set in fast update filters (FUFs).

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Improving Performance While Collecting Gate Statistics in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

file (Services PGCP)

Syntax file <*filename*> <files *files*> <match *regular-expression*> <size *maximum-file-size*> <world-readable | no-world-readable>;

Hierarchy Level [edit services pgcp traceoptions]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the trace file for tracing BGF components.

Options

filename *filename*—Name of the file to which the tracing messages are written.

Default: bsg_trace

files *number-of-files*—Number of trace files. The tracing mechanism can rotate between any given number of files, allowing for trace message inspection without interfering with the normal work of the application.

Default: 3

match *regular expression*—Regular expression to match with incoming messages. Messages that do not match the regular expression are not written to the trace file.

size *maximum-trace-file-size*—Size parameter (in bytes) to trigger rotation of files. The trace mechanism rotates files based on the current file size. When the size is bigger than the maximum configured size, the files are rotated.

Default: 1048576

world-readable | no-world-readable—Allow all users to use the log file or disallow all users from using the log file.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

flag (Services PGCP)

Syntax

```
flag {
  default trace-level;
  bgf-core {
    default trace-level;
    firewall trace-level;
    gate-logic trace-level;
    pic-broker trace-level;
    policy trace-level;
    statistics trace-level;
  }
  h248-stack {
    default trace-level;
    messages;
    control-association trace-level;
    media-gateway trace-level;
  }
  sbc-utils {
    common trace-level;
    configuration trace-level;
    device-monitor trace-level;
    ipc trace-level;
    memory-management trace-level;
    message trace-level;
    minimum trace-level;
    user-interface trace-level;
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* traceoptions]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure trace options for components of the BGF.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

gateway (Services PGCP)

Syntax

```
pgcp {
  gateway gateway-name {
    cleanup-timeout seconds;
    gateway-address gateway-address;
    gateway-controller gateway-controller-name {
      local-controller | remote-controller;
      controller-address ip-address;
      controller-port port-number;
      interim-ah-scheme {
        algorithm algorithm;
      }
    }
    gateway-port gateway-port;
    service-state (in-service | out-of-service-forced | out-of-service-graceful);
    graceful-restart {
      maximum-synchronization-mismatches number-of-mismatches;
      seconds;
    }
    data-inactivity-detection {
      inactivity-delay seconds;
      latch-deadlock-delay seconds;
      send-notification-on-delay;
      inactivity-duration seconds;
      no-rtcp-check;
      stop-detection-on-drop;
      report-service-change {
        service-change-type (forced-906 | forced-910);
      }
    }
    h248-properties {
      application-data-inactivity-detection {
        ip-flow-stop-detection (regulated-notify | immediate-notify);
      }
      base-root {
        mg-provisional-response-timer-value {
          default milliseconds;
          maximum milliseconds;
          minimum milliseconds;
        }
        mgc-provisional-response-timer-value {
          default milliseconds;
          maximum milliseconds;
          minimum milliseconds;
        }
        mg-originated-pending-limit {
          default number-of-messages;
          maximum number-of-messages;
          minimum number-of-messages;
        }
        mgc-originated-pending-limit {
          default number-of-messages;
          maximum number-of-messages;
```


Copyright © 2013, Juniper Networks, Inc.

```
      minimum number-of-messages;
    }
    normal-mg-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    normal-mgc-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
  }
  diffserv {
    dscp {
      default (dscp-value | alias | do-not-change);
    }
  }
  event-timestamp-notification {
    request-timestamp (requested | suppressed | autonomous);
  }
  hanging-termination-detection {
    timerx seconds;
  }
  ipsec-transport-security-association security-association-name;
  notification-behavior {
    notification-regulation default (once | 0-100);
  }
  platform {
    device interface-name;
    routing-engine;
  }
  segmentation {
    mg-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mgc-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mg-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
    mgc-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
  }
  traffic-management {
    max-burst-size {
```

Copyright © 2013, Juniper Networks, Inc.

```
            default bytes-per-second;
            maximum bytes-per-second;
            minimum bytes-per-second;
            rtcp {
              (fixed-value bytes-per-second | percentage percentage);
            }
          }
          peak-data-rate {
            default bytes-per-second;
            maximum bytes-per-second;
            minimum bytes-per-second;
            rtcp {
              (fixed-value bytes-per-second | percentage percentage);
            }
            rtcp-include;
          }
          sustained-data-rate (All Streams) {
            default bytes-per-second;
            maximum bytes-per-second;
            minimum bytes-per-second;
            rtcp {
              (fixed-value bytes-per-second | percentage percentage);
              }
            rtcp-include;
          }
        }
        inactivity-timer {
          inactivity-timeout {
            detect;
            maximum-inactivity-time {
              default 10-millisecond-units;
              maximum 10-millisecond-units;
              minimum 10-millisecond-units;
            }
          }
        }
      }
      h248-options {
        accept-emergency-calls-while-graceful;
        audit-observed-events-returns;
        encoding {
          no-dscp-bit-mirroring;
          use-lower-case;
        }
        h248-profile {
          profile-name profile-name;
          profile-version version-number;
        }
        implicit tcp-latch;
        implicit-tcp-source-filter;
        service-change {
          control-association-indications {
            disconnect {
              controller-failure (failover-909 | restart-902);
              reconnect (disconnected-900 | restart-902);
            }
```


Copyright © 2013, Juniper Networks, Inc.

```
            down {
              administrative (forced-905 | forced-908 | none);
              failure (forced-904 | forced-908 | none);
              graceful (graceful-905 | none );
            }
            up {
              cancel-graceful (none | restart-918);
              failover-cold (failover-920 | restart-901);
              failover-warm (failover-919 | restart-902);
            }
          }
          virtual-interface-indications {
            virtual-interface-down {
              administrative (forced-905 | forced-906 | none);
              graceful (graceful-905 | none);
            }
            virtual-interface-up {
              cancel-graceful (Virtual Interface) (none | restart-918);
              warm (none | restart-900);
            }
          }
          context-indications {
            state-loss (forced-910 | forced-915 | none);
          }
          use-wildcard-response;
        }
      }
      h248-timers {
        initial-average-ack-delay milliseconds;
        maximum-net-propagation-delay milliseconds;
        maximum-waiting-delay milliseconds;
        tmax-retransmission-delay milliseconds;
      }
      max-concurrent-calls number-of-calls;
      monitor {
        media {
          rtcp;
          rtp;
        }
      }
    service-state (in-service | out-of-service-forced | out-of-service-graceful);
      session-mirroring {
        delivery-function delivery-function-name {
          destination-address destination-address;
          destination-port destination-port;
          network-operator-id network-operator-id;
          source-address source-address;
          source-port (Mirrored BGF Packets) source-port;
        }
        disable-session-mirroring;
      }
    }
  }
```

Hierarchy Level [edit services pgcp]

Copyright © 2013, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 8.4.
graceful-restart option introduced in Junos OS Release 8.5.
h248-options option introduced in Junos OS Release 8.5.
h248-properties option introduced in Junos OS Release 8.5.
monitor option introduced in Junos OS Release 9.0.
session-mirroring option introduced in Junos OS Release 9.2.
data-inactivity-detection option introduced in Junos OS Release 9.3.
overload-control option introduced in Junos OS Release 9.3.
platform option introduced in Junos OS Release 9.6.
h248–profile option introduced in Junos OS Release 10.0.
ipsec-transport-security-association option introduced in Junos OS Release 10.0.

Description Configure a virtual BGF on the router.

Options *gateway-name*—Identifier of the virtual BGF. You can configure an IP address as the gateway name. However, the IP address is not used in the operation of the virtual BGF.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface–control—To add this statement to the configuration.

Related Documentation
- BGF VoIP Solution Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- BGF VoIP Solution Architecture in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

gateway-address

Syntax gateway-address *gateway-address*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the IP address of the virtual BGF.

Options ***gateway-address***—IP address of the virtual BGF that you are configuring on the router.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

gateway-controller

Syntax

```
gateway-controller gateway-controller-name {
  local-controller | remote-controller;
  <controller-address ip-address;>
  <controller-port port-number;>
  interim-ah-scheme {
    algorithm algorithm;
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.4.
local-controller option introduced in Junos OS Release 9.4.
remote-controller option introduced in Junos OS Release 9.4.

Description Configure a gateway controller.

Options *gateway-controller-name*—Name of the gateway controller or BSG. You can configure an IP address as the gateway controller name. However, the IP address is not used for the connection to the gateway controller.

local-controller | **remote-controller**—Type of gateway controller.

- **remote-controller**. Configure the gateway controller as a remote controller if you are using an external gateway controller. You must specify **controller-address** and **controller-port**.
- **local-controller**. Configure the gateway controller as a local controller if you are using a border signaling gateway (BSG).

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

gateway-port

Syntax gateway-port *gateway-port*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure a port number for the virtual BGF.

Options ***gateway-port***—Port number of the virtual BGF that you are configuring on the router.
Range: 0 through 65,535
Default: 2944

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

graceful

See the following sections:

- graceful (Control Association) on page 587
- graceful (Virtual Interface) on page 588

graceful (Control Association)

Syntax graceful (graceful-905 | none);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications down]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Notification ServiceChange commands that it sends to the gateway controller when the control association transitions from In-Service to Out-of-Service-Graceful.

Default If you do not specify an option, the virtual BGF does not send a ServiceChange command.

Options

graceful-905—Termination is being taken out of service. The control association has entered the Draining state.

none—The virtual BGF does not send a ServiceChange command to the gateway controller.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

graceful (Virtual Interface)

Syntax graceful (graceful-905 | none);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change virtual-interface-indications virtual-interface-down]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Notification ServiceChange commands that it sends to the gateway controller when the virtual interface transitions from In-Service to Out-of-Service-Graceful.

Default If you do not specify an option, the virtual BGF does not send a ServiceChange command.

Options **graceful-905**—Termination is being taken out of service. The interface has entered the Draining state.

none—Virtual BGF does not send a ServiceChange command.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

graceful-restart

Syntax
```
graceful-restart {
  maximum-synchronization-mismatches seconds;
  seconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure graceful restart properties that are used during synchronization between the pgcpd process and the Multiservices PIC or DPC.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Synchronization Properties in Case of Routing Engine Failure in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

h248-options

Syntax

```
h248-options {
  accept-emergency-calls-while-graceful;
  audit-observed-events-returns;
  encoding {
    no-dscp-bit-mirroring;
    use-lower-case;
  }
  h248-profile {
    profile-name profile-name;
    profile-version version-number;
  }
  implicit-tcp-latch;
  implicit-tcp-source-filter;
  service-change {
    control-association-indications {
      disconnect {
        controller-failure (failover-909 | restart-902);
        reconnect (disconnected-900 | restart-902);
      }
      down {
        administrative (forced-905 | forced-908 | none);
        failure (forced-904 | forced-908 | none);
        graceful (graceful-905 | none);
      }
      up {
        cancel-graceful (none | restart-918);
        failover-cold (failover-920 | restart-901);
        failover-warm (failover-919 | restart-902);
      }
    }
    virtual-interface-indications {
      virtual-interface-down {
        administrative (forced-905 | forced-906 | none);
        graceful (graceful-905 | none);
      }
      virtual-interface-up {
        cancel-graceful (none | restart-918);
        warm (none | restart-900);
      }
    }
    context-indications {
      state-loss (forced-910 | forced-915 | none);
    }
    use-wildcard-response;
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.
accept-emergency-calls-while-graceful option introduced in Junos OS Release 10.2.
audit-observed-events-returns option introduced in Junos OS Release 9.3.

Copyright © 2013, Juniper Networks, Inc.

encoding option introduced in Junos OS Release 9.3.
service-change option introduced in Junos OS Release 9.3.
use-lower-case option introduced in Junos OS Release 9.5.
h248-profile option introduced in Junos OS Release 10.0.
-latch option introduced in Junos OS Release 10.4.
-source-filter option introduced in Junos OS Release 10.4.

Description Configure options that affect virtual BGF H.248 behavior.

Options

accept-emergency-calls-while-graceful—Accept emergency calls when the BGF is in a draining state due to a graceful shutdown.

-latch—If explicit latching has been applied (using using ipnapt/latch) on either gate of a gate pair, implicit latching is *not* applied. If explicit latching has not been applied on either gate, latching is applied to both gates of the gate pair. When either of the gates latches, latching is automatically disabled on the other gate.

-source-filter—**-source-filter**—Applies source address (but not source port) filtering on incoming packets, using the current remote destination address if explicit source filtering has not been applied by use of gm/saf or ipnapt/latch.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Preventing Excessive Media Inactivity Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*
- Enabling Wildcards for ServiceChange Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Implicit Latching for TCP Gates in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

h248-profile

Syntax

```
h248-profile {
  profile-name profile-name;
  profile-version version-number;
}
```

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name* h248-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the profile that the BGF declares in the initial registration ServiceChange request. The profile is declared according to the H.248 standard. That is, profile-name/profile-version. For example, ETSI_BGF/1.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the H.248 Profile in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

h248-properties

Syntax

```
h248-properties {
  application-data-inactivity-detection {
    ip-flow-stop-detection (regulated-notify | immediate-notify)
  }
  base-root {
    mg-provisional-response-timer-value {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mgc-provisional-response-timer-value {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mg-originated-pending-limit {
      default number-of-messages;
      maximum number-of-messages;
      minimum number-of-messages;
    }
    mgc-originated-pending-limit {
      default number-of-messages;
      maximum number-of-messages;
      minimum number-of-messages;
    }
    normal-mg-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    normal-mgc-execution-time {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
  }
  diffserv {
    dscp {
      default (dscp-value | alias | do-not-change);
    }
  }
  event-timestamp-notification {
    request-timestamp (requested | suppressed | autonomous)
  {
  hanging-termination-detection {
    timerx seconds;
  }
  segmentation {
    mg-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    mgc-segmentation-timer {
      default milliseconds;
      maximum milliseconds;
      minimum milliseconds;
    }
    mg-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
    mgc-maximum-pdu-size {
      default bytes;
      maximum bytes;
      minimum bytes;
    }
  }
  traffic-management {
    max-burst-size {
      default bytes-persecond;
      maximum bytes-per-second;
      minimum bytes-per-second;
      rtcp {
        (fixed-value bytes-per-second | percentage percentage);
      }
    }
    peak-data-rate {
      default bytes-per-second;
      maximum bytes-per-second;
      minimum bytes-per-second;
      rtcp {
        (fixed-value bytes-per-second | percentage percentage);
      }
    }
    sustained-data-rate (All Streams) {
      default bytes-per-second;
      maximum bytes-per-second;
      minimum bytes-per-second;
      rtcp {
        (fixed-value bytes-per-second | percentage percentage);
      }
      rtcp-include;
    }
  inactivity-timer {
    inactivity-timeout {
      detect;
      maximum-inactivity-time {
        default 10-millisecond-units;
        maximum 10-millisecond-units;
        minimum 10-millisecond-units;
      }
    }
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.
diffserv option introduced in Junos OS Release 9.0.
inactivity-timer option introduced in Junos OS Release 9.2.
traffic-management option introduced in Junos OS Release 9.2.
application-data-inactivity-detection option introduced in Junos OS Release 9.3.
event-timestamp-notification option introduced in Junos OS Release 9.3.

Description Configure default values for H.248 properties.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring H.248 Segmentation Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

h248-stack

Syntax

```
h248-stack {
  default trace-level;
  messages trace-level;
  control-association trace-level;
  media-gateway trace-level;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure trace-level options for the H.248 stack component of the virtual BGF.

Options

default *trace-level*—Default trace level for all **h248-stack** messages.

messages—When this option is set, H.248 messages are written to the log file.

control-association *trace-level*—Trace level for traces relevant to the H.248 control association.

media-gateway *trace-level*—Trace level for libpgcp.

trace-level—Trace-level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

h248-timers

Syntax

```
h248-timers {
  initial-average-ack-delay milliseconds;
  maximum-net-propagation-delay milliseconds;
  maximum-waiting-delay milliseconds;
  tmax-retransmission-delay milliseconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure H.248 timers for the PGCP connection.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

hanging-termination-detection

Syntax

```
hanging-termination-detection {
  timerx seconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable and configure hanging termination detection.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

inactivity-delay

Syntax inactivity-delay *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the time after which the virtual BGF begins checking for data packets on terminations that do not include a latch event.

Options ***seconds***—Time interval before checking for media inactivity.
Range: 0 through 3600
Default: 5

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

inactivity-duration (Services PGCP)

Syntax inactivity-duration *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the time interval that determines inactivity. When the virtual BGF determines that the time since the last packet was received exceeds this duration, the virtual BGF generates an inactivity notification or service change request. The duration timer is the same for terminations with latch events and for terminations without latch events.

Options ***seconds***—Time during which no packets are received.
Range: 5 through 86400
Default: 30

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

inactivity-timeout

Syntax

```
inactivity-timeout {
  detect;
  maximum-inactivity-time {
    default 10-millisecond-units;
    maximum 10-millisecond-units;
    minimum 10-millisecond-units;
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties inactivity-timer]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the inactivity timeout event. The inactivity timeout event is used to detect that the inactivity timer has expired.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Detecting Gateway Controller Failures in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

inactivity-timer

Syntax

```
inactivity-timer {
  inactivity-timeout {
    detect;
    maximum-inactivity-time {
      default 10-millisecond-units;
      maximum 10-millisecond-units;
      minimum 10-millisecond-units;
    }
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the inactivity timer package, which allows the BGF to use message inactivity to detect that its active gateway controller has failed.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Detecting Gateway Controller Failures in *Session Border Control Solutions Guide Using BGF and IMSG*

initial-average-ack-delay

Syntax initial-average-ack-delay *milliseconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-timers]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the value of the average acknowledgment delay (AAD) that the virtual BGF uses before the first AAD is measured. The AAD is explained in Annex D of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options ***milliseconds***—Assumed initial average delay.

Range: 0 through 65,535

Default: 4000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

interim-ah-scheme

Syntax

```
interim-ah-scheme {
  algorithm hmac-null;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* gateway-controller *gateway-controller-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Set up the BGF to use the interim AH scheme.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Using the BGF to Provide VoIP Solutions in a Next-Generation Network in *Session Border Control Solutions Guide Using BGF and IMSG*

ip-flow-stop-detection

Syntax ip-flow-stop-detection (regulated-notify | immediate-notify);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties application-data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure regulated or non-regulated (immediate) notification of media inactivity events.

Options **regulated-notify**—Activate regulated notification of media inactivity events.

immediate-notify—Activate non-regulated notification of media inactivity events.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring H.248 Notification Behavior to Prevent Excessive Media Inactivity Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

ipsec-transport-security-association

Syntax ipsec-transport-security-association *security-association-name*;

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the IPsec security association to be used for this virtual BGF.

Options ***security-association-name***—Name of the IPsec security association to be used for this virtual BGF. This is a security association that you configured at the **[edit services ipsec]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Security for BGF Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring IPsec to Protect H.248 Messages in Transport Mode in *Session Border Control Solutions Guide Using BGF and IMSG*

latch-deadlock-delay

Syntax latch-deadlock-delay *seconds*

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the time after which the virtual BGF begins checking for data packets on terminations that include a latch event.

Options ***seconds***—Time interval before checking for data packets.

Range: 0 through 3600

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

max-burst-size

See the following sections:

- max-burst-size (All Streams) on page 602
- max-burst-size (RTCP Streams) on page 603

max-burst-size (All Streams)

Syntax

```
max-burst-size {
  default bytes-per-second;
  maximum bytes-per-second;
  minimum bytes-per-second;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in Junos OS Release 9.2.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Configure the maximum burst size for gate streams of any protocol, including RTP.

Default The virtual BGF uses the default value of 1000 bytes if the Policy command in H.248 messages in ON and both of the following apply:

- The maximum burst size is not set in the H.248 message.
- There is no CLI configuration for maximum burst size.

Options

default *bytes-per-second*—Default maximum burst size.
Range: 20 through 4,294,967,295

maximum *bytes-per-second*—Maximum burst size.
Range: 20 through 4,294,967,295

minimum *bytes-per-second*—Minimum maximum burst size.
Range: 20 through 4,294,967,295

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

max-burst-size (RTCP Streams)

Syntax

```
max-burst-size {
  rtcp {
    (fixed-value bytes | percentage percentage);
  }
}
```

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in Junos OS Release 9.2.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Configure the maximum burst size for for RTP/RTCP gate streams. You can configure this rate as a fixed value or as a percentage of the RTP gate's rate.

Default The virtual BGF uses the default value of 100 percent of the RTP gate's maximum burst size if the Policy command in H.248 messages in ON and both of the following apply:

- The maximum burst size is not set in the H.248 message.
- There is no CLI configuration for maximum burst size.

Options **fixed-value** —Value entered is a fixed number of bytes per second.

bytes-per-second—maximum burst size.
Range: 20 through 4,294,967,295

percentage —Value entered is a percentage of the RTP gate's rate.

percentage—Maximum burst size as a percentage of the RTP gate's rate.
Range: 1 through 1000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

max-concurrent-calls

Syntax max-concurrent-calls *number-of-calls*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the maximum number of concurrent calls on the virtual BGF. If you configure multiple virtual BGFs for one service PIC or DPC, you can use this statement to achieve a fair distribution of resources between the virtual BGFs. For example, the Multiservices 500 PIC is capable of 10,000 concurrent calls, and you can divide this number between its associated virtual BGFs.

You can overbook concurrent calls to avoid resource idleness. The configured total of all virtual BGF maximum concurrent calls can be greater than the PIC or DPC limit. For example, bgf-1 and bgf-2 are connected to single PIC. If you configure 6000 maximum concurrent calls on bgf-1 and 8000 on bgf-2, bgf-1 can open up to 6000 concurrent calls, and bgf-2 can open up to 8000 concurrent calls. However, when the total number of calls reaches 10,000, neither of the virtual BGFs will be able to open a new context.

If the resources on the PIC are exhausted and no more calls are allowed, the virtual BGF sends an H.248 error message to the gateway controller in response to new call requests.



NOTE: **You must take the virtual BGF out of service before changing max-concurrent-calls and restart the pgcpd process after returning the virtual BGF to service.**

Options ***number-of-calls***—Maximum number of concurrent calls on the virtual BGF.

Range: 0 through 10,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

maximum-fuf-percentage

Syntax maximum-fuf-percentage *percentage*

Hierarchy Level [edit services pgcp gateway *gateway-name* fast-update-filters]

Release Information Statement introduced in Junos OS Release 9.1.

Description Along with the **maximum-terms** statement, limit the number of FUF terms installed on the Packet Forwarding Engine for a virtual BGF. This limit is the maximum value of the **maximum-terms** and **maximum-fuf-percentage** statements.

Options *percentage*—Maximum percentage of gates with FUF filters relative to all gates currently installed for the virtual BGF.

Range: 0 through 100

Default: 10

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Improving Performance While Collecting Gate Statistics in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

maximum-inactivity-time

Syntax

```
maximum-inactivity-time {
    default 10-millisecond-units;
    maximum 10-millisecond-units;
    minimum 10-millisecond-units;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties inactivity-timer inactivity-timeout]

Release Information Statement introduced in Junos OS Release 9.2.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Specify default, maximum, and minimum values for the maximum inactivity time. The default value is used if the gateway controller requests that the BGF detect the inactivity timeout event, but the gateway controller does not set a value for the maximum inactivity time. The maximum and minimum values are used to set limits for the maximum inactivity time set by the gateway controller. The BGF issues an error message if the value received from the gateway controller violates the configured minimum or maximum. If the BGF does not receive a message from the gateway controller before the maximum inactivity time expires, it sends a Notify message to the gateway controller. This timer resets each time the BGF receives a message from the gateway controller.

Options

default *10–millisecond-units*—Default value for the maximum inactivity time.

Range: 100 through 65,535 (10-millisecond units)

Default: 12,000

maximum *10–millisecond-units*—Maximum value for the maximum inactivity time.

Range: 100 through 65,535 (10-millisecond units)

Default: 12,000

minimum *10–millisecond-units*—Minimum value for the maximum inactivity time.

Range: 100 through 65,535 (10-millisecond units)

Default: 12,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Detecting Gateway Controller Failures in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

maximum-net-propagation-delay

Syntax maximum-net-propagation-delay *milliseconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-timers]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the assumed maximum network propagation delay time. This value is used to calculate the LONG-TIMER as explained in Annex D of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options ***milliseconds***—Duration of the maximum network propagation delay time.

Range: 0 through 65,535 milliseconds

Default: 40,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

maximum-synchronization-mismatches

Syntax maximum-synchronization-mismatches *number-of-mismatches*;

Hierarchy Level [edit services pgcp gateway *gateway-name* graceful-restart]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the maximum number of mismatches allowed during the synchronization procedure between the pgcpd process and the PIC or DPC. If the number of mismatches exceeds this number, the pgcpd process clears the state of the PIC or DPC and the state of the pgcpd process.

Options ***number-of-mismatches***—Maximum number of mismatches allowed during the synchronization procedure with the PIC or DPC.

Range: 0 through 3000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Synchronization Properties in Case of Routing Engine Failure in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

maximum-terms

Syntax maximum-terms *number-of-terms*

Hierarchy Level [edit services pgcp gateway *gateway-name* fast-update-filters]

Release Information Statement introduced in Junos OS Release 9.1.

Description Along with the **maximum-fuf-percentage** statement, limit the number of FUF terms installed on the Packet Forwarding Engine for a virtual BGF. This limit is the maximum value of the **maximum-terms** and **maximum-fuf-percentage** statements.

Options ***number-of-terms***—Maximum number of FUF terms installed for the virtual BGF.

Range: 0 through 20000

Default: 20000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Improving Performance While Collecting Gate Statistics in *Session Border Control Solutions Guide Using BGF and IMSG*

maximum-waiting-delay

Syntax maximum-waiting-delay *milliseconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-timers]

Release Information Statement introduced in Junos OS Release 8.4.

Description Define a maximum waiting delay (MWD) value. When the virtual BGF loses its connection to a gateway controller, it attempts to reconnect to the gateway controller. If the virtual BGF cannot reconnect to the gateway controller, it traverses its list of gateway controllers and attempts to connect to one of the gateway controllers. If the virtual BGF finishes traversing its list of gateway controllers, and has not connected to a gateway controller, the virtual BGF waits for a random value between 0 and MWD milliseconds before it begins another attempt to connect to a gateway controller. See section 9.2 of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options ***milliseconds***—Maximum time the virtual BGF waits before contacting a new gateway controller when the connection to the controlling gateway controller is lost.

Range: 1 through 36,000 milliseconds

Default: 3000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

media

Syntax

```
media {
  rtcp;
  rtp;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* monitor]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable Real-Time Control Protocol (RTCP) and Real-Time Transport Protocol (RTP) application-level gateways (ALGs) for media flows and monitor packets.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Monitoring RTP and RTCP Traffic in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mg-maximum-pdu-size

Syntax

```
mg-maximum-pdu-size {
  default bytes;
  maximum bytes;
  minimum bytes;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties segmentation]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MG maximum PDU size property of the segmentation package.

Options

default *bytes*—Default maximum size of messages that the gateway controller sends to the BGF.

Range: 512 through 65,507

maximum *bytes*—Maximum maximum size of messages that the gateway controller sends to the BGF.

Range: 512 through 65,507

minimum *bytes*—Minimum maximum size of messages that the gateway controller sends to the BGF.

Range: 512 through 65,507

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Segmentation Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mg-originated-pending-limit

Syntax

```
mg-originated-pending-limit {
  default number-of-messages;
  maximum number-of-messages;
  minimum number-of-messages;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties base-root]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MG originated pending limit property of the base root package.

Options

default ***number-of-messages***—Default number of transaction pending messages that the gateway controller can receive from the virtual BGF.

Range: 1 through 512

maximum ***number-of-messages***—Maximum number of transaction pending messages that the gateway controller can receive from the virtual BGF.

Range: 1 through 512

minimum ***number-of-messages***—Minimum number of transaction pending messages that the gateway controller can receive from the virtual BGF.

Range: 1 through 512

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mg-provisional-response-timer-value

Syntax

```
mg-provisional-response-timer-value {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties base-root]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MG provisional response timer property of the base root package.

Options

default ***milliseconds***—Default time within which the gateway controller waits for a pending response from the virtual BGF if a transaction cannot be completed.

Range: 500 through 3000

maximum ***milliseconds***—Maximum time within which the gateway controller waits for a pending response from the virtual BGF if a transaction cannot be completed.

Range: 500 through 3000

minimum ***milliseconds***—Minimum time within which the gateway controller waits for a pending response from the virtual BGF if a transaction cannot be completed.

Range: 500 through 3000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mg-segmentation-timer

Syntax

```
mg-segmentation-timer {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties segmentation]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MG segmentation timer value property of the segmentation package.

Options

default *milliseconds*—Default time within which the gateway controller waits to receive outstanding message segments from the virtual BGF after it receives the SegmentationCompleteToken.

Range: 500 through 30000

maximum *milliseconds*—Maximum time within which the gateway controller waits to receive outstanding message segments from the virtual BGF after it receives the SegmentationCompleteToken.

Range: 500 through 30000

minimum *milliseconds*—Minimum time within which the gateway controller waits to receive outstanding message segments from the virtual BGF after it receives the SegmentationCompleteToken.

Range: 500 through 30000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Segmentation Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mgc-maximum-pdu-size

Syntax

```
mgc-maximum-pdu-size {
  default bytes;
  maximum bytes;
  minimum bytes;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties segmentation]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, minimum, and maximum values for the MGC maximum PDU size property of the segmentation package.

Options

default *bytes*—Default maximum size of messages that the virtual BGF sends to the gateway controller.

Range: 512 through 65,507

maximum *bytes*—Maximum size of messages that the virtual BGF sends to the gateway controller.

Range: 512 through 65,507

minimum *bytes*—Minimum maximum size of messages that the virtual BGF sends to the gateway controller.

Range: 512 through 65,507

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Segmentation Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mgc-originated-pending-limit

Syntax

```
mgc-originated-pending-limit {
  default number-of-messages;
  maximum number-of-messages;
  minimum number-of-messages;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties base-root]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MGC originated pending limit property of the base root package.

Options

default ***number-of-messages***—Default number of transaction pending messages that the virtual BGF can receive from the gateway controller.

Range: 1 through 512

maximum ***number-of-messages***—Maximum number of transaction pending messages that the virtual BGF can receive from the gateway controller.

Range: 1 through 512

minimum ***number-of-messages***—Minimum number of transaction pending messages that the virtual BGF can receive from the gateway controller.

Range: 1 through 512

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mgc-provisional-response-timer-value

Syntax

```
mgc-provisional-response-timer-value {
  default milliseconds;
  maximum milliseconds;
  minimum milliseconds;
}
```

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name* h248-properties base-root]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MGC provisional response timer value property of the base root package.

Options

default *milliseconds*—Default time within which the virtual BGF waits for a pending response from the gateway controller if a transaction cannot be completed.

Range: 500 through 3000

maximum *milliseconds*—Maximum time within which the virtual BGF waits for a pending response from the gateway controller if a transaction cannot be completed.

Range: 500 through 3000

minimum *milliseconds*—Minimum time within which the virtual BGF waits for a pending response from the gateway controller if a transaction cannot be completed.

Range: 500 through 3000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

mgc-segmentation-timer

Syntax

```
mgc-segmentation-timer {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties segmentation]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the MGC segmentation timer value property of the segmentation package.

Options

default *milliseconds*—Default time within which the virtual BGF waits to receive outstanding message segments from the gateway controller after it receives the SegmentationCompleteToken.

Range: 500 through 30000

maximum *milliseconds*—Default time within which the virtual BGF waits to receive outstanding message segments from the gateway controller after it receives the SegmentationCompleteToken.

Range: 500 through 30000

minimum *milliseconds*—Default time within which the virtual BGF waits to receive outstanding message segments from the gateway controller after it receives the SegmentationCompleteToken.

Range: 500 through 30000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Segmentation Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

monitor

Syntax

```
monitor {
  media {
    rtcp;
    rtp;
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable Real-Time Control Protocol (RTCP) and Real-Time Transport Protocol (RTP) application-level gateways (ALGs) for media flows and monitor packets.

Options The statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Monitoring RTP and RTCP Traffic in *Session Border Control Solutions Guide Using BGF and IMSG*

network-operator-id

Syntax network-operator-id *network-operator-id*;

Hierarchy Level
[edit services pgcp session-mirroring delivery-function *deliver-function-name*],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring delivery-function *deliver-function-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the network operator ID. The BGF includes the network operator ID in the header of mirrored packets that it sends to the delivery function. It is used to identify the operator.

Options ***network-operator-id***—The network operator ID can be up to five characters.

Required Privilege Level
pgcp–session-mirroring—To view this statement in the configuration.
pgcp–session-mirroring-control—To add this statement to the configuration.

Related Documentation
- Configuring Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

no-dscp-bit-mirroring

Syntax no-dscp-bit-mirroring;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options encoding]

Release Information Statement introduced in Junos OS Release 9.3.

Description Disable mirroring of DSCP bits.

Default DSCP bits are mirrored by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

no-rtcp-check

Syntax no-rtcp-check;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.5.

Description Prevent checking for inactivity on RTCP streams.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

normal-mg-execution-time

Syntax

```
normal-mg-execution-time {
  default milliseconds;
  maximum milliseconds;
  minimum milliseconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties base-root]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the normal MG execution time property of the base root package.

Options

default *milliseconds*—Default interval within which the gateway controller waits for a response to transactions from the virtual BGF.

Range: 500 through 29000

maximum *milliseconds*—Maximum interval within which the gateway controller waits for a response to transactions from the virtual BGF.

Range: 500 through 29000

minimum *milliseconds*—Minimum interval within which the gateway controller waits for a response to transactions from the virtual BGF.

Range: 500 through 29000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

normal-mgc-execution-time

Syntax

```
normal-mgc-execution-time {
  default milliseconds;
  maximum milliseconds;
  minimum milliseconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties base-root]

Release Information Statement introduced in Junos OS Release 8.5.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Set default, maximum, and minimum values for the normal MGC execution time property of the base root package.

Options

default *milliseconds*—Default interval within which the virtual BGF waits for a response to transactions from the gateway controller.

Range: 500 through 29000

maximum *milliseconds*—Maximum interval within which the virtual BGF waits for a response to transactions from the gateway controller.

Range: 500 through 29000

minimum *milliseconds*—Minimum interval within which the virtual BGF waits for a response to transactions from the gateway controller.

Range: 500 through 29000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Base Root Properties in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

notification-behavior

Syntax

```
notification-behavior {
  notification-regulation default (once | 0 – 100);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the default frequency for regulated media inactivity notifications sent by the BGF.

Options The statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

notification-rate-limit

Syntax notification-rate-limit *rate*;

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the maximum notifications sent per second by the PIC or DPC.

Options ***rate***—Maximum number of notifications per second the PIC or DPC sends to a gateway controller.

Range: 10 through 10,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

notification-regulation

Syntax notification-regulation (once | 0 – 100);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties notification-behavior]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the default frequency for sending media inactivity notifications for regulated events.

Options

once—Send only one media inactivity notification for a regulated event to the gateway controller.

0 – 100—The percentage of media inactivity notifications for regulated events to send to the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

overload-control

Syntax

```
overload-control {
  queue-limit-percentage percentage;
  reject-all-commands-threshold percentage;
  reject-new-calls-threshold percentage;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.3.
reject-all-commands-threshold and **reject-new-calls-threshold** options introduced in Junos OS Release 9.5.

Description Configure the BGF to send overload messages to the gateway controller based on the status of its work queue. The overload messages cause the gateway controller to lower the rate at which it admits packets for processing.

Options The statement is described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Overload Control for Voice Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

peak-data-rate

See the following sections:

- peak-data-rate (All Streams) on page 624
- peak-data-rate (RTCP) on page 625

peak-data-rate (All Streams)

Syntax

```
peak-data-rate {
  default bytes-per-second;
  maximum bytes-per-second;
  minimum bytes-per-second;
}
```

Hierarchy Level [edit services services pgcp gateway *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in Junos OS Release 9.2.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Configure the peak data rate for gate streams of any protocol.

Default The BGF uses the default value of 10,000 bytes per second if the Policy command in H.248 messages is ON and both of the following apply:

- The peak data rate is not set in the H.248 message.
- There is no CLI configuration for peak data rate.

Options

default *bytes-per-second*—Default peak data rate.
Range: 125 through 4,294,967,295

maximum *bytes-per-second*—Maximum peak data rate.
Range: 125 through 4,294,967,295

minimum *bytes-per-second*—Minimum peak data rate.
Range: 125 through 4,294,967,295

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

peak-data-rate (RTCP)

Syntax

```
peak-data-rate {
  rtcp (fixed-value bytes | percentage percentage);
}
```

Hierarchy Level [edit services services pgcp gateway *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the peak data rate for RTP/RTCP gate streams. You can configure this rate as a fixed value or as a percentage of the RTP gate's rate.

Default The BGF uses the default value of 5 percent of the RTP gate's rate if the Policy command in H.248 messages in ON and both of the following apply:

- The peak data rate is not set in the H.248 message.
- There is no CLI configuration for peak data rate.

Options **fixed-value** —Value entered is a fixed number of bits per second.

bytes-per-second—Peak data rate.
Range: 125 through 4,294,967,295

percentage —Value entered is a percentage of the RTP gate's rate.

percentage—Value entered is a percentage of the RTP's gate rate.
Range: 0 through 1000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

platform

Syntax

```
platform {
  device interface-name;
  routing-engine;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the platform on which the virtual BGF runs. The virtual BGF can run on the Routing Engine or on a Multiservices PIC or MS-DPC. The Multiservices 500 PIC is not supported for virtual BGFs. If you are using high availability, you can configure the virtual BGF to run on a virtual redundant Multiservices PIC (rms) interface

Options

device—Causes the virtual BGF to run on a Multiservices PIC, MS-DPC, or an **rms** interface.

interface-name—Name of the service interface. If you are using high availability, enter the **rms** interface number.

routing-engine—Causes the virtual BGF to run on the Routing Engine. By default, virtual BGFs run on the Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Virtual BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

profile-name

Syntax profile-name *profile-name*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options h248-profile]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the H.248 profile name that the BGF declares in initial registration ServiceChange requests.

Options ***profile-name***—Name of the H.248 profile.

Syntax: 1-64 bytes in length. The name must start with a letter. Allowed characters are [a-zA-Z0-9_]

Default: ETSI_BGF, which is the ETSI Ia standard (ETSI ES 283 018 v1.1.4).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the H.248 Profile in *Session Border Control Solutions Guide Using BGF and IMSG*

profile-version

Syntax profile-version *version-number*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options h248-profile]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the H.248 profile version that the BGF declares in initial registration ServiceChange requests.

Options ***version-number***—H.248 profile version number.

Range: 1 through 99

Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the H.248 Profile in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

queue-limit-percentage

Syntax queue-limit-percentage *percentage*;

Hierarchy Level [edit services pgcp gateway *gateway-name* overload-control]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the queue limit percentage (percentage of the maximum work queue size currently in use) that indicates overload. When the gateway controller activates overload control, the BGF generates an overload notification for each transaction on a gate that contains an ADD if the work queue utilization has reach this limit. When 100 percent of the queue is in use, transactions are dropped with error 510 (insufficient resources).

Options ***percentage***—Percentage of the overload control work queue in use that triggers creation of an overload notification.

Range: 25 through 100

Default: 80

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Overload Control for Voice Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

reconnect

Syntax reconnect (disconnected-900 | restart-902);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications disconnect]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Registration Request ServiceChange commands when it attempts to reregister with the gateway controller or register with a new gateway controller after the control association is disconnected.

Default If you do not specify an option, the virtual BGF includes DC/900 (**disconnected-900**).

Options

disconnected-900—Service restored. The virtual BGF is registering with the last controlling gateway controller following a disconnection of the virtual BGF and gateway controller.

restart-902—Warm boot. The virtual BGF is transitioning to In-Service, and the previously installed state is retained.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

reject-all-commands-threshold

Syntax reject-all-commands-threshold *percentage*;

Hierarchy Level [edit services (PGCP) pgcp gateway (Services PGCP) *gateway-name* overload-control]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the maximum percentage of the work queue that can be in use before the virtual BGF rejects all non-emergency transactions other than SUBTRACT transactions.

Options

percentage—Percentage of work queue space used that serves as a threshold for overload control.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Overload Control for Voice Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

reject-new-calls-threshold

Syntax reject-new-calls-threshold *percentage*;

Hierarchy Level [edit services pgcp gateway *gateway-name* overload-control]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the maximum percentage of the work queue that can be in use before the virtual BGF rejects all non-emergency ADD transactions.

Options ***percentage***—Percentage of work queue space used that serves as a threshold for overload control.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Overload Control for Voice Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

report-service-change

Syntax
```
report-service-change {
  service-change-type (forced-906 | forced-910);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Change the service state of inactive terminations to prevent continued sending of inactivity notifications.

Options The statement is described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

request-timestamp

Syntax request-timestamp (requested | suppressed | autonomous);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248–properties event-timestamp]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify whether time stamp information is made available to the gateway controller or is suppressed.

Options **requested**—Enables gateway controller access to time stamp information for notifications.

suppressed—Disables gateway controller access to time stamp information for notifications.

autonomous—Equivalent to **suppressed**.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring H.248 Notification Behavior to Prevent Excessive Media Inactivity Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*

routing-instance

Syntax

```
routing-instance instance-name {
  service-interface interface-name.unit-number;
}
```

Hierarchy Level [edit services pgcp virtual-interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.4.
service-interface option introduced in Junos OS Release 9.3.

Description Map the virtual router interface to a VPN routing and forwarding (VRF) routing instance configured on the router.

Options *instance-name*—Name of a routing instance that has been configured at the **[edit routing-instance]** hierarchy level.

The remainder of the statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

rtcp

Syntax rtcp;

Hierarchy Level [edit services pgcp gateway *gateway-name* monitor media]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable Real-Time Control Protocol (RTCP) application-level gateway (ALG) on media flows created when the gateway controller installs media gates on the virtual BGF.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

rtp

Syntax rtp;

Hierarchy Level [edit services pgcp gateway *gateway-name* monitor media]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable Real-Time Transport Protocol (RTP) application-level gateway (ALG) on media flows created when the gateway controller installs media gates on the virtual BGF.

Required Privilege Level interface-level—To view this statement in the configuration.
interface-level—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

rule

Syntax

```
rule rule-name {
  gateway gateway-name;
  nat-pool [ pool-names ];
}
```

Hierarchy Level [edit services pgcp],
[edit services service-set *service-set-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Specify the rule that the router uses when it applies the NAT pool.

Options ***rule-name***—Identifier for the rule.

pool-names—Names of one or more NAT pools to be used by the rule.

Syntax: To specify a list of NAT pools, enclose the NAT pool names in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

rule-set

Syntax

```
rule-set rule-set-name {
  [rule rule-name]
}
```

Hierarchy Level [edit services pgcp],
[edit services service-set *service-set-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that make up this rule set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sbc-utils (Services PGCP)

Syntax

```
sbc-utils {
    common trace-level;
    configuration trace-level;
    device-monitor trace-level;
    ipc trace-level;
    memory-management trace-level;
    message trace-level;
    minimum trace-level;
    user-interface trace-level;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure trace options for the Signaling Border Controller (SBC) utilities component of the virtual BGF.

Default warning

Options

minimum *trace-level*—Minimum trace level for all **sbc-util** messages.

common *trace-level*—Trace level for the common component of SBC utilities.

configuration *trace-level*—Trace level for the configuration component of SBC utilities.

device-monitor *trace-level*—Trace level for the device monitor component of SBC utilities.

ipc *trace-level*—Trace level for the IPC component of SBC utilities.

memory-management *trace-level*—Trace level for the memory management component of SBC utilities.

message *trace-level*—Trace level for the message component of SBC utilities.

user-interface *trace-level*—Trace level for the user interface component of SBC utilities.

trace-level—Trace level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

segmentation

Syntax

```
segmentation {
  mg-segmentation-timer {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
  }
  mgc-segmentation-timer {
    default milliseconds;
    maximum milliseconds;
    minimum milliseconds;
  }
  mg-maximum-pdu-size {
    default bytes;
    maximum bytes;
    minimum bytes;
  }
  mgc-maximum-pdu-size {
    default bytes;
    maximum bytes;
    minimum bytes;
  }
}
```

Hierarchy Level
[edit services pgcp gateway *gateway-name* h248-properties]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Configure default values for properties in the segmentation package defined in Annex E of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options
The statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

send-notification-on-delay

Syntax send-notification-on-delay;

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Send an inactivity notification immediately when no media packets are detected during a delay period that precedes checking for media inactivity. By default, notifications are sent after both the delay period and an additional period of inactivity have elapsed without any media packets being detected.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- inactivity-delay on page 597
- latch-deadlock-delay on page 601
- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-change

Syntax

```
service-change {
  control-association-indications {
    disconnect {
      controller-failure (failover-909 | restart-902);
      reconnect (disconnected-900 | restart-902);
    }
    down {
      administrative (forced-905 | forced-908 | none);
      failure (forced-904 | forced-908 | none);
      graceful (graceful-905 | none);
    }
    up {
      cancel-graceful (none | restart-918);
      failover-cold (failover-920 | restart-901);
      failover-warm (failover-919 | restart-902);
    }
  }
  virtual-interface-indications {
    virtual-interface-down {
      administrative (forced-905 | forced-906 | none);
      graceful (graceful-905 | none);
    }
    virtual-interface-up {
      cancel-graceful (none | restart-918);
      warm (none | restart-900);
    }
  }
  context-indications {
    state-loss (forced-910 | forced-915 | none);
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in ServiceChange commands that it sends to the gateway controller when the state of a control association, virtual interface, or context changes.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-change-type

Syntax service-change-type (forced-906 | forced-910)

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection report-service-change]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason used in changing the service state of the termination to active in order to curtail sending of inactivity messages.

Options

forced-906—Service is terminated using a forced termination method with reason code 906 (loss of lower layer connectivity).

forced-910—Service is terminated using a forced termination with reason code 910 (media capability failure).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

service-interface

Syntax service-interface *interface-name.unit-number*;

Hierarchy Level [edit services pgcp virtual-interface *virtual-interface-name* routing-instance]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the logical service interface. The NAT routes point to this service interface. This service interface must match the service interface configured in the routing instance.

Options ***interface-name.unit-number***—Name and logical interface number of the service interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Service Set for Redundant Services PICS in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-state

See the following sections:

- service-state (Virtual BGF) on page 639
- service-state (Virtual Interface) on page 640

service-state (Virtual BGF)

Syntax service-state (in-service | out-of-service-forced | out-of-service-graceful);

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Set the service state of the virtual BGF.

Options

in-service—The virtual BGF is operational and available for traffic. When the virtual BGF is in service, it attempts to connect to the gateway controller and accepts all PGCP commands from the gateway controller.

out-of-service-forced—Force the virtual BGF out of service. When the virtual BGF is forced out of service, it immediately removes all gates and disconnects from the gateway controller. The virtual BGF does not attempt to establish a new connection.

out-of-service-graceful—Cause the virtual BGF to go out of service by entering a draining mode and waiting for all terminations to be subtracted before going out of service. During the draining, the BGF accepts only subtract commands from the gateway controller.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

service-state (Virtual Interface)

Syntax service-state (in-service | out-of-service-forced | out-of-service-graceful);

Hierarchy Level [edit services pgcp virtual-interface]

Release Information Statement introduced in Junos OS Release 9.0.

Description Set the service state of the virtual interface.

Options

in-service—Virtual interface is operational and available for traffic. When the virtual interface is in service, it is connected to the physical interface and accepts all Voice calls. This is the default.

out-of-service-forced—Force the virtual interface out of service. When the virtual interface is forced out of service, it immediately removes all calls and disconnects from the physical interface. The virtual interface does not attempt to establish a new connection.

out-of-service-graceful—Cause the virtual interface goes out of service by entering a draining mode and waiting for all terminations to be subtracted before going out of service. During the draining, the virtual interface accepts only subtract commands from the gateway controller.

Required Privilege Level interface-level—To view this statement in the configuration.
interface-level—To add this statement to the configuration.

services (PGCP)

Syntax services pgcp { ... }

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 8.4.

Description Define service rules to be applied to traffic.

Options **pgcp**—Identifier for the PGCP set of rules statements.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- BGF VoIP Solution Overview and IMSG Session Border Control Solution Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

session-mirroring

Syntax

```
session-mirroring {
  delivery-function delivery-function-name {
    destination-address destination-address;
    destination-port destination-port;
    network-operator-id network-operator-id;
    source-address source-address;
    source-port (Mirrored BGF Packets) source-port;
  }
  disable-session-mirroring;
}
```

Hierarchy Level

[edit services pgcp];
[edit services pgcp gateway (Services PGCP) *gateway-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure the session mirroring feature.

Options

The statements are explained separately.

Required Privilege Level

pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation

- Session Mirroring Overview and Configuring Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

source-address

Syntax

source-address *source-address*;

Hierarchy Level

[edit services pgcp session-mirroring delivery-function *deliver-function-name*],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring delivery-function *deliver-function-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure the source address that is applied to mirrored packets.

Options

source-address—Address of the interface on which the BGF sends session-mirroring data to the delivery function.

Required Privilege Level

pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation

- Session Mirroring Overview and Configuring Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

source-port (Mirrored BGF Packets)

Syntax source-port *source-port*;

Hierarchy Level [edit services pgcp session-mirroring delivery-function *deliver-function-name*],
[edit services pgcp gateway (Services PGCP) *gateway-name* session-mirroring delivery-function *deliver-function-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the source port applied to the mirrored packets.

Options ***source-port***—Port on which the BGF sends session-mirroring data to the delivery function.

Range: 1 through 65,535

Required Privilege Level pgcp-session-mirroring—To view this statement in the configuration.
pgcp-session-mirroring-control—To add this statement to the configuration.

Related Documentation

- Session Mirroring Overview and Configuring Session Mirroring in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

state-loss

Syntax state-loss (forced-910 | forced-915 | none);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change context-indications]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Service-Interruption ServiceChange commands that it sends to the gateway controller after a state loss on a specific context.

Default If you do not specify an option, the virtual BGF includes FO/915 (**forced-915**).

Options

forced-910—State loss because of a media failure. A mismatch between the pgcpd process and the Multiservices PIC or DPC states was detected on one or more of the context's gates.

forced-915—State loss. A mismatch between the pgcpd process and the Multiservices PIC or DPC states was detected on one or more of the context's gates.

none—Virtual BGF does not send a ServiceChange command to the gateway controller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Session Mirroring Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

stop-detection-on-drop

Syntax stop-detection-on-drop;

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the BGF to stop inactivity detection when a gate action is set to drop. When the call is resumed, the BGF starts the delay time and resumes data inactivity detection.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sustained-data-rate

See the following sections:

- sustained-data-rate (All Streams) on page 644
- sustained-data-rate (RTCP Streams) on page 645

sustained-data-rate (All Streams)

Syntax

```
sustained-data-rate {
  default bytes-per-second;
  maximum bytes-per-second;
  minimum bytes-per-second;
  rtcp-include;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in Junos OS Release 9.2.
maximum and **minimum** options introduced in Junos OS Release 9.5.

Description Configure the sustained data rate for streams of any protocol, including RTP.

Default The BGF uses the default value of 10,000 bytes per second if the Policy command in H.248 messages in ON and both of the following apply:

- The sustained data rate is not set in the H.248 message.
- There is no CLI configuration for sustained data rate.

Options

default *bytes-per-second*—Default value for sustained data rate.
Range: 125 through 4,294,967,295

maximum *bytes-per-second*—Maximum value for sustained data rate.
Range: 125 through 4,294,967,295

minimum *bytes-per-second*—Minimum value for sustained data rate.
Range: 125 through 4,294,967,295

rtcp-include—Include rtcp bandwidth in the sustained data rate.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rate-Limiting for VoIP Traffic Overview and Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sustained-data-rate (RTCP Streams)

Syntax

```
sustained-data-rate {
  rtcp (fixed-value bytes-per-second | percentage percentage);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the sustained data rate for RTP/RTCP gate streams. You can configure this rate as a fixed value or as a percentage of RTP's sustained data rate.

Default The virtual BGF uses the default value of 5 percent of the RTP gates's rate if the Policy command in H.248 messages in ON and both of the following apply:

- The sustained data rate is not set in the H.248 message.
- There is no CLI configuration for sustained data rate.

Options **fixed-value** —Value entered is a fixed number of bits per second.

bytes-per-second—Sustained data rate.

Range: 125 through 4,294,967,295

percentage *bytes-per-second*—Value entered is a percentage of the RTP's gate rate.

percentage—Value entered is a percentage of the RTP's gate rate.

Range: 0 through 1000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rate-Limiting for VoIP Traffic Overview and Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

timerx

Syntax timerx *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties hanging-termination-detection]

Release Information Statement introduced in Junos OS Release 9.2.

Description Activate and configure hanging termination detection. Setting this timer to a value other than zero (0) activates hanging termination detection. If no messages are exchanged between the BGF and the gateway controller for a termination before this time expires, the BGF sends a notification to the gateway controller. The timer resets when the BGF and the gateway controller exchange a message for the termination. The timer value that you set is the default value, and can be overridden by H.248 messages sent from the gateway controller.

Your configuration takes effect on new and modified terminations.

Options ***seconds***—Number of seconds between the last message exchanged for this termination and when the BGF sends a notification to the gateway controller. Setting the timer to zero (0) deactivates hanging termination detection.

Range: 0 through 2,147,480

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Detecting Hanging Terminations in *Session Border Control Solutions Guide Using BGF and IMSG*

tmax-retransmission-delay

Syntax tmax-retransmission-delay *milliseconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-timers]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the maximum time that a transaction can be kept alive. T-Max is explained in Annex D of *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*.

Options ***milliseconds***—Duration of the delay before the BGF considers the gateway controller to be down.

Range: 0 through 65,535

Default: 25000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Services PGCP)

Syntax

```
traceoptions {
  file <filename filename> <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag {
    bgf-core {
      common trace-level;
      default trace-level;
      firewall trace-level;
      gate-logic trace-level;
      pic-broker trace-level;
      policy trace-level;
      statistics trace-level;
    }
    default trace-level;
    h248-stack {
      control-association trace-level;
      default trace-level;
      messages;
      media-gateway trace-level;
    }
    sbc-utils {
      common trace-level;
      configuration trace-level;
      default trace-level;
      device-monitor trace-level;
      ipc trace-level;
      memory-management trace-level;
      messaging trace-level;
      user-interface trace-level;
    }
  }
}
```

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in Junos OS Release 8.4.
Statement extensively revised in Junos OS Release 9.5.

Description Configure PGCP tracing operations. The messages are output to **/var/log/pgcpd**.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BGF Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

traffic-management

Syntax

```
traffic-management {
  max-burst-size {
    default bytes-per-second;
    maximum bytes-per-second;
    minimum bytes-per-second;
    rtcp {
      (fixed-value bytes-per-second | percentage percentage);
    }
  }
  peak-data-rate {
    default bytes-per-second;
    maximum bytes-per-second;
    minimum bytes-per-second;
    rtcp {
      (fixed-value bytes-per-second | percentage percentage);
    }
  }
  sustained-data-rate {
    default bytes-per-second;
    maximum bytes-per-second;
    minimum bytes-per-second;
    rtcp {
      (fixed-value bytes-per-second | percentage percentage);
    }
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure traffic management of the gate stream and the RTCP stream. The parameters for the RTCP stream take effect only when the gate is an RTP/RTCP gate.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rate-Limiting for VoIP Traffic Overview and Configuring Rate Limiting for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

up

Syntax

```
up {
  cancel-graceful (none | restart-918);
  failover-cold (failover-920 | restart-901);
  failover-warm (failover-919 | restart-902);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change control-association-indications]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Notification Messages or Registration commands in ServiceChange commands when a control association transitions to In-Service.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Control Associations in *Session Border Control Solutions Guide Using BGF and IMSG*

use-lower-case

Syntax use-lower-case;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure upper-case encoding for H.248 messages.

Default By default H.248 messages are encoded in upper case.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

use-wildcard-response

Syntax use-wildcard-response;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change]

Release Information Statement introduced in Junos OS Release 9.3.

Description Enable the virtual BGF to issue service change commands as wildcard-response commands, which trigger a short response from the gateway controller. If you do not enable the use of wildcard responses for service change commands, the gateway controller will generate an individual response for every termination that matches the service change command.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Wildcards for ServiceChange Notifications in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

virtual-interface

Syntax

```
virtual-interface number {
  nat-pool [ pool-names ];
  routing-instance instance-name {
    service-interface interface-name.unit-number;
  }
  service-state (in-service | out-of-service-forced | out-of-service-graceful);
}
```

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in Junos OS Release 8.4.
service-state option introduced in Junos OS Release 9.0.
service-interface option introduced in Junos OS Release 9.3.

Description Configure a virtual interface for the BGF.

Options ***number***—Identifier for the interface.

Range: 0 through 1023

pool-names—Names of one or more NAT pools to be used by the virtual interface.

Syntax: To specify a list of NAT pools, enclose the NAT pool names in brackets.

The remainder of the statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Virtual Interfaces with the BGF Overview and Configuring Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

virtual-interface-down

Syntax

```
virtual-interface-down {
  administrative (forced-905 | forced-906 | none);
  graceful (graceful-905 | none);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change virtual-interface-indications]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in ServiceChange commands that it sends to the gateway controller when the state of the virtual interface changes to Out-of-Service.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

virtual-interface-indications

Syntax

```
virtual-interface-indications {
  virtual-interface-down {
    administrative (forced-905 | forced-906 | none);
    graceful (graceful-905 | none);
  }
  virtual-interface-up {
    cancel-graceful (none | restart-918);
    warm(none | restart-900);
  }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in ServiceChange commands that it sends to the gateway controller when the state of the virtual interface changes.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

virtual-interface-up

Syntax

```
virtual-interface-up {
  cancel-graceful (none | restart-918);
  warm (none | restart-900);
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change virtual-interface-indications]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the ServiceChange command that the virtual BGF sends to the gateway controller when the state of the virtual interface changes to In-Service.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

warm

Syntax warm (none | restart-900);

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options service-change virtual-interface-indications virtual-interface-up]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the method and reason that the virtual BGF includes in Service-Restoration ServiceChange commands that it sends to the gateway controller when a virtual interface transitions to In-Service.

Default If you do not specify an option, the virtual BGF includes RS/900 (**restart-900**).

Options **none**—Virtual BGF does not send a ServiceChange command.

restart-900—Service restored. The virtual interface has become In-Service and is in the Forwarding state.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method and Reason in ServiceChange Commands for Virtual Interfaces in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 30

Service Interface Pools Configuration Guidelines

To configure service interface pools, include the **service-interface-pools** statement at the **[edit services]** hierarchy level:

```
[edit services]
service-interface-pools {
  pool pool-name {
    interface interface-name.unit-number;
  }
}
```

This chapter discusses the following topics that provide information about configuring service interface pools:

- Configuring Service Interface Pools on page 655

Configuring Service Interface Pools

To configure a service interface pool, include the following statements at the **[edit services service-interface-pools]** hierarchy level:

```
[edit services service-interface-pools]
pool pool-name {
  interface interface-name.unit-number;
}
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 31

Summary of Service Interface Pools Statements

The following sections explain each of the service interface pools statements. The statements are organized alphabetically.

interface

Syntax interface *interface-name.unit-number*;

Hierarchy Level [edit services service-interface-pools pool *pool-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Add logical service interfaces to the pool of service interfaces.

Options ***interface-name.unit-number***—Name and logical unit number of the service interface.

- All interfaces in a pool must belong to the same service PIC or DPC.
- All interfaces assigned to the same service must be in the same pool.
- Logical interfaces cannot be in more than one pool.
- All interfaces must have either **family inet** or **family inet6** configured.
- Logical unit 0 cannot be configured in a service interface pool.
- You can configure up to 1000 logical interfaces in a service interface pool.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

pool

Syntax

```
pool pool-name {
  interface interface-name.unit-number;
}
```

Hierarchy Level [edit services service-interface-pools]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure a service interface pool for VPN aggregation for the BGF feature.

Options *pool-name*—Name of the service interface pool.

The remaining options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

service-interface-pools

Syntax

```
service-interface-pools {
  pool pool-name {
    interface interface-name.unit-number;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure service interface pools used for VPN aggregation.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 32

Border Signaling Gateway Configuration Guidelines

To configure border signaling gateways (BSG), include the **border-signaling-gateway** statement at the **[edit services]** hierarchy level:

```
[edit services]
border-signaling-gateway {
  gateway gateway-name {
    admission-control admission-control-profile {
      dialogs {
        maximum-concurrent number;
        committed-attempts-rate dialogs-per-second;
        committed-burst-size number-of-dialogs;
      }
      transactions {
        maximum-concurrent number;
        committed-attempts-rate transactions-per-second;
        committed-burst-size number-of-transactions;
      }
    }
    embedded-spdf {
      service-class service-class-name {
        term term-name {
          from {
            media-type (any-media | audio | video);
          }
          then {
            committed-burst-size bytes;
            committed-information-rate bytes-per-second;
            dscp (alias | do-not-change | dscp-value);
            reject;
          }
        }
      }
    }
    name-resolution-cache {
      accelerations {
        initiate-alternative-queries;
        initiate-next-queries;
        no-refresh-before-ttl-expiry;
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
    blacklist-period seconds;
    maximum-records-in-cache number;
    maximum-time-in-cache (unlimited | seconds);
  }
  service-interface name;
  service-point service-point-name {
    default-media-realm
    service-interface interface-name.unit-number;
    service-point-type service-point-type;
    service-policies {
      new-call-usage-input-policies [policy-and-policy-set-names];
      new-call-usage-output-policies [policy-and-policy-set-names];
      new-transaction-input-policies [policy-and-policy-set-names];
      new-transaction-output-policies[policy-and-policy-set-names];
    }
    transport-details <port port-number> <ip-address ip-address> <tcp> <udp>;
  }
  sip {
    message-manipulation-rules {
      manipulation-rule rule-name {
        actions {
          sip-header header-field-name {
            field-value {
              modify-regular-expression regular-expression with field-value;
              add field-value;
              add-missing field-value;
              add-overwrite field-value;
              remove-regular-expression regular-expression;
              remove-all;
              reject-regular-expression regular-expression;
            }
          }
          request-uri request-uri {
            field-value {
              modify-regular-expression regular-expression with field-value;
            }
          }
        }
      }
    }
    new-call-usage-policy policy-name {
    term term-name {
      from {
        contact [ contact-fields ];
        method {
          method-invite;
        }
        request-uri [ uri-fields ];
        source-address [ ip-addresses ];
      }
      then {
        media-policy {
          data-inactivity-detection (Services PGCP) {
            inactivity-duration (Services Border Signaling Gateway) seconds;
          }
          no-anchoring;
```


Copyright © 2013, Juniper Networks, Inc.

```
        service-class service-class-name;
      }
      trace;
    }
  }
  new-call-usage-policy-set policy-set-name {
    policy-name [ policy-names ];
  }
  new-transaction-policy policy-name {
    term term-name {
      from {
        contact {
          registration-state [ registered | not-registered ];
          regular-expression [ regular-expression ];
          uri-hiding [ hidden-uri | not-hidden-uri ];
        }
        method {
          method-invite;
          method-message;
          method-options;
          method-publish;
          method-refer;
          method-register;
          method-subscribe;
        }
        request-uri {
          registration-state [ registered | not-registered ];
          regular-expression [ regular-expression ];
          uri-hiding [ hidden-uri | not-hidden-uri ];
        }
      source-address [ ip-addresses ];
      }
      then {
        (accept | reject);
        admission-control admission-control-profile;
        message-manipulation {
          forward-manipulation {
            manipulation-rule-name;
          }
          reverse-manipulation {
            manipulation-rule-name;
          }
        }
        on-3xx-response{
          recursion-limit number;
          }
        }
        route (Services) {
          egress-service-point service-point-name;
          next-hop (Services) (request-uri | address ipv4-address | <port port-number>
            <transport-protocol (udp | tcp)>);
          server-cluster cluster-name;
        }
        signaling-realm signaling-realm;
        trace;
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  new-transaction-policy-set policy-set-name {
    policy-name [ policy-names ];
  }
  routing-destinations {
    availability-check-profiles {
      profile-name;
      keepalive-interval {
        available-server seconds;
        unavailable-server seconds;
      }
      keepalive-method sip-options;
      keepalive-strategy (do-not-send <blackout-period seconds> | send-always <
        failures-before-unavailable number> < successes-before-available number |
        send-when-unavailable < successes-before-available number);
      transaction-timeout seconds;
    }
    clusters [
      cluster-name;
      server server-name {
        priority priority-level;
        weight weight-level;
      }
    }
    default-availability-check-profile profile-name;
    }
    servers {
      server-name {
      address ip4-address <port port-number> <transport (udp | tcp)>;
      admission-control profile-name;
      availability-check-profile profile-name;
      service-point service-point-name;
    }
  timers {
    inactive-callseconds;
    timer-c seconds;
  }
}
traceoptions {
  file {
    filename filename;
    files number-of-files;
    match regular-expression;
    size maximum-trace-file-size;
  }
  flag {
    datastore {
      data trace-level;
      db trace-level;
      handle trace-level;
      minimum trace-level;
    }
    framework {
      action trace-level;
      event trace-level;
```


Copyright © 2013, Juniper Networks, Inc.

```
            executor trace-level;
            freezer trace-level;
            minimum trace-level;
            memory-pool trace-level;
          }
          minimum trace-level;
          sbc-utils {
            common trace-level;
            configuration trace-level;
            device-monitor trace-level;
            ipc trace-level;
            memory-management trace-level;
            message trace-level;
            minimum trace-level;
            user-interface trace-level;
          }
          session-trace trace-level;
          signaling {
            b2b trace-level;
            b2b-wrapper trace-level;
            minimum trace-level;
            policy trace-level;
            sip-stack-wrapper trace-level;
            topology-hiding trace-level;
            ua trace-level;
          }
          sip-stack {
            dev-logging;
            event-tracing;
            ips-tracing;
            pd-log-detail (full | summary);
            pd-log-level (audit | exception | problem);
            per-tracing;
            verbose-logging;
          }
        }
      }
    }
  }
```

For information about configuring the border signaling gateway, see the *Session Border Control Solutions Guide Using BGF and IMSG*.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 33

Summary of Border Signaling Gateway Configuration Statements

The following sections explain each of the border signaling gateway statements. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

actions

Syntax

```
actions {
  sip-header header-field-name {
    field-value {
      add field-value;
      add-missing field-value;
      add-overwrite field-value;
      modify-regular-expression regular-expression with field-value;
      reject-regular-expression regular-expression;
      remove-all;
      remove-regular-expression regular-expression;
    }
  }
  request-uri request-uri {
    field-value {
      modify-regular-expression regular-expression with field-value;
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip message-manipulation-rules manipulation-rule *rule-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the actions for your manipulation rule. You can have up to 50 actions in a manipulation rule.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Manipulation of Headers and Request URIs in SIP Messages in the *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

accelerations

Syntax

```
accelerations {
  initiate-alternative-queries;
  initiate-next-queries;
  no-refresh-before-ttl-expiry;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* name-resolution-cache]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the method, if any, that the BSG uses to accelerate the process of DNS name resolution for SIP servers.

Options

initiate-alternative-queries—If this flag is on, the BSG initiates a name authority pointer (NAPTR) query, both a TCP and UDP service record (SRV) query, and an address record (A) query in parallel for each new SIP URI that it receives in a new transaction. This flag saves time if the NAPTR query fails.

Default: off

initiate-next-queries—If this flag is on, the BSG sends A record queries to all SIP servers returned in the SRV response instead of querying the first A record in the SRV response.

Default: off

no-refresh-before-ttl-expiry—if this flag is on, the BSG removes the SIP server from the cache when the TTL expires. If this flag is off, the BSG re-queries the A record and if the query is resolved, the BSG refreshes the TTL. As a result, the entry is not removed from the cache.

Default: off

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS Resolution for Locating SIP Servers in the *Multiplay Solutions Guide*

Copyright © 2013, Juniper Networks, Inc.

admission-control

See the following sections:

- admission-control (Border Signaling Gateway) on page 668
- admission-control (New Transaction Policy) on page 669

admission-control (Border Signaling Gateway)

Syntax

```
admission-control admission-control-profile {
  dialogs {
    maximum-concurrent number;
    committed-attempts-rate dialogs-per-second;
    committed-burst-size number-of-dialogs;
  }
  transactions {
    maximum-concurrent number;
    committed-attempts-rate transactions-per-second;
    committed-burst-size number-of-transactions;
  }
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure an admission control profle for a BSG.

Options ***admission-control-profile***—Name of the admission control profile.



NOTE: You can define a maximum of 100 admission control profiles for a BSG.

Other options are described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Call Admission Control in the *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

admission-control (New Transaction Policy)

Syntax admission-control *admission-control-profile*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-call-usage-policy *policy-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specifies the CAC admission controller used for this policy.

Options ***admission-control-profile***—Name of the admission control profile used for this policy.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Call Admission Control in the *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

availability-check-profiles

Syntax

```
availability-check-profiles {
  profile-name;
  keepalive-interval {
    available-server seconds;
    unavailable-server seconds;
  }
  keepalive-method sip-options;
  keepalive-strategy (do-not-send <blackout-period seconds> | send-always
    <failures-before-unavailable number> <successes-before-available number> |
    send-when-unavailable <successes-before-available number>);
  transaction-timeout seconds;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip routing-destinations]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure options used to determine that a server is available and able to receive SIP messages.

Options The options are explained separately.

Default: 600

available-server *seconds*—Number of seconds between pinging requests to an available server.

Range: 10 to 86,400

Default: 32

unavailable-server *seconds*—Number of seconds between pinging requests to an unavailable server.

Range: 10 to 86,400

Default: 32

keepalive-strategy—Specify the strategy for checking server availability.

Default: **send-when-unavailable**

send-always—Always check the availability of this server.

failures before unavailable *num*—Number of failures before the server is considered unavailable and placed in the server blacklist.

Range: 1 to 10

Default: 1

successes-before-available *number*—Number of successes before the server is considered available and removed from the server blacklist.

Range: 1 to 10

Default: 1

Copyright © 2013, Juniper Networks, Inc.

send-when-unavailable—Check server availability only if it is in the server blacklist.

successes-before-available *number*—Number of successes before the server is considered available and removed from the server blacklist.

do-not-send—Do not perform availability checking.

blackout-period *seconds*—Define the period, in seconds, during which the server is considered unavailable.

Range: 0 (no blacklisting) to 86,400 (24 hours)

transaction-timeout *seconds*—The number of seconds from the initiation of a ping request after which, if no reply, the ping is considered to have failed.

Range: 10 to 32

Default: 32

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Creating Availability Profiles for Servers in *Session Border Control Solutions Guide Using BGF and IMSG*

blacklist-period

Syntax
blacklist-period *seconds*;

Hierarchy Level
[edit services border-signaling-gateway gateway *gateway-name* name-resolution-cache]

Release Information
Statement introduced in Junos OS Release 10.0.

Description
Configure the amount of time that a SIP server remains in the black list. If the BSG finds that a server for a transaction is down, it marks the server as unavailable. The BSG does not forward SIP messages to a server on the black list until the blacklist period ends.

Options
seconds—Once the BSG marks a server as unavailable, this is the amount of time that the server remains on the blacklist.

Range: 0 through 86,400

Default: 600

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring DNS Resolution for Locating SIP Servers in the *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

clusters

Syntax

```
clusters [
  cluster-name;
  server server-name {
    priority priority-level;
    weight weight-level;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip routing-destinations]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure clusters of servers to use as routing destinations.

Options

server-name—Name of the server to include in the cluster.

priority-level—Relative priority, or redundancy order, as a choice for routing destination. 0 is the highest priority.

Range: 0 to 65,535

Default: 1

weight —Relative proportion of transactions within the specified priority level to route to this server.

Range: 0 to 65,535

Default: 1

profile—Name of the admission control profile assigned to this server.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- SIP Routing with Server Clusters Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Server Clusters in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

committed-burst-size

Syntax committed-burst-size *bytes*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* embedded-spdf service-class *service-class-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the maximum number of bytes allowed for incoming packets to burst above the committed information rate.



NOTE: When you configure committed-burst-size you must also configure committed-information-rate.

Options *bytes*—Number of bytes.

Range: 20 through 4,294967,295

Default: 10,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Admission Control Profiles in the *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

committed-information-rate

Syntax committed-information-rate *bytes-per-second*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* embedded-spdf service-class *service-class-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the maximum bandwidth that can be allocated to a packet that is flowing under normal line conditions.



NOTE: When you configure committed-information-rate you must also configure committed-burst-size.

Options *bytes-per-second*—Number of bytes per second.

Range: 125 through 4,294,967,295

Default: 2000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring QoS and Rate Limiting in the *Session Border Control Solutions Guide Using BGF and IMSG*

data-inactivity-detection (Services Border Signaling Gateway)

Syntax

```
data-inactivity-detection {
  inactivity-duration seconds;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-call-usage-policy *policy-name* term *term-name* then media-policy]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure data inactivity detection to detect latch deadlocks or other media inactivity on a gate.

Options The statement is described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Latch Deadlock and Media Inactivity Detection in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

datastore

Syntax

```
datastore {
  data trace-level;
  db trace-level;
  handle trace-level;
  minimum trace-level;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure trace-level options for the datastore component of the BSG.

Options

data *trace-level*—Trace level for the **data** subcomponent.

db *trace-level*—Trace level for the wrapper layer around the database.

handle *trace-level*—Trace level for the access API for the database.

minimum *trace-level*—Minimum trace level for all **datastore** messages.

trace-level—Trace-level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

default-media-realm

Syntax default-media-realm *realm-number*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a default value for the media-realm for each new call. The BGF uses media-realm to locate a virtual interface with the same value to determine the NAT pool for the call.

Options ***realm-number***—The realm number used to match to a virtual interface.

Range: 0 through 1023

Default: 0

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

dialogs

Syntax

```
dialogs {
  maximum-concurrent number;
  committed-attempts-ratedialogs-per-second;
  committed—burst-size number-of-dialogs
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* admission-control]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure admission control settings for dialogs.

Options

maximum-concurrent *number*—Maximum number of concurrent dialogs.

Values: 0 through 100,000

0 causes all calls to be rejected.

Default: 100000

committed-attempts-rate *dialogs-per-second*—Maximum number of attempts per second to initiate dialogs.

Values: 0 though 500

Default: 500

committed-burst-size *number-of-dialogs*—Maximum number of dialogs allowed to burst above the committed-rate and still be accepted.

Values: 0 through 1000

Default: 1000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Admission Control Profiles in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

dscp (Services Border Signaling Gateway)

Syntax dscp (*dscp-value* | *alias* | do-not-change);

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* embedded-spdf service-class *service-class-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure values for DSCP marking that the BSG uses for traffic that matches the service class term.

Default If you do not specify a DSCP value, the default value is **do-not-change**.

Options ***dscp-value***—String of six bits.

alias—Standard DSCP name. Use the **?** in the CLI to see a list of aliases.

do-not-change—Do not override the DSCP value in the packet.
Default: be

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Quality of Service for VoIP Traffic Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring QoS for the BGF in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

egress-service-point

Syntax egress-service-point *service-point-name*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then route (Services)]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the exit point of SIP requests from the BSG.

Options ***service-point-name***—Name of the service point that you want to use as the egress service point. This is a service point that you configure with the service-point statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls*Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

embedded-spdf

Syntax

```
embedded-spdf {
  service-class service-class-name {
    term term-name {
      from {
        media-type (any-media | audio | video);
      }
      then {
        committed-burst-size bytes;
        committed-information-rate bytes-per-second;
        dscp (dscp-value | alias | do-not-change);
        reject;
      }
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure an SPDF (session policy decision function). Each BSG instance consists of a single embedded SPDF that includes one or more service classes.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- IMSG Session Border Control Solution Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

file (Services Border Signaling Gateway)

Syntax file <*filename*> <files *files*> <match *regular-expression*> <size *maximum-file-size*> <world-readable | no-world-readable>;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the trace file for tracing BSG components.

Options

filename *filename*—Name of the file to which the tracing messages are written.

Default: bsg_trace

files *number-of-files*—Number of trace files. The tracing mechanism can rotate between any given number of files, allowing for trace message inspection without interfering with the normal work of the application.

Default: 3

match *regular expression*—Regular expression to match with incoming messages. Messages that do not match the regular expression are not written to the trace file.

size *maximum-trace-file-size*—Size parameter (in bytes) to trigger rotation of files. The trace mechanism rotates files based on the current file size. When the size is bigger than the maximum configured size, the files are rotated.

Default: 1048576

world-readable | no-world-readable—Allow all users to use the log file or disallow all users from using the log file.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BSG Operations *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

flag (Services Border Signaling Gateway)

Syntax

```
flag {
  datastore {
    data trace-level;
    db trace-level;
    handle trace-level;
    minimum trace-level;
  }
  framework {
    action trace-level;
    event trace-level;
    executor trace-level;
    freezer trace-level;
    minimum trace-level;
    memory-pool trace-level;
  }
  minimum trace-level;
  sbc-utils {
    common trace-level;
    configuration trace-level;
    device-monitor trace-level;
    ipc trace-level;
    memory-management trace-level;
    message trace-level;
    minimum trace-level;
    user-interface trace-level;
  }
  session-trace trace-level;
  signaling {
    b2b trace-level;
    b2b-wrapper trace-level;
    minimum trace-level;
    policy trace-level;
    sip-stack-wrapper trace-level;
    topology-hiding trace-level;
    ua trace-level;
  }
  sip-stack {
    dev-logging;
    event-tracing;
    ips-tracing;
    pd-log-detail (full | summary);
    pd-log-level (audit | exception | problem);
    per-tracing;
    verbose-logging;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway*gateway-name* traceoptions]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure trace options for components of the BSG.

Copyright © 2013, Juniper Networks, Inc.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing BSG Operations *Session Border Control Solutions Guide Using BGF and IMSG*

forward-manipulation

Syntax
```
forward-manipulation {
    manipulation-rule-name;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then message-manipulation]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the forward message manipulation rules that you want to add to your new transaction policy. Forward manipulation rules are applied to any message going from the user agent client (UAC), or the caller, to the user agent server (UAS), or the callee. They are applied to the original transaction request. If the transaction creates a dialog, the rules are also applied to other transaction requests or responses within the dialog.

Options *manipulation-rule-name*—Name of the message manipulation rule that you want to add. You can add up to five forward manipulation rules to a policy. These rules must have been configured at the **[edit** services (Border Signaling Gateway) **border-signaling-gateway** gateway (Services Border Signaling Gateway) ***gateway-name*** sip message-manipulation-rules manipulation-rule ***rule-name*** actions**]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Using New Transaction Policies to Manipulate SIP Headers or to Reject SIP Messages in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

framework

Syntax

```
framework {
    action trace-level;
    event trace-level;
    executor trace-level;
    freezer trace-level;
    minimum trace-level;
    memory-pool trace-level;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure trace options for the BSG component that provides an infrastructure that enables incremental functionality implementation.

Options

action *trace-level*—Trace level for the **framework** subcomponent that creates, initiates, and manipulates event actions.

event *trace-level*—Trace level for the **framework** subcomponent that creates, modifies, and terminates event members.

executor *trace-level*—Trace level for the **framework** subcomponent that executes configured actions for an event, handles any error states, delays processing, and so on.

freezer *trace-level*—Trace level for the **framework** subcomponent that delays the execution of an event until certain conditions are met.

minimum *trace-level*—Minimum trace level for all **framework** messages.

memory-pool *trace-level*—Trace level for the **framework** subcomponent that creates, deletes, and manipulates memory pools and pool managers, and controls the check-in and check-out of memory objects to and from memory pools.

trace-level—Trace-level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

from

See the following sections:

- from (New Call Usage Policy) on page 687
- from (New Transaction Policy) on page 688
- from (Service Class) on page 690

Copyright © 2013, Juniper Networks, Inc.

from (New Call Usage Policy)

Syntax

```
from {
  contact {
    regular-expression [ regular-expression ];
  }
  method {
    method-invite;
  }
  request-uri {
    regular-expression [ regular-expression ];
  }
  source-address [ ip-addresses ];
}
```

Hierarchy Level [edit services border-signaling-gateway gateway*gateway-name* sip new-call-usage-policy *policy-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 9.4.
regular-expression options introduced in Junos OS Release 9.5.

Description Configure match conditions for a new call usage policy.

Options

contact—Match the contents of the contact field. Contact field matching is based on regular expressions.

regular expression [*regular-expression*] —Regular expression used to match the contents of the contact field.

Syntax: To specify more than one regular expression, enclose the regular expressions in brackets.

method-invite—Match the policy to SIP INVITE methods.

request-uri—Match the contents of the uniform resource identifier (URI) in the SIP message request. Request URI matching is based on regular expressions.

regular expression [*regular-expression*]—Regular expression used to match the contents of the request URI field.

Syntax: To specify more than one regular expression, enclose the regular expressions in brackets.

source-address—Match the source address of the SIP request.

[*ip-addresses*]—IP addresses that you want to match.

Syntax: To specify more than one IP address, enclose the IP addresses in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a New Call Usage Policy in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

from (New Transaction Policy)

Syntax

```
from {
  contact {
    registration-state [ registered | not-registered ];
    regular-expression [ regular-expression ];
    uri-hiding [ hidden-uri | not-hidden-uri ];
  }
  method {
    method-invite;
    method-message;
    method-options;
    method-publish;
    method-refer;
    method-register;
    method-subscribe;
  }
  request-uri {
    registration-state [ registered | not-registered ];
    regular-expression [ regular-expression ];
    uri-hiding [ hidden-uri | not-hidden-uri ];
  }
  source-address [ ip-addresses ];
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 9.4.
regular-expression options introduced in Junos OS Release 9.5.

Description Configure match conditions for a new transaction policy.

Options

contact—Match the contents of the contact field. Contact field matching is based on regular expressions.

registration-state—Select transactions based on whether the BSG passed a SIP register message for the transaction to a SIP registrar.

Values:

- registered—Select transactions for which the BSG passed a SIP register message to a SIP registrar.
- not-registered—Select transactions for which the BSG did not pass a SIP register message to a SIP registrar.

regular expression [*regular-expression*]—Regular expression used to match the contents of the contact field.

Syntax: To specify more than one regular expression, enclose the regular expressions in brackets.

uri-hiding—Select transactions based on whether contact URIs are hidden.

Values:

Copyright © 2013, Juniper Networks, Inc.

- **hidden-uri**—Select transactions for which the contact URI is hidden.
- **unhidden-uri**—Select transactions for which the contact URI is not hidden.

method—Match the type of SIP method.

Syntax: To specify multiple SIP methods, use separate **set** statements.

request-uri—Match the contents of the uniform resource identifier (URI) in the SIP message request. Request URI matching is based on regular expressions.

registration-state—Select transactions based on whether the transactions are from registered request URIs.

Values:

- **registered**—Select transactions for which the contact URI is hidden.
- **unhidden-uri**—Select transactions for the contact URI is not hidden.

regular expression [*regular-expression*] —Regular expression used to match the contents of the request URI field.

Syntax: To specify more than one regular expression, enclose the regular expressions in brackets.

source-address—Match the source address of the SIP request.

[*ip-addresses*]—IP addresses that you want to match.

Syntax: To specify more than one IP address, enclose the IP addresses in brackets.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a New Transaction Policy in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

from (Service Class)

Syntax

```
from {
  media-type (any-media | audio | video);
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* embedded-spdf service-class *service-class-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure match conditions for a service class.

Options The options are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a New Call Usage Policy in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring QoS and Rate Limiting in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

gateway (Services Border Signaling Gateway)

Syntax

```
gateway gateway-name {
    admission-control controller-name {
      dialogs {
        maximum-concurrent number;
        committed-attempts-rate dialogs-per-second;
        committed-burst-size number-of-dialogs;
      }
      transactions {
        maximum-concurrent number;
        committed-attempts-rate transactions-per-second;
        committed-burst-size number-of-transactions;
      }
  embedded-spdf {
    service-class service-class-name {
      term term-name {
        from {
          media-type (any-media | audio | video);
        }
        then {
          committed-burst-size bytes;
          committed-information-rate bytes-per-second;
          dscp (alias | do-not-change | dscp-value);
          reject;
        }
      }
    }
  }
  name-resolution-cache {
    accelerations {
      initiate-alternative-queries;
      initiate-next-queries;
      no-refresh-before-ttl-expiry;
    }
    blacklist-period seconds;
    maximum-records-in-cache number;
    maximum-time-in-cache (unlimited | seconds);
  }
  service-interface name;
  service-point service-point-name {
    default-media-realm;
    service-interface interface-name.unit-number;
    service-point-type service-point-type;
    service-policies {
      new-call-usage-input-policies [policy-and-policy-set-names];
      new-call-usage-output-policies [policy-and-policy-set-names];
      new-transaction-input-policies [policy-and-policy-set-names];
      new-transaction-output-policies[policy-and-policy-set-names];
    }
    transport-details <port port-number> <ip-address ip-address> <tcp> <udp>;
  }
  sip {
    message-manipulation-rules {
```

Copyright © 2013, Juniper Networks, Inc.

```
    manipulation-rule rule-name {
      actions {
        sip-header header-field-name {
          field-value {
            modify-regular-expression regular-expression with field-value;
            add field-value;
            add-missing field-value;
            add-overwrite field-value;
            remove-regular-expression regular-expression;
            remove-all;
            reject-regular-expression regular-expression;
          }
        }
        request-uri request-uri {
          field-value {
            modify-regular-expression regular-expression with field-value;
          }
        }
      }
    }
  }
  new-call-usage-policy policy-name {
    term term-name {
      from {
        contact [ contact-fields ];
        method {
          method-invite;
        }
        request-uri [ uri-fields ];
        source-address [ ip-addresses ];
      }
      then {
        media-policy {
          data-inactivity-detection (Services PGCP) {
            inactivity-duration seconds;
          }
          no-anchoring;
          service-class service-class-name;
        }
        trace;
      }
    }
  new-call-usage-policy-set policy-set-name {
    policy-name [ policy-names ];
  }
  new-transaction-policy policy-name {
    term term-name {
      from {
        contact {
          registration-state [ registered | not-registered ];
          regular-expression [ regular-expression ];
          uri-hiding [ hidden-uri | not-hidden-uri ];
        }
        method {
          method-invite;
          method-message;
```


Copyright © 2013, Juniper Networks, Inc.

```
            method-options;
            method-publish;
            method-refer;
            method-register;
            method-subscribe;
          }
          request-uri {
            registration-state [ registered | not-registered ];
            regular-expression [ regular-expression ];
            uri-hiding [ hidden-uri | not-hidden-uri ];
          }
          source-address [ ip-addresses ];
        }
        then {
          (accept | reject);
          admission-control admission-control-profile;
          message-manipulation {
            forward-manipulation {
              manipulation-rule-name;
            }
            reverse-manipulation {
              manipulation-rule-name;
            }
          }
          on-3xx-response {
            recursion-limit number;
            }
          }
          route (Services) {
            egress-service-point service-point-name;
            next-hop (Services) (sip-based | address ipv4-address <port port-number>
              <transport-protocol (udp | tcp)>);
            server-cluster cluster-name;
          }
          signaling-realm signaling-realm;
          topology-hiding {
            maintain-route-headers;
          }
          trace;
        }
      }
    }
    new-transaction-policy-set policy-set-name {
      policy-name [ policy-names ];
    }
    routing-destinations {
      availability-check-profiles {
        profile-name;
        keepalive-interval {
          available-server seconds;
          unavailable-server seconds;
        }
        keepalive-method sip-options;
        keepalive-strategy (do-not-send <blackout-period seconds> | send-always <
          failures-before-unavailable number> < successes-before-available number |
          send-when-unavailable < successes-before-available number);
```

Copyright © 2013, Juniper Networks, Inc.

```
        transaction-timeout seconds;
      }
      clusters [
        cluster-name;
        server server-name {
          priority priority-level;
          weight weight-level;
        }
      }
      default-availability-check-profile profile-name;
      }
      servers {
        server-name {
        address ip4-address <port port-number> <transport (udp | tcp)>;
        admission-control profile-name;
        availability-check-profile profile-name;
        service-point service-point-name;
      }
    timers {
      inactive-call seconds;
      timer-c seconds;
    }
}
traceoptions {
  file <filename> <files files> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag {
    datastore {
      data trace-level;
      db trace-level;
      handle trace-level;
      minimum trace-level;
    }
    framework {
      action trace-level;
      event trace-level;
      executor trace-level;
      freezer trace-level;
      minimum trace-level;
      memory-pool trace-level;
    }
    minimum trace-level;
    sbc-utils {
      common trace-level;
      configuration trace-level;
      device-monitor trace-level;
      ipc trace-level;
      memory-management trace-level;
      message trace-level;
      minimum trace-level;
      user-interface trace-level;
    }
    session-trace trace-level;
    signaling {
      b2b trace-level;
      b2b-wrapper trace-level;
```


Copyright © 2013, Juniper Networks, Inc.

```
            minimum trace-level;
            policy trace-level;
            sip-stack-wrapper trace-level;
            topology-hiding trace-level;
            ua trace-level;
          }
          sip-stack {
            dev-logging;
            event-tracing;
            ips-tracing;
            pd-log-detail (full | summary);
            pd-log-level (audit | exception | problem);
            per-tracing;
            verbose-logging;
          }
        }
      }
    }
```

Hierarchy Level [edit services border-signaling-gateway]

Release Information Statement introduced in Junos OS Release 9.4.
data-inactivity-detection option introduced in Junos OS Release 9.6.
message-manipulation option introduced in Junos OS Release 9.6.
message-manipulation-rules option introduced in Junos OS Release 9.6.
name-resoltuion-cache option introduced in Junos OS Release 10.0.

Description Configure a border signaling gateway instance.

Options *gateway-name*—Identifier for the BSG.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- IMSG Session Border Control Solution Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

inactivity-duration (Services Border Signaling Gateway)

Syntax inactivity-duration *seconds*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-call-usage-policy *policy-name* term *term-name* then media-policy data-inactivity-detection]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the time interval that determines inactivity. When the virtual BGF determines that the time since the last packet was received exceeds this duration, the virtual BGF generates an inactivity notification or service change request. The duration timer is the same for terminations with latch events and for terminations without latch events.

Options ***seconds***—Time during which no packets are received.

Range: 5 through 86400

Default: 30

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

manipulation-rule

Syntax

```
manipulation-rule rule-name {
  actions {
    sip-header header-field-name {
      field-value {
        add field-value;
        add-missing field-value;
        add-overwrite field-value;
        modify-regular-expression regular-expression with field-value;
        reject-regular-expression regular-expression;
        remove-all;
        remove-regular-expression regular-expression;
      }
    }
    request-uri request-uri {
      field-value {
        modify-regular-expression regular-expression with field-value;
      }
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip message-manipulation-rules]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a rule for manipulating the header fields or the request URI in SIP messages. You can have up to 1,000 manipulation rules for a BSG.

Options ***rule-name***—Name of the manipulation rule.

The remaining options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Manipulation of Headers and Request URIs in SIP Messages in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

media-policy

Syntax

```
media-policy {
  data-inactivity-detection {
    inactivity-duration seconds;
  }
  media-release;
  no-anchoring;
  service-class service-class-name;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-call-usage-policy *policy-name* term *term-name* then]

Release Information

Statement introduced in Junos OS Release 9.4.
media-release statement introduced in Junos OS Release 10.1.
no-anchoring statement introduced in Junos OS Release 9.5.
service-class statement introduced in Junos OS Release 9.5.

Description Configure the service class to be applied to traffic that matches the new call usage policy.

Options

media-release—Disable or enable media release for the policy.

no-anchoring—Disable or enable media anchoring for the policy.

service-class *service-class-name*—Name of the service class to be applied to traffic that matches the new call usage policy. You must have configured the service class using the service-class statement at the **[edit** services (Border Signaling Gateway) **border-signaling-gateway** gateway ***gateway-name*** embedded-spdf**]** hierarchy level.

The remaining options are described separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a New Call Usage Policy in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

media-type

Syntax media-type (any-media | audio | video);

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-class *service-class-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the type of media that the service class matches.

Options **any-media**—Match all media types.

audio—Match audio traffic.

video—Match video traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

message-manipulation

Syntax

```
message-manipulation {
  forward-manipulation {
    manipulation-rule-name;
  }
  reverse-manipulation {
    manipulation-rule-name;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the forward and reverse message manipulation rules that you want to add to your new transaction policy. When the message manipulation rules in your policy match a transaction, the transaction is affected as well as any transactions that belong to a dialog that results from the transaction.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

maximum-records-in-cache

Syntax maximum-records-in-cache *number*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* name-resolution-cache]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configures the maximum number of SIP servers that can be in the DNS name resolution cache. When this number is exceeded, servers are removed from the cache starting with the least recently used entry.

Options ***number***—Number of servers that can be stored in the name resolution cache. A setting of 0 means that there is no name resolution cache.

Range: 0 through 50,000

Default: 5000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring DNS Resolution for Locating SIP Servers

maximum-time-in-cache

Syntax maximum-time-in-cache (unlimited | *seconds*);

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* name-resolution-cache]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configures the maximum time that a SIP server can be held in the DNS name resolution cache. Each server entry has a time to live (TTL) value that indicates how long the server can be saved in cache without being refreshed by a new query. You can override the TTL value to a lower value by setting the number of seconds that servers are held in cache. You can override the TTL only with a lower value. If the configured value is higher than the original TTL value, the original TTL value is applied.

Options **unlimited**—TTL value in the DNS entry is applied.

seconds—Time that an entry can remain in cache.

Range: 0 through 604,800

Default: unlimited

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring DNS Resolution for Locating SIP Servers

Copyright © 2013, Juniper Networks, Inc.

message-manipulation-rules

Syntax

```
message-manipulation-rules {
  manipulation-rule rule-name {
    actions {
      sip-header header-field-name {
        field-value {
          add field-value;
          add-missing field-value;
          add-overwrite field-value;
          modify-regular-expression regular-expression with field-value;
          reject-regular-expression regular-expression;
          remove-all;
          remove-regular-expression regular-expression;
        }
      }
      request-uri request-uri {
        field-value {
          modify-regular-expression regular-expression with field-value;
        }
      }
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure rules for manipulating the header fields or the request URI in SIP messages.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Manipulation of Headers and Request URIs in SIP Messages in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

minimum

Syntax minimum *trace-level*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag *flag*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the minimum trace level for all selected BSG trace options. This option overrides individual trace options that are set at a lower level.

Default warning

Options ***trace-level***—Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BSG Operations*Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

name-resolution-cache

Syntax

```
name-resolution-cache {
  accelerations {
    initiate-alternative-queries;
    initiate-next-queries;
    no-refresh-before-ttl-expiry;
  }
  blacklist-period seconds;
  maximum-records-in-cache number;
  maximum-time-in-cache (unlimited | seconds);
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure parameters that specify how entries are handled in the DNS name resolution cache and the type, if any, of acceleration that the BSG uses to accelerate the process of DNS name resolution for SIP servers.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS Resolution for Locating SIP Servers

new-call-usage-input-policies

Syntax new-call-usage-input-policies [*policy-and-policy-set-names*];

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name* service-policies]

Release Information Statement introduced in Junos OS Release 10.1.

Description Assign new call usage input policies or policy sets to calls that entered through the service point. All the packets arriving at the service point are matched against these policies.

Options **[*policy-and-policy-set-names*]**—Names of new call usage policies or policy sets.

Syntax: If you specify more than one policy or policy set, you must enclose all policy names in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Attaching Policies to a Service Point in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

new-call-usage-output-policies

Syntax new-call-usage-output-policies [*policy-and-policy-set-names*];

Hierarchy Level [edit services (Border Signaling Gateway) border-signaling-gateway gateway (Services Border Signaling Gateway) *gateway-name* service-point *service-point-name* service-policies]

Release Information Statement introduced in Junos OS Release 10.1.

Description Assign new call usage output policies or policy sets to calls that exited through the service point. All the packets leaving from the service point are matched against these policies.

Options [*policy-and-policy-set-names*]—Names of new call usage policies or policy sets.

Syntax: If you specify more than one policy or policy set, you must enclose all policy names in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Attaching Policies to a Service Point in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

new-call-usage-policy

Syntax

```
new-call-usage-policy policy-name {
  term term-name {
    from {
      contact [ contact-fields ];
      method {
        method-invite;
      }
      request-uri [ uri-fields ];
      source-address [ ip-addresses ];
    }
    then {
      media-policy {
        data-inactivity-detection (Services PGCP) {
          inactivity-duration seconds;
        }
        no-anchoring;
        service-class service-class-name;
      }
      trace;
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure a new call usage policy. A call is a usage that begins with a new INVITE. Dialogs can have many different usages.

Options ***policy-name***—Identifier for the new call usage policy.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

new-call-usage-policy-set

Syntax

```
new-call-usage-policy-set policy-set-name {
  policy-name [ policy-names ];
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 9.4.

Description Create a set of new call usage policies, which you can then apply to a service point. The order in which you add policies to the set determines the order in which the BSG processes the policies. The first matching policy determines which actions are taken.

Options

policy-set-name—Identifier for the new call usage policy set.

policy-names—Names of one or more new call usage policies that you want to add to the set.

Syntax: To specify a list of policies, enclose the policy names in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring New Call Usage Policy Sets in *Session Border Control Solutions Guide Using BGF and IMSG*

new-transaction-input-policies

Syntax new-transaction-input-policies [*policy-and-policy-set-names*];

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name* service-policies]

Release Information Statement introduced in Junos OS Release 10.1.

Description Assign new transaction policies or policy sets to the service point. All packets entering at the service point are matched against these policies.

Options

[*policy-and-policy-set-names*]—Names of new transaction policies or policy sets.

Syntax: To specify more than one policy or policy set, enclose the policy names in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Service Point in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

new-transaction-output-policies

Syntax new-transaction-output-policies [*policy-and-policy-set-names*];

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name* service-policies]

Release Information Statement introduced in Junos OS Release 10.1.

Description Assign new transaction policies or policy sets to the service point. All packets leaving from the service point are matched against these policies.



NOTE: You cannot assign a new transaction policy as a new transaction *output* policy if it contains route or message-manipulation statements.

Options [*policy-and-policy-set-names*]—Names of new transaction policies or policy sets.

Syntax: To specify more than one policy or policy set, enclose the policy names in brackets.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Service Point in*Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

new-transaction-policy

Syntax

```
new-transaction-policy policy-name {
  term term-name {
    from {
      contact {
        registration-state [ registered | not-registered ];
        regular-expression [ regular-expression ];
        uri-hiding [ hidden-uri | not-hidden-uri ];
      }
      method {
        method-invite;
        method-message;
        method-options;
        method-publish;
        method-refer;
        method-register;
        method-subscribe;
      }
      request-uri {
        registration-state [ registered | not-registered ];
        regular-expression [ regular-expression ];
        uri-hiding [ hidden-uri | not-hidden-uri ];
      }
      source-address [ ip-addresses ];
    }
    then {
      (accept | reject);
      admission-control admission-control-profile;
      message-manipulation {
        forward-manipulation {
          manipulation-rule-name;
        }
        reverse-manipulation {
          manipulation-rule-name;
        }
      }
      on-3xx-response {
        recursion-limit number;
        }
      }
      route (Services) {
        egress-service-point service-point-name;
        next-hop (Services) (sip-based | address ipv4-address <port port-number>
          <transport-protocol (udp | tcp)>);
        server-cluster cluster-name;
      }
      signaling-realm signaling-realm;
      topology-hiding {
        maintain-route-headers;
      }
      trace;
    }
  }
```

Copyright © 2013, Juniper Networks, Inc.

}

Hierarchy Level [edit services (Border Signaling Gateway) border-signaling-gateway gateway (Services Border Signaling Gateway) *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 9.4.
message-manipulation option introduced in Junos OS Release 9.5.

Description Specify new transaction policies for out-of-dialog transactions including dialog-opening transactions. Transaction policies are useful when the policy does not need to differentiate between events. For example, you can use new transaction policies to route all transactions according to the same rules.

A new transaction event is raised when a new SIP request, such as an INVITE, either opens a new dialog or is not related to any dialog. If the event does not match a new transaction policy, the BSG rejects the SIP request and returns a 403 (forbidden) message.

Options ***policy-name***—Identifier for the new transaction policy.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using New Transaction Policies to Manipulate SIP Headers or to Reject SIP Messages in*Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls in*Session Border Control Solutions Guide Using BGF and IMSG*
- Assigning Admission Control Profiles to New Transaction Policies in*Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

new-transaction-policy-set

Syntax

```
new-transaction-policy-set policy-set-name {
    policy-name [policy-names];
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 9.4.

Description Create a set of new transaction policies, which you can then apply to a service point. The order in which you add policies to the set determines the order in which the BSG processes the policies. The first matching policy determines which actions are taken.

Options

policy-set-name—Identifier for the new transaction policy set.

[*policy-names*]—Names of one or more new transaction policies that you want to add to the set.

Syntax: To specify a list of policies, enclose the policy names in brackets.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring New Transaction Policy Sets in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

next-hop (Services)

Syntax next-hop (sip-based | address *ipv4-address*<port *port-number*> <transport-protocol (udp | tcp)>);

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then route (Services)]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the SIP entity towards which SIP requests are sent.

Options

sip-based—All requests and responses on the dialog are routed according to SIP. If the configuration includes the **topology-hiding** option, the information in the Route header of the incoming SIP message is used. In all other cases, the **request-uri** is used. The software resolves the uniform resource identifier (URI) in the SIP message request into the IP address, port, and transport protocol of the next hop to contact.

address *ipv4–address*—Destination IPv4 address of the next hop to contact. This static address applies to all incoming requests on the dialog.

port—(Optional) Destination port of the next hop to contact.

Default: 5060

transport-protocol (udp | tcp)—(Optional) Transport protocol for routing to the next hop.

Default: udp

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Service Set in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

on-3xx-response

Syntax

```
on-3xx-response {
  recursion-limit number;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the action taken after receiving a 3XX response. When the **on-3xx-response** statement is included, the BSG sends a new, redirected request to the responding UAS, using a request URI based on the contact information in the 3XX response. When the **on-3xx-response** statement is not included, the 3XX response is passed back to the serving UAC.

Options *number*—The number of recursions allowed before sending a 408 timeout response.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Providing Redirection for Messages with 3XX Responses Overview
- Configuring Redirection for Messages with 3XX Responses

Copyright © 2013, Juniper Networks, Inc.

request-uri

Syntax

```
request-uri {
  field-value {
    modify-regular-expression regular-expression with field-value;
  }
}
```

Hierarchy Level

[edit services border-signaling-gateway gateway *gateway-name* sip message-manipulation-rules manipulation-rule *rule-name* actions]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Configure a rule for modifying the request uniform resource identifier (URI) in a SIP message.

Options

modify-regular-expression—Changes the value of a regular expression.

Syntax: modify-regular-expression *regular-expression* with *field-value*—Enter the regular expression that you want to modify followed by the value with which you want to replace the regular expression. In the following example, regular expression 1800 is replaced with 555:

```
modify-regular-expression 1800 with 555;
```

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring a New Call Usage Policy in *Session Border Control Solutions Guide Using BGF and IMSG*
- Manipulation of Headers and Request URIs in SIP Messages in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

reverse-manipulation

Syntax

```
reverse-manipulation {
    manipulation-rule-name;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then message-manipulation]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the reverse message manipulation rules that you want to add to your new transaction policy. Reverse manipulation rules are applied to any message going from the user agent server (UAS), or the callee, to the user agent client (UAC), or the caller. They are applied to the original transaction request. If the transaction creates a dialog, the rules are also applied to other transaction requests or responses within the dialog.

Options ***manipulation-rule-name***—Name of the message manipulation rule that you want to add. You can add up to five reverse manipulation rules to a policy. These rules must have been configured at the **[edit** services (Border Signaling Gateway) border-signaling-gateway gateway (Services Border Signaling Gateway) ***gateway-name*** sip message-manipulation-rules manipulation-rule ***rule-name*** actions**]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using New Transaction Policies to Manipulate SIP Headers or to Reject SIP Messages in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

route (Services)

Syntax

```
route {
  egress-service-point service-point-name;
  next-hop (Services) {
    (request-uri | address ipv4-address <port port-number> <transport-protocol (tcp | udp)>);
  }
  server-cluster cluster-name;
}
```

Hierarchy Level

[edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name* then]

Release Information

Statement introduced in Junos OS Release 9.4.
server-cluster option introduced in Junos OS Release 10.2.

Description

Configure the next-hop destination and egress service point for a new transaction policy. Alternatively, you can specify a server cluster.

Options

The options are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*
- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

routing-destinations

Syntax

```
routing-destinations {
  availability-check-profiles {
    profile-name;
    keepalive-interval {
      available-server seconds;
      unavailable-server seconds;
    }
    keepalive-method sip-options;
    keepalive-strategy (do-not-send <blackout-period seconds> | send-always <
      failures-before-unavailable number> < successes-before-available number |
      send-when-unavailable < successes-before-available number);
    transaction-timeout seconds;
  }
  clusters [
    cluster-name;
    server server-name {
      priority priority-level;
      weight weight-level;
    }
  }
  default-availability-check-profile profile-name;
  }
  servers {
    server-name {
    address ip4-address <port port-number> <transport (udp | tcp)>;
    admission-control profile-name;
    availability-check-profile profile-name;
    service-point service-point-name;
  }
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure servers, server clusters, and availability rules for routing destinations.

Options **default-availability-check-profile *profile-name***—Availability check profile that is assigned to a server when no profile is explicityly defined..

The remaining options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- SIP Routing with Server Clusters Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sbc-utils (Services Border Signaling Gateway)

Syntax

```
sbc-utils {
    common trace-level;
    configuration trace-level;
    device-monitor trace-level;
    ipc trace-level;
    memory-management trace-level;
    message trace-level;
    minimum trace-level;
    user-interface trace-level;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure trace options for the Signaling Border Controller (SBC) utilities component of the BSG.

Default warning

Options

common *trace-level*—Trace level for the common component of SBC utilities.

configuration *trace-level*—Trace level for the configuration component of SBC utilities.

device-monitor *trace-level*—Trace level for the device monitor component of SBC utilities.

ipc *trace-level*—Trace level for the IPC component of SBC utilities.

memory-management *trace-level*—Trace level for the memory management component of SBC utilities.

message *trace-level*—Trace level for the message component of SBC utilities.

minimum *trace-level*—Minimum trace level for all **sbc-util** messages.

user-interface *trace-level*—Trace level for the user interface component of SBC utilities.

trace-level—Trace-level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START and EXIT macros.
- **info**—Summary logs for normal operations, such as the policy decisions made for a call.
- **warning**—Failure recovery or failure of an external entity.
- **error**—Failure with a short-term effect, such as failed processing of a single call.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

servers

Syntax

```
servers {
  server-name {
  address ip4-address <port port-number> <transport (udp | tcp)>;
  admission-control profile-name;
  availability-check-profile profile-name;
  service-point service-point-name;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip routing-destinations]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure one or more servers for use as routing destinations.

Options

ip4-address—IP address of the server.

port-number— Port number to use on the server.

Range: 0 to 65,535

Default: 5060

tcp|udp—Transport protocol to be used on this server.

Default: udp

admission-control *profile-name*—(Optional) Name of the admission control profile used by the server.

availability-check—profile *profile-name*—Name of the availability check profile used by the server. If no availability check profile is specified, the default values are used.

service-point *service-point-name*—Name of the service point through which traffic is routed to the server.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- SIP Routing with Server Clusters Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Servers for Use in Server Clusters in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-class

Syntax

```
service-class service-class-name {
  term term-name {
    from {
      media-type (any-media | audio | video);
    }
    then {
      committed-burst-size bytes;
      committed-information-rate bytes-per-second;
      dscp (alias | do-not-change | dscp-value);
      reject;
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* embedded-spdf]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure service classes for the embedded SPDF. Service classes contain rules that pertain to the treatment of bandwidth for various media types. Each rule (or term) consists of a **from** statement and a **then** statement. The **from** statement matches traffic based on the media type. The **then** statement is a set of one or more actions that are applied if a call matches the from statement.

Options ***service-class-name***—Identifier for the service class.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Providing QoS for VoIP Traffic Overview *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring a New Call Usage Policy *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring QoS and Rate Limiting *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-interface

See the following sections:

- service-interface (Gateway) on page 721
- service-interface (Service Point) on page 721

service-interface (Gateway)

Syntax service-interface *interface-name.unit-number*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Assign the BSG to a Multiservices PIC or DPC. The PIC or DPC must have been configured at the **[edit interfaces]** hierarchy. You can assign only one BSG to a Multiservices PIC or DPC.

Options *interface-name.unit-number*—Name and logical unit number of the Multiservices PIC or DPC.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

service-interface (Service Point)

Syntax service-interface *name*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Assign the service point to a service interface. The policies attached to the service point are matched against incoming requests received on this service interface.

Options *name*—Name of the service interface. The interface must have been configured at the **[edit interfaces]** hierarchy.

Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-point

Syntax

```
service-point service-point-name {
  default-media-realm
  service-interface interface-name.unit-number;
  service-point-type service-point-type;
  service-policies {
    new-call-usage-input-policies [policy-and-policy-set-names];
    new-call-usage-output-policies [policy-and-policy-set-names];
    new-registration-input-policies [policy-and-policy-set-names];
    new-transaction-input-policies [policy-and-policy-set-names];
    new-transaction-output-policies[policy-and-policy-set-names];
  }
  transport-details <port port-number> <ip-address ip-address> <tcp> <udp>;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure a service point. Service points identify a service interface and transport parameters for incoming requests. You attach policies to the service point, and all requests that arrive at the service point are handled by these policies. Each BSG can have five service points.

You can also configure a service point to be used as an egress service point to which SIP requests are routed.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- BSG Policy Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

service-point-type

Syntax service-point-type *service-point-type*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Create the type of VoIP protocol for this service point.

Options ***service-point-type***—VoIP protocol. Currently the only protocol type supported is SIP.
Values: sip

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Service Point in *Session Border Control Solutions Guide Using BGF and IMSG*

service-policies

Syntax
```
service-policies {
  new-call-usage-input-policies [policy-and-policy-set-names];
  new-call-usage-output-policies [policy-and-policy-set-names];
  new-transaction-input-policies [policy-and-policy-set-names];
  new-transaction-output-policies [policy-and-policy-set-names];
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name*]

Release Information Statement introduced in Junos OS Release 9.4.
new-call-usage-input-policies statement added in Junos OS Release 10.1.
new-call-usage-output-policies statement added in Junos OS Release 10.1.
new-transaction-input-policies statement added in Junos OS Release 10.1.
new-transaction-output-policies statement added in Junos OS Release 10.1.

Description Specify the policies and policy sets that are applied to the service point.

Options The options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Service Point in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

services (Border Signaling Gateway)

Syntax services border-signaling-gateway { ... }

Hierarchy Level [edit]

Description Define service rules to be applied to traffic.

Options **border-signaling-gateway**—Identifier for the BSG set of statements.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

session-trace

Syntax session-trace handle *trace-level*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure tracing for transactions matching policies that have their trace flag turned on. The tracing level is effective for dialog-creating messages (such as INVITE) and out-of-dialog messages. When these message types are accepted in the policy and the policy is set to trace messages, the policy marks the dialog (and the sibling dialog) for session tracing.

Default warning

Options **minimum *trace-level***—The minimum trace level for all **session-trace** messages.

trace-level—Trace level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START, EXIT macros.
- **info**—Summary logs for normal operations e.g. the policy decisions made for a call.
- **warning**—Failure-recovery or failure of an external entity.
- **error**—Failure with short-term effect, such as failed processing of a single call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

signaling

Syntax

```
signaling {
  b2b trace-level;
  b2b-wrapper trace-level;
  minimum trace-level;
  policy trace-level;
  sip-stack-wrapper trace-level;
  topology-hiding trace-level;
  ua trace-level;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure trace options for the signaling component of the BSG.

Default warning

Options

b2b *trace-level*—Trace options for the signaling component that implements the b2b logic (translating between dialogs, associating dialogs, creating new downstream dialogs, and so on).

b2b-wrapper *trace-level*—Trace options for entry and exit to the BSG signaling application.

minimum *trace-level*—Minimum trace level for all **signaling** messages.

policy *trace-level*—Trace options for the signaling component that applies policies for call admission, routing decisions, security settings, and so on.

sip-stack-wrapper *trace-level*—Trace options for the glue layer that receives events from the SIP stack and forwards them to the application and, conversely, receives events from the application and forwards them to the SIP stack.

topology-hiding *trace-level*—Trace options for the signaling component that hides the network topology of a network by CONTACT replacement and removal or modification of certain headers.

ua *trace-level*—Trace options for the signaling subcomponent that handles RECEIVE messages.

trace-level—Trace level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:

- **debug**—Logging of all code flow of control.
- **trace**—Logging of program trace START, EXIT macros.
- **info**—Summary logs for normal operations e.g. the policy decisions made for a call.
- **warning**—Failure-recovery or failure of an external entity.

Copyright © 2013, Juniper Networks, Inc.

- **error**—Failure with short-term effect, such as failed processing of a single call.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

signaling-realms

Syntax
```
signaling-realms {
  realm realm-name;
}
```

Hierarchy Level
[edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information
Statement introduced in Junos OS Release 10.2.

Description
Define signaling realms to be assigned to registered new transactions based on new transaction policy selection criteria.

Options
realm *realm-name*—Name of a signaling realm.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a New Transaction Policy in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sip

Syntax

```
sip {
    message-manipulation-rules {
        manipulation-rule rule-name {
            actions {
                sip-header header-field-name {
                    field-value {
                        modify-regular-expression regular-expression with field-value;
                        add field-value;
                        add-missing field-value;
                        add-overwrite field-value;
                        remove-regular-expression regular-expression;
                        remove-all;
                        reject-regular-expression regular-expression;
                    }
                }
                request-uri request-uri {
                    field-value {
                        modify-regular-expression regular-expression with field-value;
                    }
                }
            }
        }
    }
    new-call-usage-policy policy-name {
        term term-name {
            from {
                contact [ contact-fields ];
                method {
                    method-invite;
                }
                request-uri [ uri-fields ];
                source-address [ ip-addresses ];
            }
            then {
                media-policy {
                    data-inactivity-detection (Services PGCP) {
                        inactivity-duration seconds;
                    }
                    no-anchoring;
                    service-class service-class-name;
                }
                trace;
            }
        }
    }
    new-call-usage-policy-set policy-set-name {
        policy-name [ policy-names ];
    }
    new-transaction-policy policy-name {
        term term-name {
            from {
                contact {
                    registration-state [ registered | not-registered ];
```


Copyright © 2013, Juniper Networks, Inc.

```
            regular-expression [ regular-expression ];
            uri-hiding [ hidden-uri | not-hidden-uri ];
          }
          method {
            method-invite;
            method-message;
            method-options;
            method-publish;
            method-refer;
            method-register;
            method-subscribe;
          }
          request-uri {
            registration-state [ registered | not-registered ];
            regular-expression [ regular-expression ];
            uri-hiding [ hidden-uri | not-hidden-uri ];
          }
          source-address [ ip-addresses ];
        }
        then {
          (accept | reject);
          admission-control admission-control-profile;
          message-manipulation {
            forward-manipulation {
              manipulation-rule-name;
            }
            reverse-manipulation {
              manipulation-rule-name;
            }
          }
          on-3xx-response {
            recursion-limit number;
            }
          }
          route (Services) {
            egress-service-point service-point-name;
            next-hop (Services) (sip-based | address ipv4-address <port port-number>
              <transport-protocol (udp | tcp)>);
            server-cluster cluster-name;
          }
          signaling-realm signaling-realm;
          topology-hiding {
            maintain-route-headers;
          }
          trace;
        }
      }
    }
    new-transaction-policy-set policy-set-name {
      policy-name [ policy-names ];
    }
    routing-destinations {
      availability-check-profiles {
        profile-name;
        keepalive-interval {
          available-server seconds;
```

Copyright © 2013, Juniper Networks, Inc.

```
        unavailable-server seconds;
      }
      keepalive-method sip-options;
      keepalive-strategy (do-not-send <blackout-period seconds> | send-always <
        failures-before-unavailable number> < successes-before-available number |
        send-when-unavailable < successes-before-available number);
      transaction-timeout seconds;
    }
    clusters [
      cluster-name;
      server server-name {
        priority priority-level;
        weight weight-level;
      }
    }
    default-availability-check-profile profile-name;
    }
    servers {
      server-name {
      address ip4-address <port port-number> <transport (udp | tcp)>;
      admission-control profile-name;
      availability-check-profile profile-name;
      service-point service-point-name;
    }
  timers {
    inactive-call seconds;
    timer-c seconds;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 9.4.
data-inactivity-detection option introduced in Junos 9.6.
message-manipulation-rules option introduced in Junos OS Release 9.6.

Description Configure SIP policies and timers.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- SIP Routing Overview in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sip-header

Syntax

```
sip-header header-field-name {
  field-value {
    add field-value;
    add-missing field-value;
    add-overwrite field-value;
    modify-regular-expression regular-expression with field-value;
    reject-regular-expression regular-expression;
    remove-all;
    remove-regular-expression regular-expression;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip message-manipulation-rules manipulation-rule *rule-name* actions]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the values of the header fields that you want to manipulate in SIP messages. You can have up to five of each of the field value definitions in each SIP header configuration. For example, up to five **modify-regular-expression** field values or up to five **add-missing** field values.

Options

header-field-name—Name of the header field in SIP headers for which you want to define field values.

add ***field-value***—Adds an instance of the header field with the field value that you define. If the header field already exists, the software creates a new instance of the header field and inserts it before any existing instance of the header field. Having more than one field value is not allowed for some header fields.

add-missing ***field-value***—Adds a new header field with the field value that you define if the header field is missing from the SIP header.

add-overwrite ***field-value***—Adds a new header field with the field value that you define if the header field is missing from the SIP header. If the header field already exists, its field value is overwritten with the new field value. The software overwrites the field value in all instances of the header field.

modify-regular-expression—Changes the value of a regular expression.

Syntax: modify-regular-expression ***regular-expression*** **with** ***field-value***—Enter the regular expression that you want to modify followed by the value with which you want to replace the regular expression. In the following example, regular expression 1800 is replaced with 555:

```
modify-regular-expression 1800 with 555;
```

remove-all—Removes all instances of the header field.

remove-regular-expression ***regular-expression***—Removes all of the header fields that have field values that match this regular expression.

Copyright © 2013, Juniper Networks, Inc.

reject-regular-expression *regular-expression*—Rejects SIP messages and terminates the usage that the message is part of if the header field contains the regular expression.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Manipulation of Headers and Request URIs in SIP Messages in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Message Manipulation Rules in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

sip-stack

Syntax

```
sip-stack {
  dev-logging;
  event-tracing;
  ips-tracing;
  pd-log-detail (full | summary);
  pd-log-level (audit | exception | problem);
  per-tracing;
  verbose-logging;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in Junos OS Release 9.4.

Description Set trace options for the SIP stack component of the BSG.

Options

dev-logging—Configure development tracing for the stack.

event-tracing—Activate or deactivate the stack's event tracing.

ips-tracing—Activate or deactivate the stack's IPS tracing.

pd-log-detail—Specify the amount of detail to be sent to the log file.

- **full**—All available information is sent to the log file.
- **summary**—The type of logging, the identifier and the first line of the log message are sent to the log file.

pd-log-level—Specifies which types of PD logs are printed to the log file. This can be set to:

- **problem**—Problem log messages are sent to the log file.
- **exception**—Exception and problem log messages are sent to the log file.
- **audit**—All log messages are sent to the log file.

This option determines the levels of log messages to be sent to the log file. Selecting a level causes messages at that level and any higher levels to be sent to the log file. For example, setting this option to **exception** causes both exception and problem logs to be sent to the log file. Setting it to **audit** causes all logs to be sent to the log file. The default value is **audit**.

per-tracing—Activate or deactivate the stack's performance tracing.

verbose-logging—Configure verbose tracing for the stack.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

term

See the following sections:

- term (New Call Usage Policy) on page 735
- term (New Transaction Policy) on page 736
- term (Service Class) on page 737

term (New Call Usage Policy)

Syntax

```
term term-name {
  from {
    contact [ contact-fields ];
    method {
      method-invite;
    }
    request-uri [ uri-fields ];
    source-address [ ip-addresses ];
  }
  then {
    media-policy {
      data-inactivity-detection {
        inactivity-duration seconds;
      }
      media-release;
      no-anchoring;
      service-class service-class-name;
    }
    trace;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-call-usage-policy *policy-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Define the new call usage policy term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

term (New Transaction Policy)

Syntax

```
term term-name {
  from {
    contact {
      registration-state [ registered | not-registered ];
      regular-expression [ regular-expression ];
      uri-hiding [ hidden-uri | not-hidden-uri ];
    }
    method {
      method-invite;
      method-message;
      method-options;
      method-publish;
      method-refer;
      method-register;
      method-subscribe;
    }
    request-uri {
      registration-state [ registered | not-registered ];
      regular-expression [ regular-expression ];
      uri-hiding [ hidden-uri | not-hidden-uri ];
    }
  source-address [ ip-addresses ];
  }
  then {
    (accept | reject);
    admission-control admission-control-profile;
    message-manipulation {
      forward-manipulation {
        manipulation-rule-name;
      }
      reverse-manipulation {
        manipulation-rule-name;
      }
    }
    on-3xx-response{
      recursion-limit number;
      }
    }
    route (Services) {
      next-hop (Services) (request-uri | address ipv4-address | <port port-number> |
        <transport-protocol (udp | tcp)>);
      egress-service-point service-point-name;
      server-cluster cluster-name;
    }
    signaling-realm signaling-realm;
    trace;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Copyright © 2013, Juniper Networks, Inc.

message-manipulation introduced in Junos OS Release 9.6.

Description Define the new transaction policy term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*
- Using New Transaction Policies to Route SIP Requests (next-hop) in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring Routing of VPN Calls in *Session Border Control Solutions Guide Using BGF and IMSG*

term (Service Class)

Syntax
```
term term-name {
  from {
    media-type (any-media | audio | video);
  }
  then {
    committed-information-rate bytes-per-second;
    committed-burst-size bytes;
    dscp (alias | 6-bit-pattern | decimal-value | hexadecimal-value);
    reject;
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* embedded-spdf service-class *service-class-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the service class term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

then

See the following sections:

- then (New Call Usage Policy) on page 738
- then (New Transaction Policy) on page 739
- then (Service Class) on page 740

then (New Call Usage Policy)

Syntax

```
then {
  media-policy {
    data-inactivity-detection (Services PGCP) {
      inactivity-duration seconds;
    }
    media-release;
    no-anchoring;
    service-class service-class-name;
  }
  trace;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-call-usage-policy *policy-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Define the actions performed on incoming requests that match the new call usage policy.

Options **trace**—Trace messages are accepted by this policy.

The remaining options are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

then (New Transaction Policy)

Syntax

```
then {
  (accept | reject);
  admission-control admission-control-profile;
  message-manipulation {
    forward-manipulation {
      manipulation-rule-name;
    }
    reverse-manipulation {
      manipulation-rule-name;
    }
  }
  on-3xx-response {
    recursion-limit number;
    }
  }
  route (Services) {
    egress-service-point service-point-name;
    next-hop (Services) (request-uri | address ipv4-address | <port port-number> |
      <transport-protocol (udp | tcp)>);
    server-cluster cluster-name;
  }
  signaling-realm signaling-realm;
  trace;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 9.4.
message-manipulation option introduced in Junos OS Release 9.6.
on-3xx-response option introduced in Junos OS Release 10.2.
signaling-realm option introduced in Junos OS Release 10.2.

Description Define the actions performed on incoming requests that match this policy.

Options

accept—Accept the traffic and send it to its destination.

admission-control ***controller-name***—Accept or reject the traffic based on admission control configured for ***controller-name***.

signaling-realm *realm-name*—Name of the signaling realm to be assigned to traffic that matches the **from** selection criteria.

reject—Do not accept the traffic and return a rejection message. You can log or sample rejected traffic.

trace—Trace messages are accepted by this policy.

The remaining options are explained separately.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a New Transaction Policy in *Session Border Control Solutions Guide Using BGF and IMSG*

then (Service Class)

Syntax

```
then {
  committed-burst-size bytes;
  committed-information-rate bytes-per-second;
  dscp (dscp-value | alias | do-not-change);
  reject;
}
```

Hierarchy Level
[edit services border-signaling-gateway gateway *gateway-name* embedded-spdf service-class *service-class-name* term *term-name*]

Release Information
Statement introduced in Junos OS Release 9.4.

Description
Define the actions performed on traffic that matches the service class.

Options
reject—Do not accept the traffic and return a rejection message. Rejected traffic can be logged or sampled.

The remaining options are described separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring QoS and Rate Limiting in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

timer-c

Syntax timer-c *seconds*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip timers]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure Timer C, an INVITE transaction timeout. The timer tracks the duration of time waiting for a final response to an INVITE request, ensuring that resources are released if the timer expires. When Timer C expires, a CANCEL is sent to the caller and a 408 error message (Request timeout) is sent to the call recipient.

Options *seconds*—Duration of the timeout period.

Range: 180 to 300 seconds

Default: 180 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- SIP Timers Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring SIP Timers in *Session Border Control Solutions Guide Using BGF and IMSG*

timers

Syntax

```
timers {
  inactive-call seconds;
  timer-c seconds;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure timers used to issue SIP timeouts.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- SIP Timers Overview in *Session Border Control Solutions Guide Using BGF and IMSG*
- Configuring SIP Timers in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Services Border Gateway Signaling)

Syntax

```
traceoptions {
  file <filename> <files files> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag {
    datastore {
      data trace-level;
      db trace-level;
      handle trace-level;
      minimum trace-level;
    }
    framework {
      action trace-level;
      event trace-level;
      executor trace-level;
      freezer trace-level;
      minimum trace-level;
      memory-pool trace-level;
    }
    minimum trace-level;
    sbc-utils {
      common trace-level;
      configuration trace-level;
      device-monitor trace-level;
      ipc trace-level;
      memory-management trace-level;
      message trace-level;
      minimum trace-level;
      user-interface trace-level;
    }
    session-trace trace-level;
    signaling {
      b2b trace-level;
      b2b-wrapper trace-level;
      minimum trace-level;
      policy trace-level;
      sip-stack-wrapper trace-level;
      topology-hiding trace-level;
      ua trace-level;
    }
    sip-stack {
      dev-logging;
      event-tracing;
      ips-tracing;
      pd-log-detail (full | summary);
      pd-log-level (audit | exception | problem);
      per-tracing;
      verbose-logging;
    }
  }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Copyright © 2013, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure border signaling gateway tracing operations. The messages are output to **/var/log/**.

Options Options are described separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing BSG Operations in *Session Border Control Solutions Guide Using BGF and IMSG*

transactions

Syntax

```
transactions {
  maximum-concurrent number;
  committed-attempts-ratetransactions-per-second;
  committed-burst-size number-of-transactions;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* admission-control]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure admission control settings for out-of-dialog-transactions.

Options

maximum-concurrent *number*—Maximum number of concurrent transactions. 0 causes all calls to be rejected.

Values: 0 through 100,000

Default: 100000

committed-attempts-rate *transactions-per-second*—Maximum number of attempts per second to initiate an out-of-dialog transaction.

Values: 0 though 100

Default: 100

committed-burst-size*number-of-transactions*—Maximum number of transactions allowed to burst above the committed-rate and still be accepted.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Admission Control Profiles in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

transport-details

Syntax transport-details port *port-number* ip-address *ip-address* [tdp |udp];

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* service-point *service-point-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the transport parameters for a service point. The transport parameters consist of a combination of port number, IP address, and transport protocol. Policies are applied only to incoming requests that match the transport parameters. You can configure only one set of transport parameters for each service point.

Options ***port-number***—Port number on which you want to match incoming traffic.

Range: 0 through 65,535

Default: 5060

ip-address—IP address on which you want to match incoming messages. If you do not define an IP address, the software uses the IP address of the service interface assigned to this service point.

Values: upd or tcp

Default: udp

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Service Point in *Session Border Control Solutions Guide Using BGF and IMSG*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 34

PTSP Configuration Guidelines

To configure the static policies for the packet-triggered subscribers and policy control (PTSP) feature, include the **ptsp** statement at the **[edit services]** hierarchy level:

```
[edit services]
ptsp {
  rule rule-name {
    count-type (application | rule);
    demux (destination-address | source-address);
    forward-rule forward-rule-name;
    match-direction (input | input-output | output);
    term precedence {
      from {
        application-group-any;
        application-groups [ application-group-name ];
        applications [ application-name ];
        local-port-range low low-value high high-value;
        local-ports [ value-list ];
        protocol protocol-number;
        remote-address address <except>;
        remote-address-range low low-value high high-value <except>;
        remote-port-range low low-value high high-value;
        remote-ports [ value-list ];
        remote-prefix-list prefix-list-name <except>;
      }
      then {
        (accept | discard);
        count (application | application-group | application-group-any | rule | none);
        forwarding-class forwarding-class;
        police policer-name;
      }
    }
  }
  rule-set rule-set-name {
    rule rule-name;
  }
  forward-rule rule-name {
    term precedence {
      from {
        application-groups [ application-group-name ];
        applications [ application-name ];
        local-address address <except>;
        local-address-range low low-value high high-value <except >;
```

Copyright © 2013, Juniper Networks, Inc.

```
            local-prefix-list prefix-list-name <except >;
        }
        then {
            forwarding-instance forwarding-instance unit-number unit-number;
        }
    }
  }
}
```



NOTE: Starting with Junos OS Release 12.1, offloading is supported by the packet-triggered subscribers and policy control (PTSP) plug-in. Offloading is accomplished using the Juniper Forwarding Mechanism (JFM) on all MX Series routers with Modular Port Concentrators (MPCs). The **show services subscriber flows client-id** *client-id* command displays the offload status. For more information about offloading, see "Local Policy Decision Function Configuration Guidelines" on page 859

For information about using the PTSP statements to configure the PTSP feature, see the Junos OS Subscriber Management, Release 13.1.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 35

Summary of PTSP Configuration Statements

application-group-any

Syntax application-group-any;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify that any application group defined in the database is considered a match.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

application-groups

Syntax application-group [*application-group-name*];

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence* from]
[edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify one or more application groups defined in the application identification configuration for inclusion as a match condition.

Options ***application-group-name***—Identifier of the application group.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

applications (Services PTSP)

Syntax applications [*application-name*];

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence* from]
[edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify one or more applications defined in the application identification configuration for inclusion as a match condition.

Options *application-name*—Identifier of the application.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

count-type

Syntax count-type (application | rule);

Hierarchy Level [edit services ptsp rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the statistics aggregation, collection, and reporting style for this rule. Terms and rules cannot mix and match different styles. All service rules attached to a given service set must have the same style.

Options **application**—Report statistics in a flat file and aggregate them by application for one of the following:

- An application, where the count action **application** is specified in the term.
- An application group, where the count action **application-group** is specified in the term.
- All application groups, where the count action **application-group-any** is specified in the term.

rule—Aggregate statistics for the service rule. The statistics are reported by Diameter. All count actions in all terms for the rule must specify **rule**.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

demux

Syntax demux (destination-address | source-address);

Hierarchy Level [edit services ptsp rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the IP address used to establish the subscriber context. Subscriber instantiation is always triggered for ingress packets, so this value indicates which IP address in the ingress packets for the flow is used. If the IP address does not correspond to a known subscriber, then a new subscriber context is created. All service rules attached to a given service set must have the same setting.

Options

destination-address—Use the destination IP address field of the ingress packet header for the flow.

source-address—Use the source IP address field of the ingress packet header for the flow.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

forward-rule (Configuring)

Syntax

```
forward-rule forward-rule-name {
  term precedence {
    from {
      application-groups [ application-group-name ];
      applications [ application-name ];
      local-address address <except>;
      local-address-range low low-value high high-value <except >;
      local-prefix-list prefix-list-name <except >;
    }
    then {
      forwarding-instance forwarding-instance;
      unit-number unit-number;
    }
  }
}
```

Hierarchy Level [edit services ptsp]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the forwarding instance for a specific subscriber or set of subscribers based on the IP address, network, or prefix list. The rule match is applied on the input side.

Options *forward-rule-name*—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

forward-rule (Including in Rule)

Syntax forward-rule *forward-rule-name*;

Hierarchy Level [edit services ptsp rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify the forwarding instance for inclusion in a rule.

Options ***forward-rule-name***—Identifier for the forward rule that specifies the forwarding instance.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

from (Forward Rule)

Syntax
```
from {
  application-groups [ application-group-name ];
  applications [ application-name ];
  local-address address <except >;
  local-address-range low low-value high high-value <except >;
  local-prefix-list prefix-list-name <except >;
}
```

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify match conditions for the PTSP term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

from (Rule)

Syntax

```
from {
  application-group-any;
  application-groups [ application-group-name ];
  applications [ application-name ];
  local-port-range low low-value high high-value;
  local-ports [ value-list ];
  protocol protocol-number;
  remote-address address <except >;
  remote-address-range low low-value high high-value <except >;
  remote-port-range low low-value high high-value;
  remote-ports [ value-list ];
  remote-prefix-list prefix-list-name <except >;
}
```

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify match conditions for the PTSP term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

local-address

Syntax local-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the address for rule matching. Local address values are matched against a source or destination IP address for the flow depending on the configured value for the **demux** statement. If you do not specify an address, then any local address matches this term. If you do not specify a prefix value, then a host mask is the default.

Options ***address***—IPv4 or IPv6 address or prefix value.

any-unicast—Match all unicast addresses.

except—(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1
- demux on page 749

Copyright © 2013, Juniper Networks, Inc.

local-address-range

Syntax local-address-range low *low-value* high *high-value* <except>;

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the address range for rule matching. Local address values are matched against a source or destination IP address for the flow depending on the configured value for the **demux** statement. If you do not specify an address, then any local address matches this term.

Options ***low-value***—Lower boundary for the IPv4 or IPv6 address range.

high-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exclude the specified address range from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1
- demux on page 749

local-port-range

Syntax local-port-range low *low-value* high *high-value*;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the port range for rule matching.

Options ***low-value***—Lower boundary for the port range.

high-value—Upper boundary for the port range.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

local-ports

Syntax local-ports [*port-numbers*];

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify one or more ports for inclusion as a match condition.

Options ***port-numbers***—Port number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

local-prefix-list

Syntax local-prefix-list *prefix-list-name* <except>;

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***prefix-list-name***—Prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

match-direction (Services PTSP)

Syntax match-direction (input | input-output | output);

Hierarchy Level [edit services ptsp rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the direction in which the rule match is applied.

Options **input**—Apply the rule match on the input side of the interface.

input-output—Apply the rule match bidirectionally.

output—Apply the rule match on the output side of the interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

protocol

Syntax protocol *protocol-number*;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify the protocol for inclusion as a match condition.

Options ***protocol-number***—Protocol number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

remote-address

Syntax remote-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the address for rule matching. Remote address values are matched against a destination or source IP address for the flow depending on the configured value for the **demux** statement. If you do not specify an address, then any remote address matches this term. If you do not specify a prefix value, then a host mask is the default.

Options ***address***—IPv4 or IPv6 address or prefix value.

any-unicast—Match all unicast addresses.

except—(Optional) Exclude the specified address, prefix, or unicast packets from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1
- demux on page 749

Copyright © 2013, Juniper Networks, Inc.

remote-address-range

Syntax remote-address-range low *low-value* high *high-value* <except>;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the address range for rule matching. Remote address values are matched against a destination or source IP address for the flow depending on the configured value for the **demux** statement. If you do not specify an address, then any remote address matches this term.

Options ***low-value***—Lower boundary for the IPv4 or IPv6 address range.

high-value—Upper boundary for the IPv4 or IPv6 address range.

except—(Optional) Exclude the specified address range from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1
- demux on page 749

remote-port-range

Syntax remote-port-range low *low-value* high *high-value*;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the port range for rule matching.

Options ***low-value***—Lower boundary for the port range.

high-value—Upper boundary for the port range.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

remote-ports

Syntax remote-ports [*port-numbers*];

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify one or more ports for inclusion as a match condition.

Options ***port-numbers***—Port number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

remote-prefix-list

Syntax remote-prefix-list *prefix-list-name* <except>;

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence* from]

Release Information Statement introduced in Junos OS Release 10.2.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***prefix-list-name***—Prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

rule (Configuring)

Syntax

```
rule rule-name {
  count-type (application | rule);
  demux (destination-address | source-address);
  forward-rule forward-rule-name;
  match-direction (input | input-output | output);
  term precedence {
    from {
      application-group-any;
      application-groups [ application-group-name ];
      applications [ application-name ];
      local-port-range low low-value high high-value;
      local-ports [ value-list ];
      protocol protocol-number;
      remote-address address <except>;
      remote-address-range low low-value high high-value <except>;
      remote-ports [ value-list ];
      remote-port-range low low-value high high-value;
      remote-prefix-list prefix-list-name <except>;
    }
    then {
      (accept | discard);
      count (application | application-group | application-group-any | rule | none);
      forwarding-class forwarding-class;
      police policer-name;
    }
  }
}
```

Hierarchy Level [edit services ptsp]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the rule the router uses when applying this service.

Options *rule-name*—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

rule (Including in Rule Set)

Syntax rule *rule-name*;

Hierarchy Level [edit services ptsp rule-set *rule-set-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the rule the router uses when applying this service.

Options ***rule-name***—Identifier for the collection of terms that constitute this rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

rule-set (Services PTSP)

Syntax
```
rule-set rule-set-name {
    [rule rule-names ];
}
```

Hierarchy Level [edit services ptsp]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

services (PTSP)

Syntax services ptsp { ... }

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define the services to be applied to traffic.

Options **ptsp**—Identify the values configured for PTSP matching rules.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

term (Forward Rule)

Syntax

```
term precedence {
  from {
    application-groups [ application-group-name ];
    applications [ application-name ];
    local-address address <except>;
    local-address-range low low-value high high-value <except>;
    local-prefix-list prefix-list-name <except>;
  }
  then {
    forwarding-instance forwarding-instance;
    unit-number unit-number;
  }
}
```

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define the term properties for the forward rule.

Options ***precedence***—Precedence value for this term in relation to other terms. Term with lowest precedence is evaluated first.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

term (Rule)

Syntax

```
term precedence {
  from {
    application-group-any;
    application-groups [ application-group-name ];
    applications [ application-name ];
    local-port-range low low-value high high-value;
    local-ports [ value-list ];
    protocol protocol-number;
    remote-address address <except>;
    remote-address-range low low-value high high-value <except>;
    remote-port-range low low-value high high-value;
    remote-ports [ value-list ];
    remote-prefix-list prefix-list-name <except>;
  }
  then {
    (accept | discard);
    count (application | application-group | application-group-any | rule);
    forwarding-class forwarding-class;
    police policer-name;
  }
}
```

Hierarchy Level [edit services ptsp rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define the term properties for the PTSP rule.

Options ***precedence***—Precedence value for this term in relation to other terms. Term with lowest precedence is evaluated first.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

then (Forward Rule)

Syntax

```
then {
  forwarding-instance forwarding-instance;
  unit-number unit-number;
}
```

Hierarchy Level [edit services ptsp forward-rule *forward-rule-name* term *precedence*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define the term actions for the forward rule.

Options

forwarding-instance—Identifier for the forwarding instance for packet flows accepted under this policy.

unit-number—Unit number associated with the forwarding instance.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

then (Rule)

Syntax

```
then {
  (accept | discard);
  count (application | application-group | application-group-any | rule);
  forwarding-class forwarding-class;
  police policer-name;
}
```

Hierarchy Level [edit services ptsp rule *rule-name* term *precedence*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define the term actions. You can configure the router to accept or discard the targeted traffic. The action modifiers (count and forwarding-class) are optional.

Options You can configure one of the following actions:

- **accept**—Accept the packets and all subsequent packets in flows that match the rules.
- **discard**—Discard the packet and all subsequent packets in flows that match the rules.

When you select **accept** as the action, you can optionally configure one or both of the following action modifiers. No action modifiers are allowed with the **discard** action.

- **count (application | application-group | application-group-any | rule | none)**—For all accepted packets that match the rules, record a packet count using PTSP statistics practices. You can specify one of the following options; there is no default setting:
  - **application**—Count the application that matched in the **from** clause.
  - **application-group**—Count the application group that matched in the **from** clause.
  - **application-group-any**—Count all application groups that match **from application-group-any** under the **any** group name.
  - **rule**—Count the rule that matched in the **from** clause.
  - **none**—Same as not specifying **count** as an action.
- **forwarding-class *forwarding-class***—Specify the forwarding class name for outgoing packets.

When you include a policer, the only allowed action is **discard**. For more information on policers, see the Routing Policy Configuration Guide.

- **police *policer-name***—Apply rate-limiting properties to the traffic as configured at the **[edit firewall policer *policer-name*]** hierarchy level. This configuration allows bit-rate and burst-size attributes to be applied to the traffic that are not supported by PTSP rules.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Static PTSP Rules in Junos OS Subscriber Management, Release 13.1

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 3

Dynamic Application Awareness for Junos OS

- Dynamic Application Awareness for Junos OS Overview on page 771
- Application Identification Configuration Guidelines on page 781
- Summary of Application Identification Configuration Statements on page 799
- Application-Aware Access List Configuration Guidelines on page 839
- Summary of AACL Configuration Statements on page 847
- Local Policy Decision Function Configuration Guidelines on page 859
- Summary of L-PDF Configuration Statements on page 867

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 36

Dynamic Application Awareness for Junos OS Overview

This chapter describes several related features that support application-level filtering and per-subscriber, per-application group bandwidth control as an extension of Intrusion Detection and Prevention (IDP). In addition to IDP, the main components are application identification (APPID), application-aware access list (AACL) services, and local policy decision functionality for application-related services (L-PDF).



NOTE: Starting with Junos OS Release 12.1, all interface-style services are supported for dynamic Point-to-Point Protocol over Ethernet (PPPoE) subscribers on all MX Series routers with modular Modular Port Concentrators (MPCs).

Copyright © 2013, Juniper Networks, Inc.

NOTE: Because the Services SDK framework lacks aggressive constraint checks, you should not set the **policy-db-size** statement at the **[edit chassis fpc** *slot-number* **pic** *pic-number* **adaptive-services service-package extension-provider]** hierarchy level to a high value. For dynamic application awareness configurations, the recommended values for the Services SDK options at this hierarchy level are as follows:

- **control-cores** = 1
- **data-cores** = 7
- **object-cache-size** = 1280 (for Multiservices 400 PIC and Multiservices DPC)
- **policy-db-size** = 200
- Include these **package** values: **jservices-idp**, **jservices-appid**, **jservices-llpdf**, **jservices-aacl**

For more information about this configuration, see the following topics in the *SDK Applications Configuration Guide and Command Reference*:

- Configuring Control and Data Cores
- Configuring Memory Settings
- Configuring Packages on the PIC

NOTE: In the export version of JUNOS, signature download is not expected to work for AppID and IDP features in the Dynamic Application Awareness (DAA) suite. In order to make it work, you must additionally install the Crypto Software Suite.

This chapter includes the following:

- IDP Overview on page 772
- APPID Overview on page 774
- AACL Overview on page 775
- L-PDF Overview on page 776
- Configuring Multiple IDP Detectors on page 777
- Best-Effort Application Identification of DPI-Serviced Flows on page 777

IDP Overview

The Dynamic Application Awareness for the Junos OS set of services adds support for the intrusion detection and prevention (IDP) functionality using Deep Packet Inspection (DPI) technology to Juniper Networks MX Series 3D Universal Edge Routers equipped with Multiservices Dense Port Concentrators (MS-DPCs) and M120 or M320 Multiservices Edge Routers equipped with Multiservices 400 PICs and Aggregated Multiservices (AMS)

Copyright © 2013, Juniper Networks, Inc.

PICs. Aggregated Multiservices PICs (ams- interfaces) enable multiple ms- interfaces to be grouped together in a single bundle and cause the traffic destined for this AMS group to be distributed over the member services PICs of the group. Junos OS Trio chipsets enable the calculation of a symmetric hash for the forward and reverse flows, and support a microcode map in the forwarding plane. This capability enables load-balancing of traffic across various services PICs in an AMS group. Starting with Junos OS Release 12.1, ams- interfaces enable an N:1 redundancy mechanism to cluster together N number of ms- interfaces in an AMS group that supports load sharing.



NOTE: For ams- interfaces and rms- interfaces, the statistics data in the bulk statistics file is collected using the reports received from the MS PICs. For the ams- interfaces, the retrieval and storage of statistics is not possible because of multiple PICs containing statistics data for the same subscriber. For interfaces in an AMS group, statistics data from different MS PICs in the AMS group are collected and aggregated on the Routing Engine where a timer control is activated and the data is saved in the bulkstats file based on this timer. This method of collection causes the statistics data in the bulkstats file to be displayed with a small delay period.

The IDP functionality is already supported on Juniper Networks J Series Services Routers and SRX Series Services Gateways running the Junos OS and is described in the *Junos OS Security Configuration Guide*. Starting with Junos OS Release 11.3, support for the IDP functionality is extended to T320, T640, and T1600 routers. In addition, multiple IDP detectors are now supported on the M120, M320, and MX Series routers with Enhanced III Flexible PIC Concentrators (FPCs).

The same CLI statements and commands are used on all platforms with the following caveats:

- **Service sets**—IDP is incorporated as a component of service sets only on the specified Juniper Networks T Series, M Series and MX Series routers. IDP depends on application identification services (APPID) for definition and detection of some Layer 7 applications. Before configuring an IDP policy, you must download the APPID application package. Only one service set can be applied to a single interface when the APPID functionality is used.
- **Multiple IDP detectors**—Except for the maximum number of decoder binary instances (4) that are loaded into the process space, multiple IDP detectors on the M120, M320, and MX Series routers function in a similar way to the existing IDP detector support on J Series and SRX Series devices. To view the current policy and the corresponding detector version, use the **show security idp status detail** command.

To configure IDP properties, include statements at the **[edit security idp]** hierarchy level. In general, you configure IDP processes by including the **idp-policy** statement at the **[edit system processes]** hierarchy level. For use in T Series, M Series and MX Series applications, you then reference this configuration by including the **idp-profile** statement at the **[edit services service-set]** hierarchy level. To configure SNMP IDP objects, include the **idp** statement at the **[edit snmp health-monitor]** hierarchy level. The operational commands

Copyright © 2013, Juniper Networks, Inc.

for monitoring and regulating IDP activity are the **clear security idp**, **request security idp**, and **show security idp** commands.

To configure the source IP address for downloading security packages, use the command **set security idp security-package source-address *ip-address*** because it is not possible to download security packages if the router uses private addressing on its outgoing interface. The source address should be a valid IP address on the node.

NOTE: On T Series, M Series and MX Series routers, the IDP **ip-action** statement is supported on TCP, UDP, and ICMP flows. When the **ip-action target is service**, the ip-action flow is applied if the traffic matches the values specified for the source port, destination port, source address, and destination address. However, for ICMP flows, the destination port is 0, so that any ICMP flow matching the source port, source address, and destination address would be blocked. For more information about the **ip-action** statement, see the *Junos OS CLI Reference*.

When the Multiservices PIC configured for a service set is either administratively taken offline or undergoes a failure, all the traffic entering the configured interface with an IDP service set would be dropped without notification. To avoid this traffic loss, include the **bypass-traffic-on-pic-failure** statement at the **[edit services service-set *service-set-name* service-set-options]** hierarchy level. When this statement is configured, the affected packets are forwarded in the event of a Multiservices PIC failure or offlining, as though interface-style services were not configured.

NOTE: Data channel applications for protocols such as FTP, TFTP, RTSP, and SIP are not in the same application group as their control channel applications. For example, control channel application **junos:ftp** is in the group **junos:file-server** but the corresponding data application **junos:system:ftp-data** is not in any group.

Related Documentation

- Configuring Multiple IDP Detectors on page 777

APPID Overview

The APPID feature identifies applications as constituents of application groups in TCP/UDP/ICMP traffic. It is supported on MX Series routers equipped with Multiservices DPCs and on M120 or M320 routers equipped with Multiservices 400 PICs and Aggregated Multiservices (AMS) PICs. Aggregated Multiservices PICs (ams- interfaces) enable multiple ms- interfaces to be grouped together in a single bundle and cause the traffic destined for this AMS group to be distributed over the member services PICs of the group. Junos OS Trio chipsets enable the calculation of a symmetric hash for the forward and reverse flows, and support a microcode map in the forwarding plane. This capability enables load-balancing of traffic across various services PICs in an AMS group. Starting

Copyright © 2013, Juniper Networks, Inc.

with Junos OS Release 12.1, ams- interfaces enable an N:1 redundancy mechanism to cluster together N number of ms- interfaces in an AMS group that supports load sharing.



NOTE: For ams- interfaces and rms- interfaces, the statistics data in the bulk statistics file is collected using the reports received from the MS PICs. For the ams- interfaces, the retrieval and storage of statistics is not possible because of multiple PICs containing statistics data for the same subscriber. For interfaces in an AMS group, statistics data from different MS PICs in the AMS group are collected and aggregated on the Routing Engine where a timer control is activated and the data is saved in the bulkstats file based on this timer. This method of collection causes the statistics data in the bulkstats file to be displayed with a small delay period.

To configure APPID, include statements at the **[edit services application-identification]** hierarchy level to specify parameter values for defining applications, enable or disable application rules, and gather the applications and rules into groups.

The following are related operational commands:

- **show/clear application-identification application-system-cache**
- **show/clear application-identification counters**

For more information on the CLI configuration, see the Application Identification. For more information on the operational commands, see the Junos OS Operational Mode Commands.

AACL Overview

The application-aware access list (AACL) service adds support for a new service that uses application names and groups as matching criteria for filtering traffic. AACL is a stateless, rules-based service that must be combined with application identification to enable policies to be applied to flows based on application and application group membership in addition to traditional packet matching rules. It is supported on MX Series routers equipped with Multiservices DPCs and on M120 or M320 routers equipped with Multiservices 400 PICs. Starting with Junos OS Release 11.3, AACL is supported on T320, T640, and T1600 routers also.

AACL is configured in a similar way to other rules-based services such as Network Address Translation (NAT), class of service (CoS), and stateful firewall. To configure AACL, include rule specifications for match criteria and actions at the **[edit services aacl]** hierarchy level. You can chain AACL rules along with other service rules by including them in a service-set definition at the **[edit services service-set]** hierarchy level, as previously documented.

There is one pair of related operational commands, **show/clear application-aware-access-list statistics**.

For more information on the CLI configuration, see the Application-Aware Access List. For more information on the operational command, see the Junos OS Operational Mode Commands.

Copyright © 2013, Juniper Networks, Inc.

L-PDF Overview

Local policy decision functionality for application-related services adds support for a new process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces. This functionality is collectively named the local policy decision function (L-PDF). L-PDF is supported on:

- MX Series routers equipped with Multiservices DPCs.
- M120 or M320 routers equipped with Multiservices 400 PICs.
- Aggregated Multiservices (AMS) PICs.

Multiple ms- interfaces can be bundled together in an AMS PIC interface, which causes the traffic destined for this AMS group to be distributed over the member services PICs of the group. Junos OS Trio chipsets enable the calculation of a symmetric hash for the forward and reverse flows, and support a microcode map in the forwarding plane. This capability enables load-balancing of traffic across various services PICs in an AMS group. Starting with Junos OS Release 12.1, ams- interfaces enable an N:1 redundancy mechanism to cluster together N number of **ms- interfaces** in an AMS group that supports load sharing.

Starting with Junos OS Release 11.3, local L-PDF that resides on the services PIC is supported on T320, T640, and T1600 routers. The application identification (APPID) service defines the applications and how they are grouped. The application-aware access list (AACL) service defines the applications and application groups for which statistics are collected for a specific user or interface. The L-PDF configuration defines the way in which the statistics are output.

To configure properties for statistics output, include the **policy-decision-statistics-profile** statement at the **[edit accounting-options]** hierarchy level. A new **traceoptions** configuration is available at the **[edit system services local-policy-decision-function]** hierarchy level. To configure a dynamic profile to attach a specified service set to an interface, include the **service** statement at the **[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level. To attach a service set to a static interface, include the **service-set *service-set-name*** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output)]** hierarchy level. For more information on service sets, see Service Set Properties.

The following related operational commands are supported:

- **show services local-policy-decision-function flows**
- **show/clear services local-policy-decision-function statistics**
- **show/clear services application-aware-access-list statistics**

For more information on the CLI configuration, see the Local Policy Decision Function. For more information on the operational commands, see the Junos OS Operational Mode Commands.

Copyright © 2013, Juniper Networks, Inc.

Configuring Multiple IDP Detectors

To configure multiple IDP detectors:

1. In configuration mode, go to the **[edit security]** hierarchy level:

   user@host# edit security

2. Go to the **[edit security idp sensor-configuration flow]** hierarchy level:

   [edit security]
   user@host# edit idp sensor-configuration flow

3. Configure the **no-reset-on-policy** statement:

   [edit security idp sensor-configuration flow]
   user@host# set no-reset-on-policy

4. Verify the configuration:

```
[edit security]
user@host# show security
idp {
    sensor-configuration {
        flow {
            no-reset-on-policy;
        }
    }
}
```

Related Documentation

- IDP Overview on page 772

Best-Effort Application Identification of DPI-Serviced Flows

This topic describes the following information:

- Features that Support Application-Level Filtering on page 777
- Best-Effort Application Determination on page 778
- APPID, AACL, and L-PDF Processing in Preconvergence Scenarios on page 778

Features that Support Application-Level Filtering

On MX Series routers equipped with Multiservices DPCs and M120 or M320 routers equipped with Multiservices 400 PICs, Intrusion Detection and Prevention (IDP) is accomplished by Deep Packet Inspection (DPI) of TCP, UDP, and ICMP flows. The application identification (APPID) feature defines applications as members of application groups in TCP/UDP/ICMP traffic. IDP depends on APPID for identification and detection of some Layer 7 applications.

The application-aware access list (AACL) service uses application names and groups as matching criteria for filtering traffic. The service defines the applications and application groups for which statistics are collected for a specific user or interface.

Copyright © 2013, Juniper Networks, Inc.

The local policy decision function (L-PDF) enables you to configure properties for statistics output. L-PDF supports a process that regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

Best-Effort Application Determination

Typically, APPID conclusively determines the Layer 7 application associated with a given DPI-serviced flow. In these cases, the application identification is final. Occasionally, APPID is only able to make an initial, inconclusive determination of the Layer 7 application associated with a given flow. This is referred to as a "best-effort" application identification. In such cases, the APPID process continues processing packets on that flow and might subsequently make a conclusive determination of the application associated with that flow. In some cases of best-effort application identification, the flow ends before a final application determination can be made.

APPID, AACL, and L-PDF Processing in Preconvergence Scenarios

The following sections describe APPID, AACL, and L-PDF processing in various stages of application identification for a DPI-serviced flow of TCP/UDP/ICMP traffic.

- Prior to a Final or Best-Effort Application Identification on page 778
- Upon Best-Effort Application Identification on page 778
- While Application Identification Is on a Best-Effort Basis on page 779
- If a Flow Ends Before an Application Identification Is Made on page 779
- If a Flow Ends While Application Identification on a Best-Effort Basis on page 779

Prior to a Final or Best-Effort Application Identification

During the time that APPID has not yet made either a final or best-effort determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has not yet made either a final or best-effort determination of the associated application:

- **show services local-policy-decision-function flows (interface** ***interface-name*** **| subscriber** ***subscriber-name*****)**
- **show services application-aware-access-list flows (interface** ***interface-name*** **| subscriber** ***subscriber-name*****)**

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays “unknown” for a flow for which APPID has not yet made either a final or best-effort determination of the associated application.

Upon Best-Effort Application Identification

When a best-effort application determination is made, AACL does not apply any AACL term actions configured for that flow. There are a number of reasons for this, one being that the action itself (such as "discard") could make a final application determination impossible. Instead, AACL or L-PDF tracks the flow and accepts all packets for that flow

Copyright © 2013, Juniper Networks, Inc.

until a final determination is made, at which time the normal AACL or L-PDFL actions are fully applied to the flow.

While Application Identification Is on a Best-Effort Basis

During the time that APPID identification of the application associated with a given flow is on a best-effort basis, the flow does not contribute to any per-subscriber or per-application statistics collection.

The output of the following operational mode commands includes flows for which APPID has only made a best-effort determination of the associated application:

- **show services local-policy-decision-function flows (interface *interface-name* | subscriber *subscriber-name*)**
- **show services application-aware-access-list flows (interface *interface-name* | subscriber *subscriber-name*)**

In the command output, the **Action** field displays "accept" and the **Application** or **Application group** field displays “unknown” for a flow for which APPID has only made a best-effort determination of the associated application.

If a Flow Ends Before an Application Identification Is Made

If a flow ends before APPID has made either a final or a best-effort application identification, AACL or L-PDF uses the "unknown" application ID as a final determination and performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** AACL term action is configured for the "application-group-any" application, then the statistics for that flow will be collected and aggregated against the count bucket type, and reported as such.

If a Flow Ends While Application Identification on a Best-Effort Basis

If a flow ends while the application identification is on a best-effort basis, AACL or L-PDF uses that best-effort determination as a final determination. AACL or L-PDF performs any necessary collection, aggregation, and reporting of statistics based on that Layer 7 application. In particular, if the **count** AACL term action is configured for that Layer 7 application, then the statistics for the flow will be collected and aggregated against the AACL or L-PDF statistics. However, in the case of nested applications, AACL and L-PDF will not consider the best-effort determination as final and the nested application will be reported as an unknown application.

Related Documentation

- Configuring AACL Rules on page 840
- Configuring Statistics Profiles on page 859
- aacl-fields on page 868
- aacl-statistics-profile on page 869
- rule on page 852
- services on page 853
- term on page 856

Copyright © 2013, Juniper Networks, Inc.

- then on page 857

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 37

Application Identification Configuration Guidelines

To configure application identification services (APPID), include the **application-identification** statement at the **[edit services]** hierarchy level:

```
[edit services]
application-identification {
  application application-name {
    disable;
    idle-timeout seconds;
    index number;
    session-timeout seconds;
    type type;
    type-of-service service-type;
    port-mapping {
      port-range {
        tcp (port | range);
        udp (port | range);
      }
      disable;
    }
  }
  application-group group-name {
    application-groups {
      name [application-group-name];
    }
    applications {
      name [application-name];
    }
    index number;
    disable;
  }
  application-system-cache-timeout seconds;
  enable-heuristics
  max-checked-bytes bytes;
  min-checked-bytes bytes;
  nested-application
  nested-application-settings
  no-application-identification;
  no-application-system-cache;
  no-clear-application-system-cache;
```

Copyright © 2013, Juniper Networks, Inc.

```
    no-protocol-method;
    no-signature-based;
    profile profile-name {
      [ rule-set rule-set-name ];
    }
    rule rule-name {
      disable;
      address address-name {
        destination {
          ip address</prefix-length>;
          port-range {
            tcp [ ports-and-port-ranges ];
            udp [ ports-and-port-ranges ];
          }
        }
        source {
          ip address</prefix-length>;
          port-range {
            tcp [ ports-and-port-ranges ];
            udp [ ports-and-port-ranges ];
          }
        }
        order number;
      }
      application application-name;
    }
    rule-set rule-set-name {
      rule application-rule-name;
    }
    signature-method-all-ports
    traceoptions {
      file filename <files number> <match regex> <size size> <world-readable |
        no-world-readable>;
      flag flag;
      no-remote-trace;
    }
}

[edit services]
    hcm {
      url-rule url-rule-name {
        term term-num {
          from {
            url-list url-list-name ;
            url url_identifier {
              host hostname ;
              request-url page-name ;
            }
          }
          then {
            discard;
            accept;
            count;
            log-request;
          }
        }
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
        url-rule-set url-rule-set-name {
          url-rule rule1 ;
          url-rule rule2 ;
        }
      }
    }
```

This chapter contains the following sections:

- Defining an Application Identification on page 783
- Configuring APPID Rules on page 785
- Using Stateful Firewall Rules to Identify Data Sessions on page 786
- Configuring Application Profiles on page 788
- Configuring Application Groups on page 788
- Application Identification for Nested Applications on page 789
- Disabling Application Identification for Nested Applications on page 790
- Configuring Global APPID Properties on page 791
- Configuring Automatic Download of Application Package Updates on page 792
- Configuring APPID Support for Heuristics on page 792
- Configuring APPID Support for Unidirectional Traffic on page 793
- Tracing APPID Operations on page 794
- Examples: Configuring Application Identification Properties on page 796

Defining an Application Identification

To configure a specific IP address or port-based application identification, include the **application *application-name*** statement at the **[edit services application-identification]** hierarchy level:

```
application application-name {
  disable;
  idle-timeout seconds;
  index number;
  session-timeout seconds;
  type type;
  type-of-service service-type;
  port-mapping {
    port-range {
      tcp [ ports-and-port-ranges ];
      udp [ ports-and-port-ranges ];
    }
    disable;
  }
}
```

You can include the following general properties in the configuration:

Copyright © 2013, Juniper Networks, Inc.

- **application**—Application name, a required statement; maximum 31 characters. Predefined applications have the prefix **junos-** to avoid conflict with user-defined ones.
- **idle-timeout**—Amount of time that a session remains idle before it is deleted.
- **index**—Application index number in the range from 1 through 65,534, with integers 1 through 1024 reserved for predefined applications.
- **session-timeout**—Lifetime of a session.
- **type**—Well known applications, such as HTTP or FTP.
- **type-of-service**—Type of service, defined by service objective. There is no default value; options are **maximize-reliability**, **maximize-throughput**, **minimize-delay**, and **minimize-monetary-cost**.
- **disable**—Disable this application definition in the APPID service.



NOTE: You can also specify session and idle timeout values globally for a Multiservices interface by including the following statements at the **[edit interfaces *interface-name* services-options]** hierarchy level:

- **inactivity-non-tcp-timeout**—Inactivity timeout period for non-TCP established sessions.
- **inactivity-tcp-timeout**—Inactivity timeout period for TCP established sessions.
- **session-timeout**—Lifetime of a session.
- **disable-global-timeout-override**—Disallow overriding a global inactivity or session timeout.

You can include the following port-mapping properties at the **[edit services application-identification port-mapping]** hierarchy level:

- **port-range**—TCP or UDP port number or numeric range, entered as **[*minimum-value* – *maximum-value*]**. For port-mapping configurations, this entry is required if the parent node exists.
- **disable**—Disable port-mapping properties for this application.



NOTE: For applications with signatures for both client-to-server and server-to-client directions, the APPID for Dynamic Application Awareness must accept the data packets in both directions on the same session to complete the identification process.

For a configuration example, see "Examples: Configuring Application Identification Properties" on page 796.

Copyright © 2013, Juniper Networks, Inc.

Configuring APPID Rules

This configuration specifies the properties for identifying an application for which a source or destination IP address and port is used for a known application, without the requirement of an application signature. For example, the Session Initiation Protocol (SIP) server initiates a session from its identified port, 5060. You can therefore specify the SIP server IP address and port 5060 in the port mapping configuration for the SIP application. The advantage of using this method is to provide efficiency and accuracy of application identification for your network.

To configure application rule properties, include the **rule** statement at the **[edit services application-identification]** hierarchy level:

```
rule rule-name {
  address address-name {
    destination {
      ip address</prefix-length>;
      port-range {
        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
      }
    }
    source {
      ip address</prefix-length>;
      port-range {
        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
      }
    }
    order number;
  }
  application application-name;
  disable;
}
```

You can include the following application rule properties:

- **address**—Address properties for APPID rule processing. This statement is mandatory; you must specify either destination or source properties.
- **destination**—Destination address and port information. The **ip** statement defines the IP address and netmask (IPv4 only), and the **port-range** statement defines the TCP or UDP port number or numeric range, entered as **[*minimum-value – maximum-value*]**.
- **source**—Source address and port information. The **ip** statement defines the IP address and netmask (IPv4 only), and the **port-range** statement defines the TCP or UDP port number or numeric range, entered as **[*minimum-value – maximum-value*]**.
- **order**—Application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session; the lower the number, the higher the priority. This statement is mandatory and must contain a unique value.

Copyright © 2013, Juniper Networks, Inc.

- **application**—Name of the application to be included in the rule.
- **disable**—Disable processing for this application rule.

The **rule-set** statement defines a collection of APPID rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services application-identification]** hierarchy level with a **rule** statement for each rule:

```
rule-set rule-set-name {
  rule application-rule-name;
}
```

For a configuration example, see "Examples: Configuring Application Identification Properties" on page 796.

Using Stateful Firewall Rules to Identify Data Sessions

The APPID configuration properties enable the Junos OS to detect applications based on signatures, ports, and addresses. For signature-based detection, most of the protocol control sessions are identified, but data sessions are not identified. For example, APPID identifies FTP connections to port 21 (FTP control sessions); however, FTP can open child/data sessions to transfer files and data. These sessions are not identified by signature-based APPID because they do not have well-defined signatures.

Application-level gateways (ALGs) configured using stateful firewall rules can assist APPID in identifying these data sessions. These sessions include file and video transfers that are heavy consumers of bandwidth, so a mechanism for policing and classifying this traffic effectively is a useful tool. In addition to FTP, this mechanism applies to TFTP and RTSP traffic.

To incorporate the stateful firewall rules into Dynamic Application Awareness for Junos OS sessions, include the following configurations:

1. Include the stateful firewall package at the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]** hierarchy level:

   ```
   package jservices-sfw;
   ```

2. Define two stateful firewall rules as shown in the following example, one to identify the appropriate ALGs for FTP, TFTP, or RTSP traffic and the other to allow all traffic:

   

   NOTE: Session Initiation Protocol (SIP) is already covered by APPID and the SIP ALG is not supported by stateful firewall, hence a SIP configuration is not needed.

   ```
   [edit services]
   stateful-firewall {
     rule rule1 {
       match-direction input-output;
       term term1 {
   ```


Copyright © 2013, Juniper Networks, Inc.

```
        from {
          applications [ junos-ftp junos-tftp junos-rtsp ];
        }
        then {
          accept;
        }
      }
    }
    rule rule2 {
      match-direction input-output;
      term term1 {
        then {
          accept;
        }
      }
    }
    rule-set rs1 {
      rule rule1;
      rule rule2;
    }
}
```



NOTE: The existing AACL and L-PDF operational mode commands should report the new applications when they are identified.

3. Attach the stateful firewall rule set to a service set, as shown in the following example:

```
service-set test-chaining {
  application-identification-profile add-based;
  stateful-firewall-rule-sets rs1;
  idp-profile idp1;
  aacl-rules rule1;
  interface-service {
    service-interface ms-2/0/0.0;
  }
}
```

4. Include *no-drop* settings for stateful firewall and TCP, as needed.

   Stateful firewall processing drops packets in a number of scenarios:

   - TCP sessions do not start with a SYN flag. (This prevents sessions from resuming; otherwise, when the PIC starts for the first time, all existing TCP sessions in flight will be dropped).
   - If the TCP tracker detects SYN but no SYN/ACK or only an ACK, then the ACK is dropped. There are a number of similar checks to verify the TCP connection, window checks, and so forth.
   - TCP checks for stateful firewall are aggressive when ALGs are run. It is not possible to ignore TCP errors when an ALG is run on a session.
   - If an ALG detects malformed packets (for example, if the FTP PORT command is not RFC-compliant), it drops packets. If an ALG is not able to allocate resources, it drops packets.

Copyright © 2013, Juniper Networks, Inc.

You can include the settings shown in the following example to assist in controlling these packet drops:

```
[edit interfaces]
ms-1/2/0 {
   services-options {
      ignore-errors {
         tcp;
         alg;
      }
   }
}
```

The **tcp** statement mediates the first two issues listed, with reference to TCP SYN detection. The **alg** statement handles the fourth issue. ALGs require strict TCP processing, which cannot be relaxed.

Configuring Application Profiles

You can define an application profile for use in a service set. The profile consists of one or more rule sets, but only one profile can be included per service set.

To specify the application profile constituents, include the **profile** statement at the **[edit services application-identification]** hierarchy level:

```
profile profile-name {
   [ rule-set rule-set-name ];
}
```

You assign a profile name and include one or more predefined rule sets. For more information on rule sets, see "Configuring APPID Rules" on page 785. You can then include the profile in a service-set definition:

```
[edit services]
service-set service-set-name {
   profile profile-name;
}
```

The definitions specific to Dynamic Application Awareness include the APPID and IDP profiles and the AACL rule set. For more information on service sets, see Service Set Properties.

Configuring Application Groups

You can define an application group to process a number of applications or subgroups at the same time. To configure application group properties, include the **application-group** statement at the **[edit services application-identification]** hierarchy level:

```
application-group group-name {
   application-groups {
      application-group-name;
   }
   applications {
      application-name;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    index number;
    disable;
}
```

You can include the following application group properties:

- **applications**—List of applications to include in this application group. The **name** statement is mandatory and must include at least one entry.
- **application-groups**—List of application groups to include in a larger application group. The **name** statement is mandatory and must include at least one entry.
- **index**—Application group index number in the range from 1 through 65,534. This mandatory value must be unique.
- **disable**—Disable processing for this application group.

For a configuration example, see "Examples: Configuring Application Identification Properties" on page 796.

Application Identification for Nested Applications

The application identification feature is used by intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports. *Nested applications* are protocols running over the parent application. For example, both Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications. To do this, the application layer is split into two layers: Layer 7 applications and Layer 7 protocols.

The predefined application signatures included with Junos OS have been created to detect the Layer 7 nested applications. Predefined application signatures can be used in attack objects.

To configure nested application properties, include the **nested-application** statement at the **[edit services application-identification]** hierarchy level:

```
nested-application name {
    index number;
    protocol protocol;
    signature name {
        chain-order ;
        maximum-transactions number;
        member name {
            context (http-header-content-type | http-header-host | http-url-parsed |
                http-url-parsed-param-parsed);
            direction (any | client-to-server | server-to-client);
            pattern dfa-pattern;
        }
        order number;
    }
    type type;
}
```

You can include the following application rule properties:

Copyright © 2013, Juniper Networks, Inc.

- **chain-order**—Signatures can contain multiple members. If the chain order feature is on, those members are read in order. The default for this option is no chain order. If a signature contains only one member, this option is ignored.
- **context**—Define a service specific context. The options are **http-header-content-type** , **http-header-host** , **http-url-parsed**, **http-url-parsed-param-parsed**. This statement is mandatory.
- **direction**—The connection direction of the packets to apply pattern matching. The options are **client-to-server**, **server-to-client**, or **any**. This statement is mandatory.
- **index**—A number that is a one-to-one mapping to the application name that is used to ensure that each signature definition is unique. The index range for predefined applications is 1 through 32767. The index range for custom applications and custom nested applications is 32768 through 65534.
- **maximum transactions**—The maximum number of transactions that should occur before a match is made. This statement is mandatory.
- **member**—Define a member name for a custom nested application signature definition. Custom definitions can contain multiple members that define attributes for an application.
- **order**—Define application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session. The lower number has higher priority. This statement is mandatory.
- **pattern**—Define an attack pattern to be detected. This statement is mandatory.
- **protocol**—The protocol that will be monitored to identify nested applications. The value **http** is supported. This statement is mandatory.
- **signature**—Name of the custom nested application signature definition. Must be a unique name with a maximum length of 32 characters. This statement is mandatory.
- **type**—Well- known application name for this application definition, such as Facebook or Kazza. This application name must be unique with a maximum length of 32 characters. This statement is mandatory.

Disabling Application Identification for Nested Applications

Sometimes there is a need to identify multiple different applications running on the same Layer 7 protocols. For example, both Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications. To do this, the application layer is split into two layers: Layer 7 applications and Layer 7 protocols. Application identification for nested applications is turned on by default. You can manually turn it off by using the CLI.

To disable nested application identification:

- Set the **no-nested-application** statement.

```
[edit services application-identification nested-application-settings]
user@host# no-nested-application
```


Copyright © 2013, Juniper Networks, Inc.

To verify the configuration, issue the **show services application-identification nested-application-settings** command.

To reenable nested application identification:

- Delete the **no-nested-application** statement.

```
[edit services application-identification nested-application-settings]
user@host# delete services application-identification nested-application-settings
  no-nested-application
```

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

Related Documentation

- Application Identification for Nested Applications on page 789

Configuring Global APPID Properties

You can define additional properties that apply on a global basis to APPID processing and are not part of a specific application, group, rule, or profile definition. To configure these global APPID properties, include the following statements at the **[edit services application-identification]** hierarchy level:

```
application-identification {
  application-system-cache-timeout seconds;
  max-checked-bytes bytes;
  min-checked-bytes bytes;
  nested-applicationname
  nested-application-settings
  no-application-identification
  no-application-system-cache;
  no-clear-application-system-cache;
  no-protocol-method;
  no-signature-based;
  signature-method-all-ports;
}
```

The global application properties have the following effect:

- **application-system-cache-timeout**—Lifetime for system cache entries, in seconds.
- **max-checked-bytes**—The maximum number of bytes to be inspected in APPID processing, in the range from 0 through 100,000 bytes.
- **min-checked-bytes**—The minimum number of bytes to be inspected in APPID processing, in the range from 0 through 2000 bytes.
- **nested-application**—Configure a custom nested application definition for the desired application name that will be used by the system to identify the nested application as it passes through the device. For more information see nested-application.
- **nested-application-settings**—Configure nested application options for application identification services. For more information seenested-application-settings.
- **no-application-identification**—Disable all application identification methods.

Copyright © 2013, Juniper Networks, Inc.

- **no-application-system-cache**—Disable storing application identification results in the application system cache.
- **no-clear-application-system-cache**—Disable clearing the application system cache.
- **no-protocol-method**—Disable the protocol-based application identification method, which is enabled by default.
- **no-signature-based**—Disable the signature-based application identification method.
- **signature-method-all-ports**—Run signature matching on all traffic.

Configuring Automatic Download of Application Package Updates

You can set up automatic downloading of application package updates. To configure downloads, include the **download** statement at the **[edit services application-identification]** hierarchy level:

```
download {
  automatic {
    interval hour;
    start-time time;
  }
  url url;
}
```

You can include the following download statements:

- **download**—Define download properties.
- **automatic**—Set **start-time** value and **interval** in hours for automatic downloads. The default **start-time** is **0:00** and the range is from 0:00 through 24:00. The default **interval** is **24** and the range is from 1 through 168.
- **url**—Specify the download URL.

Configuring APPID Support for Heuristics

Heuristics methodology provides a mechanism for identifying encrypted data packets in point-to-point applications. These packets are not normally detected by the existing application signatures.

To enable APPID to employ heuristics in traffic identification:

1. Include the **enable-heuristics** statement:

```
[edit services application-identification]
user@host# enable-heuristics
```

The **show services application-identification counter** operational command includes additional output fields that report the number of encrypted sessions.

Copyright © 2013, Juniper Networks, Inc.

NOTE: When you enable heuristics, performance and scaling values might be negatively affected. This mechanism assists the APPID module in identifying encrypted traffic, but only if the identifications are supported by the current signature package.

Configuring APPID Support for Unidirectional Traffic

With asymmetrical routing, a networking device sees only one side of the network sessions, either from client to server or from server to client. Additional functionality is required to support application identification with unidirectional traffic. This addition enables a session for a specified service set to support an asymmetrical routing environment, and allows complete application matches using existing application signatures for traffic in the client-to-server direction only.

To enable APPID to support application matching on unidirectional traffic:

1. Include the **support-uni-directional-traffic** statement:

   ```
   [edit services service-set service-set-name service-set-options]
   user@host# support-uni-directional-traffic
   ```

   This enables the session belonging to the specified service set to support the asymmetrical routing environment. The APPID module then reports complete matches for the unidirectional traffic.

2. Include the **enable-asymmetric-traffic-processing** statement:

   ```
   [edit services service-set service-set-name service-set-options]
   user@host# enable-asymmetric-traffic-processing
   ```

   This enables the framework and plug-in to handle unidirectional traffic at a service-set level.

When you enable these settings, APPID treats unidirectional TCP traffic like a UDP connection. UDP traffic itself does not receive any special treatment because the service PIC cannot determine whether UDP traffic is unidirectional or bidirectional. The settings do not affect processing of sessions created with bidirectional traffic.

If the traffic includes both unidirectional and bidirectional sessions, the APPID module uses heuristics to decide whether to change the reporting logic.

NOTE: This feature does not change the processing for any services except APPID. However, other services, including stateful firewall, AACL, and IDP, can process unidirectional traffic in a limited manner.

Copyright © 2013, Juniper Networks, Inc.

Tracing APPID Operations

Tracing operations track all adaptive services operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit services application-identification]** hierarchy level, the default tracing behavior is as follows:

- Important events are logged in a file called **serviced** located in the **/var/log** directory.
- When the file **serviced** reaches 128 kilobytes (KB), it is renamed **serviced.0**, then **serviced.1**, and so on, until there are three trace files. Then the oldest trace file (**serviced.2**) is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Only the user who configures the tracing operation can access the log files.
- To display the end of the log, issue the **show log serviced | last** operational mode command:

```
[edit]
user@host# run show log serviced | last
```

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements:

```
file filename <files number> <match regex> <size size> <(world-readable |
  no-world-readable>;
flag {
  all;
}
```

You configure these statements at the **[edit services application-identification traceoptions]** hierarchy level.

These statements are described in the following sections:

- Configuring the APPID Log Filename on page 794
- Configuring the Number and Size of APPID Log Files on page 795
- Configuring Access to the Log File on page 795
- Configuring a Regular Expression for Lines to Be Logged on page 795
- Configuring the Tracing Flags on page 795

Configuring the APPID Log Filename

By default, the name of the file that records trace output is **serviced**. You can specify a different name by including the **file** statement at the **[edit services application-identification traceoptions]** hierarchy level:

```
file filename;
```


Copyright © 2013, Juniper Networks, Inc.

Configuring the Number and Size of APPID Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed ***filename*.0**, then ***filename*.1**, and so on, until there are three trace files. Then the oldest trace file (***filename*.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit services application-identification traceoptions]** hierarchy level:

```
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (***filename***) reaches 2 MB, ***filename*** is renamed ***filename*.0**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0** is renamed ***filename*.1** and ***filename*** is renamed ***filename*.0**. This process repeats until there are 20 trace files. Then the oldest file (***filename*.19**) is overwritten by the newest file (***filename*.0**).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, only the user who configures the tracing operation can access log files.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit services application-identification traceoptions]** hierarchy level:

```
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit services application-identification traceoptions]** hierarchy level:

```
file no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the **match** statement at the **[edit services application-identification traceoptions file *filename*]** hierarchy level and specifying a regular expression (regex) to be matched:

```
file filename match regex;
```

Configuring the Tracing Flags

By default, if the **traceoptions** configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit services application-identification traceoptions]** hierarchy level:

```
flag {
   all;
}
```

Copyright © 2013, Juniper Networks, Inc.

Currently, the only supported flag is **all**, which instructs the router to trace all operations.

Examples: Configuring Application Identification Properties

The following examples show an address-based application identification configuration:

```
[edit services application-identification]
rule rule1 {
   application-name test2;
   address 1 {
      source {
         ip 10.110.1.1/16;
         port-range {
            tcp 1110-1150;
         }
      }
      destination {
         ip 10.11.1.1/16;
         port-range {
            tcp 111-1100;
         }
      }
      order 1;
   }
}
}
```

```
[edit services application-identification]
rule-set rs1 {
   rule rule1;
}
profile pf1 {
   rule-set rs1;
}
[edit services]
service-set sset1 {
   application-identification-profile pf1;
}
```

The following examples show application group configuration:

```
[edit services application-identification]
application-group junos:peer-to-peer {
   index 5;
   application-groups {
      junos:chat;
      junos:file-sharing;
      junos:voip;
   }
}
```

```
[edit services application-identification]
application-group junos:voip {
   index 14;
   applications {
      junos:h225ras;
      junos:h225sgn;
```


Copyright © 2013, Juniper Networks, Inc.

```
        junos:mgcp;
        junos:sip;
    }
}
```

The following examples show application identification for nested application configuration:

```
nested-application nested1 {
    type nested1;
    index 65345;
protocol HTTP;
    signature nestedcust001 {
        member m01 {
            context http-url-parsed;
            pattern .*nested.*;
            direction any;
        }
        maximum-transactions 2;
    order 3825;
```

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 38

Summary of Application Identification Configuration Statements

The following sections explain each of the application identification configuration statements. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

address

Syntax

```
address address-name {
  destination {
    ip address</prefix-length>;
    port-range {
      tcp [ ports-and-port-ranges ];
      udp [ ports-and-port-ranges ];
    }
  }
  source {
    ip address</prefix-length>;
    port-range {
      tcp [ ports-and-port-ranges ];
      udp [ ports-and-port-ranges ];
    }
  }
  order number;
}
```

Hierarchy Level [edit services application-identification rule *rule-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define address properties for application-identification rule processing. This statement is mandatory; you must specify either the destination or source properties.

Options ***address-name***—Identifier for address information.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring APPID Rules on page 785

Copyright © 2013, Juniper Networks, Inc.

application

See the following sections:

- application (Defining) on page 801
- application (Including in Rule) on page 802

application (Defining)

Syntax

```
application application-name {
  disable;
  idle-timeout seconds;
  index number;
  port-mapping {
    disable;
    port-range {
      tcp [ ports-and-port-ranges ];
      udp [ ports-and-port-ranges ];
    }
  }
  session-timeout seconds;
  type type;
  type-of-service service-type;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the application and its properties.

The remaining statements are explained separately.

Options ***application-name***—Identifier for the application. This is a mandatory value and has a maximum length of 32 characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

application (Including in Rule)

Syntax application *application-name*;

Hierarchy Level [edit services application-identification rule *rule-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Identify the application for inclusion in a rule.

Options ***application-name***—Identifier for the application.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Rules on page 785

application-group

Syntax

```
application-group group-name {
  disable;
  application-groups {
    application-group-name;
  }
  applications {
    application-name;
  }
  index number;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the properties and contents of the application group.

Options ***group-name***—Unique identifier for the group.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Application Groups on page 788

Copyright © 2013, Juniper Networks, Inc.

application-groups (Services Application Identification)

Syntax
```
application-groups {
  application-group-name;
}
```

Hierarchy Level [edit services application-identification application-group *group-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Identify the list of application groups for inclusion in a larger application group. An ***application-group-name*** statement is mandatory.

Options ***application-group-name***—Identifier for the application group. Maximum length is 32 characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Application Groups on page 788

application-system-cache-timeout

Syntax application-system-cache-timeout *seconds*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the lifetime for entries in the application system cache.

Options ***seconds***— Lifetime for system cache entries, in seconds.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

applications (Services Application Identification)

Syntax

```
applications {
  application-name;
}
```

Hierarchy Level [edit services application-identification application-group *group-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Identify the list of applications for inclusion in the application group.

Options ***application-name***—Identifier for the application. Maximum length is 32 characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Groups on page 788

automatic

Syntax

```
automatic {
  interval hour;
  start-time time;
}
```

Hierarchy Level [edit services application-identification download]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define automatic download properties.

Options **interval** ***hour***—Download interval in hours. The default is **24** and the range is from 1 through 168.

start-time ***time***—Start-time value. The default is **0:00** and the range is from 0:00 through 24:00.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Automatic Download of Application Package Updates on page 792

Copyright © 2013, Juniper Networks, Inc.

chain-order

Syntax chain-order;

Hierarchy Level [edit services application-identification nested-application *name* signature *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Signatures can contain multiple members. If the chain order feature is on, those members are read in order. By default, chain ordering is turned off. If a signature contains only one member, this option is ignored.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789

context

Syntax context (http-header-content-type | http-header-host | http-url-parsed | http-url-parsed-param-parsed);

Hierarchy Level [edit services application-identification nested-application *name* signature *name* member *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define a service-specific context, such as **http-url**.

Options *value*—Service-specific context.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789

Copyright © 2013, Juniper Networks, Inc.

destination (Services)

Syntax

```
destination {
  ip address</prefix-length>;
  port-range {
    tcp [ ports-and-port-ranges ];
    udp [ ports-and-port-ranges ];
  }
}
```

Hierarchy Level [edit services application-identification rule *rule-name* address *address-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define destination properties for application-identification rule processing.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring APPID Rules on page 785

direction

Syntax direction (any | client-to-server | server-to-client) ;

Hierarchy Level [edit services application-identification nested-application *name* signature *name* member *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the connection direction of the packets to apply pattern matching.

Options ***direction***—The directions of packets are **client-to-server**, **server-to-client**, or **any**.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation

- Application Identification for Nested Applications on page 789.

Copyright © 2013, Juniper Networks, Inc.

disable

See the following sections:

- disable (APPID Application) on page 807
- disable (APPID Application Group) on page 807
- disable (APPID Port Mapping) on page 808

disable (APPID Application)

Syntax disable;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable this application definition.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

disable (APPID Application Group)

Syntax disable;

Hierarchy Level [edit services application-identification application-group *group-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable application group properties.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Application Groups on page 788

Copyright © 2013, Juniper Networks, Inc.

disable (APPID Port Mapping)

Syntax disable;

Hierarchy Level [edit services application-identification application *application-name* port-mapping]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable port-mapping properties for application identification.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

disable-global-timeout-override

Syntax disable-global-timeout-override;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Disallow overriding a global inactivity or session timeout.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

download

Syntax

```
download {
  automatic {
    interval hour;
    start-time time;
  }
  url url;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define application download properties.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Automatic Download of Application Package Updates on page 792

enable-heuristics

Syntax enable-heuristics;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 11.2.

Description Enables APPID to identify encrypted data packets in point-to-point applications by using heuristics methodology.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Support for Heuristics on page 792

Copyright © 2013, Juniper Networks, Inc.

enable-asymmetic-traffic-processing

Syntax enable-asymmetic-traffic-processing;

Hierarchy Level [edit services service-set *service-set-name* service-set-options]

Release Information Statement introduced in Junos OS Release 11.2.

Description Enables APPID to perform application matching on unidirectional traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Support for Unidirectional Traffic on page 793

enable-heuristics

Syntax enable-heuristics;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 11.2.

Description Enables APPID to identify encrypted data packets in point-to-point applications by using heuristics methodology.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Support for Heuristics on page 792

Copyright © 2013, Juniper Networks, Inc.

idle-timeout

Syntax idle-timeout *seconds*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define idle timeout for an application in seconds. When the timeout period expires, the session ends if no packets have been received.

Options ***seconds***—Idle timeout period.
Default: 30
Range: 1 through 604,800

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

ignore-errors

Syntax ignore-errors <alg> <tcp>;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description Define settings for minimizing TCP packet drops during stateful firewall processing.



NOTE: **ignore-errors option is not supported on adaptive services interfaces (sp-x/y/z).**

Options **alg**—Mediate ALG behavior that results in dropping malformed packets or random packets when the software is unable to allocate resources.

tcp—Prevent software from dropping packets that fail TCP SYN checks.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

inactivity-non-tcp-timeout

Syntax inactivity-non-tcp-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Define the inactivity timeout period for non-TCP established sessions in seconds.

Options *seconds*—Timeout period.
Range: 4 through 86,400

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

inactivity-tcp-timeout

Syntax inactivity-tcp-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Define the inactivity timeout period for TCP established sessions in seconds.

Options *seconds*—Timeout period.
Range: 4 through 86,400

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

index (Nested Applications)

Syntax index *number*;

Hierarchy Level [edit services application-identification nested-application *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Set a number that is a one-to-one mapping to the application name. The application name is used to ensure that each signature definition is unique.

Options ***number***—Numeric value associated with an application name. The index range for predefined applications is from 1 through 32767. The index range for custom applications and custom nested applications is from 32768 through 65534.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789.

index

Syntax index *number*;

Hierarchy Level [edit services application-identification application *application-name*],
[edit services application-identification application-group *group-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Assign an application or application-group index number. This is a mandatory value.

Options ***number***—Index number; must be a unique, unsigned value.

Range: 0 through 65535

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783
- Configuring Application Groups on page 788

Copyright © 2013, Juniper Networks, Inc.

ip

Syntax ip *address</prefix-length>*;

Hierarchy Level [edit services application-identification rule *rule-name* address destination],
[edit services application-identification rule *rule-name* address source]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define an IP address and netmask for identifying the traffic destination or source.

Options ***address</prefix-length>***—IP address and netmask.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Rules on page 785

max-checked-bytes

Syntax max-checked-bytes *bytes*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the maximum number of bytes to be inspected.

Options ***bytes***—Maximum number of bytes.
Range: 0 through 100,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

maximum-transactions

Syntax maximum-transactions *number*;

Hierarchy Level [edit services application-identification nested-application *name* signature *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Set the maximum number of transactions required before a match is made.

Options ***number***—Maximum number of transactions.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789

member

Syntax member *name*;

Hierarchy Level [edit services application-identification nested-application *name* signature *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define a member name for a custom nested application signature definition. Custom definitions can contain multiple members that define attributes for an application.

Options ***name***—Name of member for a custom nested application signature definition.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789

Copyright © 2013, Juniper Networks, Inc.

min-checked-bytes

Syntax min-checked-bytes *bytes*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the minimum number of bytes to be inspected.

Options ***bytes***—Minimum number of bytes.
Range: 0 through 2000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

nested-application

Syntax

```
nested-application name {
  index number;
  protocol protocol ;
  signature name {
    chain-order ;
    maximum-transactions number;
    member name {
      context (http-header-content-type | http-header-host | http-url-parsed |
        http-url-parsed-param-parsed);
      direction (any | client-to-server | server-to-client);
      pattern dfa-pattern;
    }
    order number;
  }
  type type;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure a custom nested application definition for the desired application name that will be used by the system to identify the nested application as it passes through the device. Custom nested application definitions can be used for nested applications that are not part of the Juniper Networks predefined nested application database.

Options *name*—Name of nested application.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation

- Application Identification for Nested Applications on page 789

nested-application-settings

Syntax

```
nested-application-settings {
  no-application-system-cache;
  no-nested-application;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure nested application options for application identification services.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789.

no-application-identification

Syntax

```
no-application-identification;
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable all application identification methods.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

no-application-system-cache

Syntax no-application-system-cache;

Hierarchy Level [edit services application-identification],
[edit services application-identification nested-application-settings]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable storing application identification results in the application system cache. Nested application identification information is saved in the application system cache to improve performance. This cache is updated when a different application is identified. This caching is turned on by default. Use the no-application-system-cache statement to turn it off.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791
- Application Identification for Nested Applications on page 789.

no-clear-application-system-cache

Syntax no-clear-application-system-cache;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable clearing the application system cache.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

no-nested-application

Syntax no-nested-application;

Hierarchy Level [edit services application-identification nested-application-settings]

Release Information Statement introduced in Junos OS Release 10.2.

Description Sometimes there is a need to identify multiple different applications running on the same Layer 7 protocols. For example, both Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications. To do this, the application layer is split into two layers: Layer 7 applications and Layer 7 protocols. This function is turned on by default. Use the **no-nested-application** statement to turn it off.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789

no-protocol-method

Syntax no-protocol-method;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.1.

Description Disable the protocol-based application identification method.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

no-signature-based

Syntax no-signature-based;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable the signature-based application identification method.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Global APPID Properties on page 791

order (Services Application Identification)

Syntax order *number*;

Hierarchy Level [edit services application-identification nested-application *name* signature *name* member *name*]
[edit services application-identification rule *rule-name* address]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session. The lower number has higher priority.

Options *number*—Order number. This value is mandatory and must be unique.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Rules on page 785
- Application Identification for Nested Applications on page 789

Copyright © 2013, Juniper Networks, Inc.

pattern

Syntax pattern *dfa-pattern*;

Hierarchy Level [edit services application-identification nested-application *name* signature *name* member *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define an attack pattern to be detected.

Options ***dfa-pattern***—Pattern of attack to match. Deterministic Finite Automata (DFA) is a powerful pattern matching engine.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation
- Application Identification for Nested Applications on page 789

port-mapping

Syntax
```
port-mapping {
  disable;
  port-range {
    tcp [ ports-and-port-ranges ];
    udp [ ports-and-port-ranges ];
  }
}
```

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define port-mapping properties for application identification.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

port-range

Syntax

```
port-range {
  tcp [ ports-and-port-ranges ];
  udp [ ports-and-port-ranges ];
}
```

Hierarchy Level [edit services application-identification application *application-name* port-mapping],
[edit services application-identification rule *rule-name* address destination],
[edit services application-identification rule *rule-name* address source]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define TCP and UDP port numbers or numeric ranges. For port-mapping configurations, this entry is required if the parent node exists.

Options ***ports-and-port-ranges***—Individual port numbers, numeric port ranges, or both. Separate the values with spaces. The format for numeric port ranges is *minimum-value–maximum-value*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Defining an Application Identification on page 783
- Configuring APPID Rules on page 785

profile

Syntax

```
profile profile-name {
  rule-set rule-set-name;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define members of application profile, which consists of one or more rule sets.

Options ***profile-name***—Identifier for application profile.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Application Profiles on page 788

Copyright © 2013, Juniper Networks, Inc.

protocol

Syntax protocol *protocol*;

Hierarchy Level [edit services application-identification nested-application *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify the protocol that will be monitored to identify nested applications. HTTP is supported.

Options ***protocol***—An agreed-upon or standardized method for transmitting data and establishing communications between different devices. The value **http** is supported.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation

- Application Identification for Nested Applications on page 789

Copyright © 2013, Juniper Networks, Inc.

hcm

See the following sections:

- from on page 825
- host on page 826
- request-url on page 826
- term on page 827
- then on page 827
- url on page 828
- url-list on page 828
- url-rule on page 829
- url-rule-set on page 829

from

Syntax

```
from {
  url-list url-list-name;
  url url_identifier {
    host hostname;
    request-url page-name;
  }
}
```

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify input conditions for the HCM term.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

host

Syntax host *hostname*;

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num* from url *url_identifier*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify a hostname for the matching URL for the **term**. A match for that term is considered when a URL matches the **hostname** and the **request-URL** within the same term.

Options ***hostname***—Name of the host for the URL rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

request-url

Syntax request-url *page-name* ;

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num* from url *url_identifier*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify a request-URL to match the **term**. A match for the term is considered when a URL matches any hostname and any request-URL within a term.

Options *page-name*—Page name of the request URL.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

term

Syntax

```
term term-num {
  from {
    url-list url-list-name;
    url url_identifier {
      host hostname;
      request-url page-name;
    }
  }
}
```

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify a numbered identity for each term inside a rule.

Options *term-num*—Identifier value for the term.

Range: 1 through 255

Default: If no value is entered, the default value is 1.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

then

Syntax

```
then {
  discard;
  accept;
  count;
  log-request;
}
```

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Define the HCM term actions.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

url

Syntax url *url_identifier*;

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num* from]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify an integer that uniquely identifies a particular URL definition within a term.

Options ***url_identifier***—URL identifier number.
Range: 1 through 32,767
Default: If no value is added, the default value is 1.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

url-list

Syntax url-list *url-list-name* ;

Hierarchy Level [edit services hcm url-rule *url-rule-name* term *term-num* from]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify the name of a previously defined URL list to be included as a matching condition. A match for the term is considered when a URL matches any hostname and any request-URL within the same term.

Options ***url-list-name***—Name of the previously defined URL list.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

url-rule

Syntax

```
url-rule url-rule-name {
  term term-num {
    from {
      url-list url-list-name;
      url url_identifier {
        host hostname;
        request-url page-name;
      }
    }
    then {
      discard;
      accept;
      count;
      log-request;
    }
  }
}
```

Hierarchy Level [edit services hcm]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify the name of the URL rule.

Options ***url-rule-name***—Name of the URL rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

url-rule-set

Syntax

```
url-rule-set url-rule-set-name {
  url-rule rule1;
  url-rule rule2;
}
```

Hierarchy Level [edit services hcm url-rule *url-rule-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the name of the rule set. A rule set is a collection of rules ordered in the sequence in which they are entered.

Options ***url-rule-set-name***—Name of the collection of URL rules that constitute this rule set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rule

See the following sections:

- rule (Configuring) on page 830
- rule (Including in Rule Set) on page 831

rule (Configuring)

Syntax

```
rule rule-name {
  address {
    destination {
      ip address</prefix-length>;
      port-range {
        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
      }
    }
    source {
      ip address</prefix-length>;
      port-range {
        tcp [ ports-and-port-ranges ];
        udp [ ports-and-port-ranges ];
      }
    }
    order number;
  }
  application application-name;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define properties for application-identification rule processing.

Options *rule-name*—Unique identifier for the rule.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring APPID Rules on page 785

Copyright © 2013, Juniper Networks, Inc.

rule (Including in Rule Set)

Syntax rule *rule-name*;

Hierarchy Level [edit services application-identification rule-set *rule-set-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Identify rules for inclusion in application rule set.

Options ***rule-name***—Unique identifier for the rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Rules on page 785

rule-set (Services Application Identification)

Syntax
```
rule-set rule-set-name {
    rule application-rule-name;
}
```

Hierarchy Level [edit services application-identification],
[edit services application-identification profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define members of rule set.

Options ***rule-set-name***—Unique identifier for the rule set.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Rules on page 785

Copyright © 2013, Juniper Networks, Inc.

services (Application Identification)

Syntax services application-identification { ... }

Hierarchy Level [edit]

Release Information **services** statement introduced before Junos OS Release 7.4.
application-identification statement introduced in Junos OS Release 9.5.

Description Define the services to be applied to traffic.

Options **application-identification**—The values configured for application-identification properties.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Application Identification

Copyright © 2013, Juniper Networks, Inc.

session-timeout

See the following sections:

- session-timeout (Interfaces) on page 833
- session-timeout (Application Identification) on page 833

session-timeout (Interfaces)

Syntax session-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in Junos OS Release 10.0.

Description Define session lifetime globally for the Multiservices interface in seconds.

Options ***seconds***—Duration of session.
Range: 4 through 86,400

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

session-timeout (Application Identification)

Syntax session-timeout *seconds*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define session lifetime for the specified application in seconds.

Options ***seconds***—Duration of session.
Default: 3600
Range: 1 through 604,800

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

signature

Syntax

```
signature name {
  chain-order;
  maximum-transactions number;
  member name {
    context value;
    direction (any | client-to-server | server-to-client);
    pattern dfa-pattern;
  }
  order number;
}
```

Hierarchy Level [edit services application-identification nested-application *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Identify the name of the custom nested application signature definition. The name must be unique with a maximum length of 32 characters.

Options *name*—Name of the signature definition.

The remaining statements are described separately.

Required Privilege Level system—To view this statement in the configuration.
system control—To add this statement to the configuration.

Related Documentation

- Application Identification for Nested Applications on page 789

signature-method-all-ports

Syntax signature-method-all-ports

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 11.2.

Description Runs signature matching on all traffic in application-identification. This is called the signature-match mode.

In the default mode, or fast-port-match mode, all traffic destined to well-known ports (up to 1024) immediately returns the final port match. However, the device runs signature matching for all traffic destined for port 80,

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Global APPID Properties on page 791

Copyright © 2013, Juniper Networks, Inc.

source

Syntax

```
source {
  ip address</prefix-length>;
  port-range {
    tcp [ ports-and-port-ranges ];
    udp [ ports-and-port-ranges ];
  }
}
```

Hierarchy Level [edit services application-identification rule *rule-name* address *address-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define source properties for application-identification rule processing.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Rules on page 785

support-uni-directional-traffic

Syntax support-uni-directional-traffic;

Hierarchy Level [edit services service-set *service-set-name* service-set-options]

Release Information Statement introduced in Junos OS Release 11.2.

Description Enables APPID to perform application matching on unidirectional traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring APPID Support for Unidirectional Traffic on page 793

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Application Identification)

Syntax

```
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure application identification tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000 files

Default: 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag—Tracing operation to perform. **all** is the only valid completion.

- **all**—Trace all events.

match *regex*—(Optional) Regular expression for lines to be logged.

no-world-readable—(Optional) Disallow any user to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10240 through 1073741824 or the maximum file size supported on your system

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—(Optional) Allow any user to read the log file.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing APPID Operations on page 794

type

Syntax type *type*;

Hierarchy Level [edit services application-identification application *application-name*]
[edit services application-identification nested-application *name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define type of application, such as HTTP or FTP.

Options *type*—Application type. This is a mandatory value and has a maximum length of 32 characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783
- Application Identification for Nested Applications on page 789

type-of-service

Syntax type-of-service *service-type*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the type of service by service objective. There is no default value.

Options The following ***service-type*** options are available:

- **maximize-reliability**—Service designed for maximum reliability in packet transmission.
- **maximize-throughput**—Service designed for maximum throughput.
- **minimize-delay**—Service designed for minimum delay in packet transmission.
- **minimize-monetary-cost**—Service designed for minimum monetary cost.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining an Application Identification on page 783

Copyright © 2013, Juniper Networks, Inc.

url

Syntax url *url*;

Hierarchy Level [edit services application-identification download]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the URL for application package downloads.

Options ***url***—Download URL.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Automatic Download of Application Package Updates on page 792

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 39

Application-Aware Access List Configuration Guidelines

To configure application-aware access list (AACL) services, include the **aacl** statements at the **[edit services]** hierarchy level:

```
[edit services]
aacl {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        application-group-any;
        application-groups [ application-group-names ];
        application-unknown;
        applications [ application-names ];
        destination-address address <any-unicast>;
        destination-address-range low minimum-value high maximum-value;
        destination-prefix-list list-name;
        source-address address <any-unicast>;
        source-address-range low minimum-value high maximum-value;
        source-prefix-list list-name;
      }
      then {
        (accept | discard);
        count (application | application-group | application-group-any | none);
        forwarding-class class-name;
        policer policer-name;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
}
```

This chapter contains the following sections:

- Configuring AACL Rules on page 840
- Configuring AACL Rule Sets on page 844

Copyright © 2013, Juniper Networks, Inc.

- Configuring Logging of AACL Flows on page 845
- Example: Configuring AACL Rules on page 845

Configuring AACL Rules

To configure an AACL rule, include the **rule *rule-name*** statement at the **[edit services aacl]** hierarchy level:

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-group-any;
      application-groups [ application-group-names ];
      applications [ application-names ];
      destination-address address <any-unicast>;
      destination-address-range low minimum-value high maximum-value;
      destination-prefix-list list-name;
      nested-applications [ nested-application-names ];
      nested-application-unknown
      source-address address <any-unicast>;
      source-address-range low minimum-value high maximum-value;
      source-prefix-list list-name;
    }
    then {
      (accept | discard);
      count (application | application-group | application-group-any | nested-application
        | none);
      forwarding-class class-name;
      policer policer-name;
    }
  }
}
```

Each AACL rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how to configure the components of AACL rules:

- Configuring Match Direction for AACL Rules on page 841
- Configuring Match Conditions in AACL Rules on page 841
- Configuring Actions in AACL Rules on page 842
- Logging AACL Flows Based on Application on page 843

Copyright © 2013, Juniper Networks, Inc.

Configuring Match Direction for AACL Rules

Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services aacl rule *rule-name*]** hierarchy level:

```
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, bidirectional rule creation is allowed.

The match direction is used with respect to the traffic flow through the services PIC or DPC. When a packet is sent to the PIC or DPC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the services PIC or DPC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC or DPC, the packet direction is output. For more information on inside and outside interfaces, see "Configuring Service Sets to be Applied to Services Interfaces" on page 468.

On the PIC or DPC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

Configuring Match Conditions in AACL Rules

To configure AACL match conditions, include the **from** statement at the **[edit services aacl rule *rule-name* term *term-name*]** hierarchy level:

```
from {
    application-group-any;
    application-groups [ application-group-names ];
    applications [ application-names ];
    destination-address address <any-unicast>;
    destination-address-range low minimum-value high maximum-value;
    destination-prefix-list list-name;
    nested-applications [ nested-application-names ];
    nested-application-unknown
    source-address address <any-unicast>;
    source-address-range low minimum-value high maximum-value;
    source-prefix-list list-name;
}
```

IPv4 and IPv6 source and destination addresses are supported. You can use either the source address or the destination address as a match condition, in the same way that you configure a firewall filter; for more information, see the Routing Policy Configuration Guide.

Alternatively, you can specify a list of source or destination prefixes by configuring the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either

Copyright © 2013, Juniper Networks, Inc.

the **destination-prefix-list** or the **source-prefix-list** statement in the AACL rule. For an example, see "Example: Configuring AACL Rules" on page 845.

If you omit the **from** term, the AACL rule accepts all traffic and the default protocol handlers take effect:

- User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP) create a bidirectional flow with a predicted reverse flow.
- IP creates a unidirectional flow.

You can also include application and application group definitions you have configured at the **[edit services application-identification]** hierarchy level; for more information, see the topics in Application Identification.

- To apply one or more specific application protocol definitions, include the **applications** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level.
- To apply one or more sets of application group definitions you have defined, include the **application-groups** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level.



NOTE: If you include one of the statements that specifies application protocols, the router derives port and protocol information from the corresponding configuration at the **[edit services application-identification]** hierarchy level; you cannot specify these properties as match conditions.

- To consider any application group defined in the database as a match, include the **application-group-any** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level.
- To consider any nested application defined in the database a match, include the **nested-applications** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level. Nested applications are protocols that run on a parent application. For example, if the Facebook application runs on the parent application junos:http, the nested application will be junos:http:facebook.

Configuring Actions in AACL Rules

To configure AACL actions, include the **then** statement at the **[edit services aacl rule *rule-name* term *term-name*]** hierarchy level:

```
then {
    (accept | discard);
    (count (application | application-group | application-group-any | nested-application |
        none) | forwarding-class class-name);
}
```

You must include one of the following actions:

Copyright © 2013, Juniper Networks, Inc.

- **accept**—The packet is accepted and sent on to its destination.
- **discard**—The packet is not accepted and is not processed further.

When you select **accept** as the action, you can optionally configure one or both of the following action modifiers. No action modifiers are allowed with the **discard** action.

- **count (application | application-group | application-group-any | nested-application | none)**—For all accepted packets that match the rules, record a packet count using AACL statistics practices. You can specify one of the following options; there is no default setting:
  - **application**—Count the application that matched in the **from** clause.
  - **application-group**—Count the application group that matched in the **from** clause.
  - **application-group-any**—Count all application groups that match **from application-group-any** under the **any** group name.
  - **nested-application**—Count all nested applications that matched in the **from** clause.
  - **none**—Same as not specifying **count** as an action.



NOTE:

- When a session closes before APPID has identified nested applications, the session is treated as a best-effort session and AACL does not get the nested application information. In such cases, nested applications will be reported as unknown applications.
- During the time that the application identification (APPID) feature has not yet made a final determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection. For more information, see "Best-Effort Application Identification of DPI-Serviced Flows" on page 777.

- **forwarding-class *class-name***—Specify the packets' forwarding-class name.

You can optionally include a **policer** that has been specified at the **[edit firewall]** hierarchy level. Only the bit-rate and burst-size properties specified for the policer are applied in the AACL rule set. The only action application when a policer is configured is **discard**. For more information on policer definitions, see the Routing Policy Configuration Guide.

Logging AACL Flows Based on Application

You can now log AACL flows based on application. You can select a specific application or request information on unknown applications.

You can now configure AACL rules to match unknown applications. All existing actions that can apply to recognized applications can also apply to unknown applications. You can use the following statements at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

- application-group-any
- application-groups
- application-unknown
- applications
- nested-application-unknown
- nested-applications

The addition of matching "application unknown" enables the specific logging of the input flows associated with applications that cannot be identified. Because logging is triggered by an input event, you must specify **match-direction** as **input-output** or **input**.

To configure logging of flows for AACL, include the **match-direction input** or **match-direction input-output** statement at the **[edit services aacl rule *rule-name*]** hierarchy level, include an **applications** or **application-unknown** statement at the **[edit services aacl rule *rule-name* term *term-name* from]** hierarchy level, and include only one **log** statement at the **[edit services aacl rule *rule-name* term *term-name* then]** hierarchy level. The log statements can include any of the following options:

- **session-start**
- **session-end**
- **session-start-end-no-stats**
- **session-start-interim-end**
- **session-interim-end**
- **session-end**

Configuring AACL Rule Sets

The **rule-set** statement defines a collection of AACL rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services aacl]** hierarchy level with a **rule** statement for each rule:

```
rule-set rule-set-name {
    rule rule-name;
}
```

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

Copyright © 2013, Juniper Networks, Inc.

Configuring Logging of AACL Flows

You can configure logging of AACL flows for a given application or for all unknown applications using AACL rules. You must set **match-direction** to **input** or **input-output** for logging to occur.

1. Create a rule and term.

   user@host# **edit services aacl rule** *rule-name* **term** *term-name*

2. Specify selection of an application.

   [edit services aacl rule *rule-name* term *term-name*]
   user@host# **set from applications** *application-name*]

   OR

   Specify selection of all unknown applications.

   [edit services aacl rule <variable>rule-name</variable > term
   <variable>term-name</variable>]
   set from application-unknown

3. In the **then** statement, specify logging of input flow.

   [edit services aacl rule *rule-name* term *term-name*]
   user@host# **set then log input-flows**]

Example—Configuration of Logging of Input Flows for Unknown Applications

```
[edit services aacl rule aacl_rule5]
match-direction input-output;
  term t0 {
    from {
      application-unknown;
    }
    then {
      count application;
      log input-flow;
      accept;
    }
  }
```

Example—Setup of a Specific Log File

The following example shows how to direct the aacl flow log to a file other than the default syslog file on the Routing Engine file system.

```
[edit system syslog]
file aacl_log {
  external any;
  match aacl-flow-log;
}
```

Example: Configuring AACL Rules

The following example shows an AACL configuration containing a rule with three terms using a variety of match conditions and actions:

Copyright © 2013, Juniper Networks, Inc.

```
[edit services aacl]
rule aacl-test {
   match-direction input;
   term term1 {
      from {
         source-address 10.0.1.1
         application test1;
      }
      then {
         accept;
      }
   }
   term term2 {
      from {
         source-address {
            any-unicast;
         }
         application test1;
      }
      then {
         discard;
      }
   }
   term term3 {
      from {
         source-address {
            any-unicast;
         }
         application test1 test2;
      }
      then {
         accept;
         count application;
      }
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 40

Summary of AACL Configuration Statements

The following sections explain each of the application-aware access list (AACL) services statements. The statements are organized alphabetically.

applications (Services AACL)

Syntax applications [*application-names*];

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.

Description Identify one or more applications defined in the application identification configuration for inclusion as a match condition.

Options ***application-names***—Identifiers of the applications.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Match Conditions in AACL Rules on page 841

Copyright © 2013, Juniper Networks, Inc.

application-groups (Services AACL)

Syntax application-groups [*application-group-names*];

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.

Description Identify one or more application groups defined in the application identification configuration for inclusion as a match condition.

Options ***application-group-names***—Identifiers of the application groups.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

application-group-any

Syntax application-group-any;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.

Description Indicates that any application group defined in the database is considered a match.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

application-unknown

Syntax application-unknown

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 11.2.

Description Enable AACL logging of flows for unknown applications.

Copyright © 2013, Juniper Networks, Inc.

destination-address

Syntax destination-address *address*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the destination address for rule matching.

Options ***address***—Destination IPv4 or IPv6 address or prefix value.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

destination-address-range

Syntax destination-address-range low *minimum-value* high *maximum-value*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the destination address range for rule matching.

Options ***minimum-value***—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

Copyright © 2013, Juniper Networks, Inc.

destination-prefix-list

Syntax destination-prefix-list *list-name*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the destination prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options *list-name*—Destination prefix list.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

from

Syntax

```
from {
  application-group-any;
  application-groups [ application-group-names ];
  application-unknown;
  applications [ application-names ];
  destination-address address <any-unicast>;
  destination-address-range low minimum-value high maximum-value;
  destination-prefix-list list-name;
  nested-application-unknown;
  source-address address <any-unicast>;
  source-address-range low minimum-value high maximum-value;
  source-prefix-list list-name;
}
```

Hierarchy Level [edit services aacl rule *rule-name* term *term-name*]

Release Information Statement introduced before Junos OS Release 9.5.

Description Specify match conditions for the AACL term.

Options For information on match conditions, see the description of firewall filter match conditions in the Routing Policy Configuration Guide.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring AACL Rules on page 840

Copyright © 2013, Juniper Networks, Inc.

match-direction

Syntax match-direction (input | output | input-output);

Hierarchy Level [edit services aacl rule *rule-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the direction in which the rule match is applied.

Options

input—Apply the rule match on the input side of the interface.

output—Apply the rule match on the output side of the interface.

input-output—Apply the rule match bidirectionally.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Match Direction for AACL Rules on page 841

Copyright © 2013, Juniper Networks, Inc.

rule

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application-group-any;
      application-groups [ application-group-names ];
      application-unknown;
      applications [ application-names ];
      destination-address address <any-unicast>;
      destination-address-range low minimum-value high maximum-value;
      destination-prefix-list list-name;
      nested-application-unknown;
      source-address address <any-unicast>;
      source-address-range low minimum-value high maximum-value;
      source-prefix-list list-name;
    }
    then {
      (accept | discard);
      count (application | application-group | application-group-any | nested-application |
        none);
      forwarding-class class-name;
      policer policer-name;
    }
  }
}
```

Hierarchy Level

[edit services aacl],
[edit services aacl rule-set *rule-set-name*]

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring AACL Rules on page 840

Copyright © 2013, Juniper Networks, Inc.

rule-set (Services AACL)

Syntax
```
rule-set rule-set-name {
  [rule rule-names ];
}
```

Hierarchy Level [edit services aacl]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring AACL Rule Sets on page 844

services (AACL)

Syntax services aacl { ... }

Hierarchy Level [edit]

Release Information **aacl** statement introduced in Junos OS Release 9.5.

Description Define the services to be applied to traffic.

Options **aacl**—The values configured for application-aware-access-list matching rules.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Application-Aware Access List

Copyright © 2013, Juniper Networks, Inc.

source-address (AACL)

Syntax source-address *address*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the source address for rule matching.

Options ***address***—Source IPv4 or IPv6 address or prefix value.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

source-address-range

Syntax source-address-range low *minimum-value* high *maximum-value*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.
IPv6 support introduced in Junos OS Release 12.2

Description Specify the source address range for rule matching.

Options ***minimum-value***—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Match Conditions in AACL Rules on page 841

Copyright © 2013, Juniper Networks, Inc.

source-prefix-list

Syntax source-prefix-list *list-name*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the source prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Source prefix list.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Match Conditions in AACL Rules on page 841

Copyright © 2013, Juniper Networks, Inc.

term

Syntax

```
term term-name {
  from {
    application-group-any;
    application-groups [ application-group-names ];
    application-unknown;
    applications [ application-names ];
    destination-address address <any-unicast>;
    destination-address-range low minimum-value high maximum-value;
    destination-prefix-list list-name;
    nested-application-unknown;
    source-address address <any-unicast>;
    source-address-range low minimum-value high maximum-value;
    source-prefix-list list-name;
  }
  then {
    (accept | discard);
    count (application | application-group | application-group-any | nested-application |
      none);
    forwarding-class class-name;
    policer policer-name;
  }
}
```

Hierarchy Level [edit services aacl rule *rule-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the AACL term properties.

Options ***term-name***—Identifier for the term.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring AACL Rules on page 840

Copyright © 2013, Juniper Networks, Inc.

then

Syntax

```
then {
  (accept | discard);
  count (application | application-group | application-group-any | nested-application | none);
  forwarding-class class-name;
  log event-type;
  policer policer-name;
}
```

Hierarchy Level [edit services aacl rule *rule-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 9.5.
policer statement added in Junos OS Release 9.6.
The **nested-application** option for the **count** statement introduced in Junos OS Release 11.1.

Description Define the AACL term actions. You can configure the router to accept or discard the targeted traffic. The action modifiers (count and forwarding-class) are optional.

Options You can configure one of the following actions:

- **accept**—Accept the packets and all subsequent packets in flows that match the rules.
- **discard**—Discard the packet and all subsequent packets in flows that match the rules.

When you select **accept** as the action, you can optionally configure one or both of the following action modifiers. No action modifiers are allowed with the **discard** action.

- **count (application | application-group | application-group-any | nested-application | none)**—For all accepted packets that match the rules, record a packet count using AACL statistics practices. You can specify one of the following options; there is no default setting:
  - **application**—Count the application that matched in the **from** clause.
  - **application-group**—Count the application group that matched in the **from** clause.
  - **application-group-any**—Count all application groups that match **from application-group-any** under the **any** group name.
  - **nested-application**—Count all nested applications that matched in the **from** clause.
  - **none**—Same as not specifying **count** as an action.
- **forwarding-class *class-name***—Specify the packets' forwarding-class name.

policer *policer-name*—Apply rate-limiting properties to the traffic as configured at the **[edit firewall policer *policer-name*]** hierarchy level. This configuration allows bit-rate and burst-size attributes to be applied to the traffic that are not supported by AACL rules. When you include a policer, the only allowed action is **discard**. For more information on policers, see the Routing Policy Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring AACL Rules on page 840
- Routing Policy Configuration Guide

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 41

Local Policy Decision Function Configuration Guidelines

Starting with Junos OS Release 12.1, the local policy descision function (L-PDF) plug-in can offload flows to the Packet Forwarding Engine. Offloading is supported only on MX Series routers with Modular Port Concentrators (MPCs) and accomplished using the Juniper Forwarding Mechanism (JFM). JFM allows services flows to be offloaded to the Packet Forwarding Engine. However, 5-tuple flows cannot be offloaded. Apart from the local L-PDF plug-in, offloading is supported on the packet-triggered subscribers and policy control (PTSP) plug-in. The **show services application-aware-access-list flows subscriber *subscriber-name*** command displays offload status.

This chapter includes the following sections:

- Configuring Statistics Profiles on page 859
- Applying L-PDF Profiles to Service Sets on page 863
- Tracing L-PDF Operations on page 864

Configuring Statistics Profiles

The local policy decision function (L-PDF) enables you to configure properties for statistics output. To do this, you create a statistics profile, which configures the files to which statistics records are exported and the format that is exported. There are two configurations you can use to specify the profile, as described in the following subsections:

- Configuring an L-PDF Statistics Profile on page 860
- Configuring an AACL Statistics Profile on page 861



NOTE: You must use the same configuration stanza for specifying the profile and the file selection. If configurations are committed in both hierarchies, the one at the **[edit system services local-policy-decision-function]** hierarchy level takes precedence.

Copyright © 2013, Juniper Networks, Inc.

NOTE:

- When a session closes before APPID has identified nested applications, the session is treated as a best-effort session and L-PDF does not get the nested application information. In such cases, nested applications will be reported as unknown applications.
- During the time that the application identification (APPID) feature has not yet made a final determination of the application associated with a given flow, the flow does not contribute to any per-subscriber or per-application statistics collection. For more information, see "Best-Effort Application Identification of DPI-Serviced Flows" on page 777.

NOTE: For rms- interfaces, the statistics received from the active Multiservices PICs in the RMS group are combined with the statistics of the reported ended flows kept on the Routing Engine. The aggregated value is written to the statistics file. In the case of AMS interfaces, all the Multiservices PICs consisting of the AMS group reports statistics independently. These statistics are aggregated on the Routing Engine. The Routing Engine runs an independent timer, which on expiry writes the aggregated entry in the statistics file. This method of collection causes the statistics data in the statistics file to be displayed with a small delay.

Configuring an L-PDF Statistics Profile

You can specify an L-PDF statistics profile by including the following configuration at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
policy-decision-statistics-profile profile-name {
   application-aware-access-list-fields [ field-name ];
   file filename;
   files number;
   size bytes;
}
```

NOTE: This configuration method is not the preferred method for configuring Dynamic Application Awareness statistics. It is only maintained for backwards compatibility and may be deprecated in a future software release and does not support the use of IPv6 address and prefix length. The new, preferred configuration is found at the **[edit system services local-policy-decision-function]** hierarchy level, as described in "Configuring an AACL Statistics Profile" on page 861. We encourage you to migrate to the new configuration method.

You specify a profile name to identify the profile and other properties as needed by including the **policy-decision-statistics-profile** statement. The **aacl-fields** statement

Copyright © 2013, Juniper Networks, Inc.

specifies which statistics to collect in an accounting-data log file. This log file is located on the **/var/log** directory on the router. You specify the log file by including the **file** ***filename*** statement. The filename is prefixed by the **aacl_statistics_** prefix; for example, if you specify the filename **lpdfd**, the log file will be **/var/log/aacl_statistics_lpdfd**.

The **application-aware-access-list-fields** statement supports the following options:

- **address**—IP Address
- **application**—Application name
- **application-group**—Application group name
- **input-bytes**—Number of input bytes
- **input-interface**—Input interface name
- **input-packets**—Number of input packets
- **mask**—Netmask
- **output-bytes**—Number of output bytes
- **output-packets**—Number of output packets
- **subscriber-name**—Subscriber name
- **timestamp**—Timestamp
- **vrf-name**—VPN routing and forwarding (VRF) name

For more information on configuring profiles, see the Network Management Configuration Guide.

Configuring an AACL Statistics Profile

You can specify an AACL statistics profile by including the following configuration at the **[edit system services]** hierarchy level:

```
local-policy-decision-function {
    statistics {
        file filename {
            archive-sites [ url ];
            files number;
            size bytes;
            transfer-interval minutes;
        }
        aacl-statistics-profile profile-name {
            aacl-fields [ field-name ];
            file filename;
            report-interval minutes;
            record-mode (interim-active-only | interim-full);
        }
        record-type (delta | interim);
    }
}
```

To specify the file properties, include the **file** statement at the **[edit system services local-policy-decision-function statistics** hierarchy level with a unique filename:

Copyright © 2013, Juniper Networks, Inc.

- The **archive-sites** statement specifies one or more URLs for archiving the files. Archiving can be done by using FTP or SCP.
- The **files** statement specifies the maximum number of files that are maintained at one time.
- The **size** statement specifies the maximum size of each file.
- The **transfer-interval** statement specifies the interval between data transfers in minutes.

You specify a profile name to identify the profile and other properties as needed by including the **aacl-statistics-profile** statement. The **aacl-fields** statement specifies which statistics to collect in an accounting-data log file. This log file is located on the **/var/stats/aacl** directory on the router. You specify the log file by including the **file *filename*** statement.

The **aacl-fields** statement supports the following options:

- **address**—IP Address
- **all-fields**—All available fields
- **application**—Application name
- **application-group**—Application group name
- **input-bytes**—Number of input bytes
- **input-interface**—Input interface name
- **ipv6-address**—IPv6 address
- **ipv6-prefix-length**—Prefix length associated with the displayed IPv6 address
- **input-packets**—Number of input packets
- **mask**—Netmask
- **output-bytes**—Number of output bytes
- **output-packets**—Number of output packets
- **subscriber-name**—Subscriber name
- **timestamp**—Timestamp
- **vrf-name**—VPN routing and forwarding (VRF) name

The **record-type** statement specifies whether a record is **delta** or **interim**; **delta** is the default setting. The **report-interval** statement specifies the reporting interval in minutes; the default setting is 15 minutes and the range is 5 through 1440 minutes. The **record-mode** statement specifies how the statistics are reported for each reporting interval; the default setting is **interim-full** and reports all available statistics. To report only statistics that have changed for the reporting interval, use the **interim-active-only** setting.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The IPv6 fields (**ipv6-address** and **ipv6-prefix-length**) are not supported for **record-type delta**. The IPv6 fields are supported for **record-type interim** only,meaning that the fields are restricted to the S- (Login) record.

For more information on configuring profiles, see the Network Management Configuration Guide.

Applying L-PDF Profiles to Service Sets

You can optionally apply policy decision statistics profiles as part of a service-set definition. To do this, you include the **policy-decision-statistics-profile** statement at the **[edit services service-set *service-set-name*]** hierarchy level:

```
policy-decision-statistics-profile profile-name;
```

NOTE: To provide high availability for the policy decision statistics, associate the service-set definition with a redundant services PIC (rsp) interface.

You can include only one profile name in the specification for the **application-aware access-list** statement.

The following example shows a sample configuration for attachment of an L-PDF statistics profile:

```
services {
   service-set test_aacl_sset {
      aacl-rules aacl_rule;
      policy-decision-statistics-profile {
         pdf_stats_prof;
      }
      interface-service {
         service-interface ms-0/3/0.0;
      }
   }
}
```

NOTE: Only one service set can be applied to a single interface when L-PDF functionality is used.

The following example shows a sample configuration for attachment of a service set to a static interface:

```
interfaces {
   fe-0/0/0 {
      vlan-tagging;
      unit 1 {
         vlan-id 1;
         family inet {
```

Copyright © 2013, Juniper Networks, Inc.

```
            service {
              input {
                service-set test_aacl_sset;
              }
              output {
                service-set test_aacl_sset;
              }
            }
            address 10.1.1.1/24;
          }
        }
      }
    }
```



NOTE: The **session-offload** statement at the **[edit chassis fpc** *slot-number* **pic** *number* **adaptive-services service-package extension-provider]** hierarchy level controls session offload behavior for Multiservices DPCs on MX Series routers. It controls session offload on a per-device basis, where a device is a Multiservices interface (**ms-***fpc-pic-port*). Currently, the session offload function is supported for at most one Multiservices interface. When the offload function is enabled, it is strongly recommended that you limit Dynamic Application Awareness features to that Multiservices interface.

The default is to not offload any sessions. For more information on chassis configuration, see the Junos OS System Basics Configuration Guide.

Tracing L-PDF Operations

Tracing operations track L-PDF operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit system services local-policy-decision-function]** hierarchy level, you can customize the trace file settings:

```
traceoptions {
  file filename <files number> <size size>;
  flag flag;
}
```

The flags track the following information:

- **all**—Everything
- **configuration**—Configuration traces
- **database**—Database traces
- **general**—Miscellaneous traces
- **gres**—Graceful Routing Engine switchover (GRES) traces
- **ptsp-statistics**—PTSP statistics traces

Copyright © 2013, Juniper Networks, Inc.

- **rtsock**—Routing socket traces
- **statistics**—Statistics traces
- **subscriber**—Subscriber traces

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 42

Summary of L-PDF Configuration Statements

The following sections explain each of the local policy decision function (L-PDF) statements. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

aacl-fields

Syntax

```
aacl-fields {
  field-name;
}
```

Hierarchy Level [edit system services local-policy-decision-function statistics aacl-statistics-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0.
IPv6 support introduced in Junos OS Release 12.2

Description Define the statistics to collect in a data log file.

Options ***field-name***—Name of the field:

- **address**—IPv4 address
- **all-fields**—All available fields
- **application**—Application name
- **application-group**—Application group name
- **input-bytes**—Number of input bytes
- **input-interface**—Input interface name
- **input-packets**—Number of input packets
- **ipv6-address**—IPv6 address
- **ipv6-prefix-length**—Prefix length associated with the displayed IPv6 address
- **mask**—Netmask
- **output-bytes**—Number of output bytes
- **output-packets**—Number of output packets
- **subscriber-name**—Subscriber name
- **timestamp**—Timestamp
- **vrf-name**—VPN routing and forwarding (VRF) name

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Statistics Profiles on page 859

Copyright © 2013, Juniper Networks, Inc.

aacl-statistics-profile

Syntax

```
aacl-statistics-profile profile-name {
  aacl-fields {
    field-name;
  }
  file filename;
  record-mode (interim-active-only | interim-full);
  report-interval minutes;
}
```

Hierarchy Level [edit services service-set *service-set-name*],
[edit system services local-policy-decision-function statistics]

Release Information Statement introduced in Junos OS Release 10.0.
record-mode option introduced in Junos OS Release 10.2.

Description Create an AACL statistics profile, which configures the files to which statistics records are exported and the format that is exported.

Options

file *filename*—Name of the file to receive the statistics data output. Enclose the name within quotation marks. All files are placed in the directory **/var/stats/aacl**.

record-mode—Record mode for the reporting interval; possible values are **interim-active-only**, which reports only statistics that have changed, or **interim-full**, which reports all available statistics.

report-interval *minutes*—Frequency at which statistics are recorded, in minutes.

Default: 15 minutes

Range: 5 through 1440 minutes

The remaining statements are explained separately.

Usage Guidelines See "Configuring Statistics Profiles" on page 859.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- For more information on profiles, see the Network Management Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

application-aware-access-list-fields

Syntax

```
application-aware-access-list-fields {
  field-name;
}
```

Hierarchy Level [edit accounting-options policy-decision-statistics-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define the statistics to collect in a data log file.

Options ***field-name***—Name of the field:

- **address**—IP address
- **application**—Application name
- **application-group**—Application group name
- **input-bytes**—Number of input bytes
- **input-interface**—Input interface name
- **input-packets**—Number of input packets
- **mask**—Netmask
- **output-bytes**—Number of output bytes
- **output-packets**—Number of output packets
- **subscriber-name**—Subscriber name
- **timestamp**—Timestamp
- **vrf-name**—VPN routing and forwarding (VRF) name

Usage Guidelines See "Configuring Statistics Profiles" on page 859.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

file

Syntax

```
file file-name {
  archive-sites url;
  files file-number;
  size bytes;
  transfer-interval minutes;
}
```

Hierarchy Level [edit system services local-policy-decision-function statistics]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a file to which statistics records are exported and the format that is exported.

Options

archive-sites [*url*]—One or more destinations for archiving data.

filename—Name of the file to receive the statistics data output.

files *number*—(Optional) Maximum number of accounting files.

Range: 3 through 1000 files

Default: 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 262144 through 1073741824 or the maximum file size supported on your system

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

transfer-interval *minutes*—Frequency at which to transfer files to archive sites, in minutes.

Usage Guidelines See "Configuring Statistics Profiles" on page 859.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

local-policy-decision-function

Syntax

```
local-policy-decision-function {
  statistics {
    aacl-statistics-profile profile-name {
      aacl-fields {
        field-name;
      }
      file filename;
      report-interval minutes;
    }
    file file-name {
      archive-sites url;
      files file-number;
      size bytes;
      transfer-interval minutes;
    }
    record-type (delta | interim);
  }
  traceoptions {
    file filename <files number> <size size>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify L-PDF properties.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Statistics Profiles" on page 859.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

policy-decision-statistics-profile

Syntax

```
policy-decision-statistics-profile profile-name {
  aacl-fields {
    field-name;
  }
  file filename;
  files file-number;
  size bytes;
}
```

Hierarchy Level

[edit accounting-options],
[edit services service-set *service-set-name*]

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Create a policy decision statistics profile, which configures the files to which statistics records are exported and the format that is exported.

Options

file *filename*—Name of the file to receive the accounting-data output. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of accounting files.

Range: 2 through 1000 files

Default: 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

profile-name—Name of the policy decision statistics profile.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10240 through 1073741824 or the maximum file size supported on your system

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring Statistics Profiles" on page 859.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- For more information on profiles, see the Network Management Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

statistics (System Services)

Syntax

```
statistics {
  aacl-statistics-profile profile-name {
    aacl-fields {
      field-name;
    }
    file filename;
    report-interval minutes;
  }
  file file-name {
    archive-sites [ url ];
    files file-number;
    size bytes;
    transfer-interval minutes;
  }
  record-type (delta | interim);
}
```

Hierarchy Level [edit system services local-policy-decision-function]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure file and data specifications for recording AACL statistics.

Options **record-type**—Record type; possible values are **delta** or **interim**.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Statistics Profiles" on page 859.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Services Local Policy Decision Function)

Syntax

```
traceoptions {
  file filename <files number> <size size>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit services local-policy-decision-function],
[edit system services local-policy-decision-function]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure local policy decision function (L-PDF) tracing options.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000 files

Default: 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag—Tracing operation to perform. To specify more than one flag, include multiple **flag** statements.

- **all**—Everything
- **configuration**—Configuration traces
- **database**—Database traces
- **general**—Miscellaneous traces
- **gres**—Graceful Routing Engine switchover (GRES) traces
- **ptsp-statistics**—PTSP statistics traces
- **rtsock**—Routing socket traces
- **statistics**—Statistics traces
- **subscriber**—Subscriber traces

no-remote-trace—Disable remote tracing.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed

Copyright © 2013, Juniper Networks, Inc.

trace-file*.1** and ***trace-file is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10240 through 1073741824 or the maximum file size supported on your system

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Usage Guidelines See "Tracing L-PDF Operations" on page 864.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

PART 4

Encryption Services

- Encryption Overview on page 879
- Encryption Interfaces Configuration Guidelines on page 881
- Summary of Encryption Configuration Statements on page 891

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 43

Encryption Overview

This chapter discusses the following topics:

- Encryption Overview on page 879

Encryption Overview

The IP Security (IPsec) architecture provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates.

IPsec defines a security association (SA) and key management framework that can be used with any network layer protocol. The SA specifies what protection policy to apply to traffic between two IP-layer entities. For more information, see the Junos OS System Basics Configuration Guide. The standards are defined in the following RFCs:

- RFC 2401, *Security Architecture for the Internet Protocol*
- RFC 2406, *IP Encapsulating Security Payload (ESP)*

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 44

Encryption Interfaces Configuration Guidelines

To enable encryption interfaces, you can configure the following properties:

- Configuring Encryption Interfaces on page 881
- Configuring Filters for Traffic Transiting the ES PIC on page 883
- Configuring an ES Tunnel Interface for a Layer 3 VPN on page 888
- Configuring ES PIC Redundancy on page 888
- Configuring IPsec Tunnel Redundancy on page 889

Configuring Encryption Interfaces

When you configure the encryption interface, you associate the configured SA with a logical interface. This configuration defines the tunnel, including the logical unit, tunnel addresses, maximum transmission unit (MTU), optional interface addresses, and the name of the IPsec SA to apply to traffic. To configure an encryption interface, include the following statements at the **[edit interfaces es-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
family inet {
   ipsec-sa ipsec-sa; # name of security association to apply to packet
   address address; # local interface address inside local VPN
   destination address; # destination address inside remote VPN
}
tunnel {
   source source-address;
   destination destination-address;
}
```

The addresses configured as the tunnel source and destination are the addresses in the outer IP header of the tunnel.

Copyright © 2013, Juniper Networks, Inc.



NOTE: You must configure the tunnel source address locally on the router, and the tunnel destination address must be a valid address for the security gateway terminating the tunnel.

The ES Physical Interface Card (PIC) is supported on M Series and T Series routers.

The SA must be a valid tunnel-mode SA. The interface address and destination address listed are optional. The destination address allows the user to configure a static route to encrypt traffic. If a static route uses that destination address as the next hop, traffic is forwarded through the portion of the tunnel in which encryption occurs.

Specifying the Security Association Name for Encryption Interfaces

The security association is the set of properties that defines the protocols for encrypting Internet traffic. To configure encryption interfaces, you specify the SA name associated with the interface by including the **ipsec-sa** statement at the **[edit interfaces es-*fpc/pic/port* unit *logical-unit-number* family inet]** hierarchy level:

```
ipsec-sa sa-name;
```

For information about configuring the security association, see "Configuring Filters for Traffic Transiting the ES PIC" on page 883.

Configuring the MTU for Encryption Interfaces

The protocol MTU value for encryption interfaces must always be less than the default interface MTU value of 3900 bytes; the configuration fails to commit if you select a greater value. To set the MTU value, include the **mtu** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
mtu bytes;
```

For more information, see the Junos® OS Network Interfaces.

Example: Configuring an Encryption Interface

Configure an IPsec tunnel as a logical interface on the ES PIC. The logical interface specifies the tunnel through which the encrypted traffic travels. The **ipsec-sa** statement associates the security profile with the interface.

```
[edit interfaces]
es-0/0/0 {
  unit 0 {
    tunnel {
      source 10.5.5.5; # tunnel source address
      destination 10.6.6.6; # tunnel destination address
    }
    family inet {
      ipsec-sa manual-sa1; # name of security association to apply to packet
      mtu 3800;
      address 10.1.1.8/32 { # local interface address inside local VPN
      destination 10.2.2.254;  # destination address inside remote VPN
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
    }
}
```

Configuring Filters for Traffic Transiting the ES PIC

This section contains the following topics:

- Traffic Overview on page 883
- Configuring the Security Association on page 884
- Configuring an Outbound Traffic Filter on page 885
- Applying the Outbound Traffic Filter on page 886
- Configuring an Inbound Traffic Filter on page 886
- Applying the Inbound Traffic Filter to the Encryption Interface on page 887

Traffic Overview

Traffic configuration defines the traffic that must flow through the tunnel. You configure outbound and inbound firewall filters, which identify and direct traffic to be encrypted and confirm that decrypted traffic parameters match those defined for the given tunnel. The outbound filter is applied to the LAN or WAN interface for the incoming traffic you want to encrypt. The inbound filter is applied to the ES PIC to check the policy for traffic coming in from the remote host. Because of the complexity of configuring a router to forward packets, no automatic checking is done to ensure that the configuration is correct.

NOTE: **The valid firewall filters statements for IPsec are destination-port, source-port, protocol, destination-address, and source-address.**

In Figure 7 on page 883, Gateway A protects the network **10.1.1.0/24**, and Gateway B protects the network **10.2.2.0/24**. The gateways are connected by an IPsec tunnel. For more information about firewalls, see the Routing Policy Configuration Guide.

Figure 7: Example: IPsec Tunnel Connecting Security Gateways

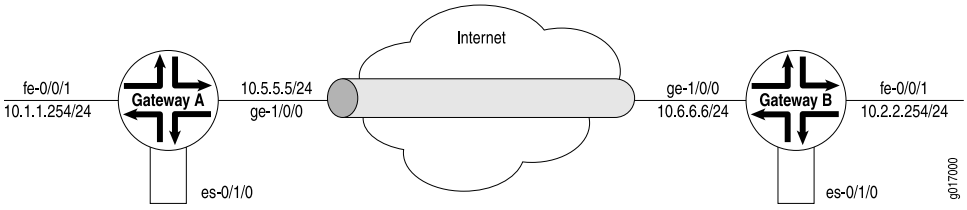

The SA and ES interface for security Gateway A are configured as follows:

```
[edit security ipsec]
security-association manual-sa1 {
  manual {
    direction bidirectional {
      protocol esp;
      spi 2312;
      authentication {
```

```
            algorithm hmac-md5-96;
            key ascii-text 1234123412341234;
          }
          encryption {
            algorithm 3des-cbc;
            key ascii-text 123456789009876543211234;
          }
        }
      }
    }
    [edit interfaces es-0/1/0]
    unit 0 {
      tunnel {
        source 10.5.5.5;
        destination 10.6.6.6;
      }
      family inet {
        ipsec-sa manual-sa1;
        address 10.1.1.8/32 {
          destination 10.2.2.254;
        }
      }
    }
```

Configuring the Security Association

To configure the SA, include the **security-association** statement at the **[edit security]** hierarchy level:

```
security-association name {
  mode (tunnel | transport);
  manual {
    direction (inbound | outbound | bi-directional) {
      auxiliary-spi auxiliary-spi-value;
      spi spi-value;
      protocol (ah | esp | bundle);
      authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      encryption {
        algorithm (des-cbc | 3des-cbc);
        key (ascii-text key | hexadecimal key);
      }
    }
    dynamic {
      replay-window-size (32 | 64);
      ipsec-policy policy-name;
    }
  }
}
```

For more information about configuring an SA, see the Junos OS System Basics Configuration Guide. For information about applying the SA to an interface, see "Specifying the Security Association Name for Encryption Interfaces" on page 882.

Copyright © 2013, Juniper Networks, Inc.

Configuring an Outbound Traffic Filter

To configure the outbound traffic filter, include the **filter** statement at the **[edit firewall]** hierarchy level:

```
filter filter-name {
   term term-name {
      from {
         match-conditions;
      }
      then {
         action;
         action-modifiers;
      }
   }
}
```

For more information, see the Routing Policy Configuration Guide.

Example: Configuring an Outbound Traffic Filter

Firewall filters for outbound traffic direct the traffic through the desired IPsec tunnel and ensure that the tunneled traffic goes out the appropriate interface (see Figure 7 on page 883). Here, an outbound firewall filter is created on security Gateway A; it identifies the traffic to be encrypted and adds it to the input side of the interface that carries the internal virtual private network (VPN) traffic:

```
[edit firewall]
filter ipsec-encrypt-policy-filter {
   term term1 {
      from {
         source-address { # local network
         10.1.1.0/24;
      }
      destination-address { # remote network
      10.2.2.0/24;
   }
}
then ipsec-sa manual-sa1; # apply SA name to packet
term default {
   then accept;
}
```



NOTE: The source address, port, and protocol on the outbound traffic filter must match the destination address, port, and protocol on the inbound traffic filter. The destination address, port, and protocol on the outbound traffic filter must match the source address, port, and protocol on the inbound traffic filter.

Copyright © 2013, Juniper Networks, Inc.

Applying the Outbound Traffic Filter

After you have configured the outbound firewall filter, you apply it by including the **filter** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
filter {
  input filter-name;
}
```

Example: Applying the Outbound Traffic Filter

Apply the outbound traffic filter. The outbound filter is applied on the Fast Ethernet interface at the **[edit interfaces fe-0/0/1 unit 0 family inet]** hierarchy level. Any packet matching the IPsec action term (**term 1**) on the input filter (**ipsec-encrypt-policy-filter**), configured on the Fast Ethernet interface, is directed to the ES PIC interface at the **[edit interfaces es-0/1/0 unit 0 family inet]** hierarchy level. So, if a packet arrives from the source address **10.1.1.0/24** and goes to the destination address **10.2.2.0/24**, the Packet Forwarding Engine directs the packet to the ES PIC interface, which is configured with the **manual-sa1** SA. The ES PIC receives the packet, applies the **manual-sa1** SA, and sends the packet through the tunnel.

The router must have a route to the tunnel end point; add a static route if necessary.

```
[edit interfaces]
fe-0/0/1 {
  unit 0 {
    family inet {
      filter {
        input ipsec-encrypt-policy-filter;
      }
      address 10.1.1.254/24;
    }
  }
}
```

Configuring an Inbound Traffic Filter

To configure an inbound traffic filter, include the **filter** statement at the **[edit firewall]** hierarchy level:

```
filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
```

For more information, see the Routing Policy Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring an Inbound Traffic Filter

Configure an inbound firewall filter. This filter performs the final IPsec policy check and is created on security gateway A. The policy check ensures that only packets that match the traffic configured for this tunnel are accepted.

```
[edit firewall]
filter ipsec-decrypt-policy-filter {
  term term1 { # perform policy check
  from {
    source-address { # remote network
    10.2.2.0/24;
  }
  destination-address { # local network
  10.1.1.0/24;
}
then accept;
```

Applying the Inbound Traffic Filter to the Encryption Interface

After you create the inbound firewall filter, you can apply it to the ES PIC. To apply the filter to the ES PIC, include the **filter** statement at the **[edit interfaces es-*fpc/pic/port* unit *logical-unit-number* family inet filter]** hierarchy level:

```
filter {
  input filter;
}
```

The input filter is the name of the filter applied to received traffic. For a configuration example, see "Example: Configuring an Inbound Traffic Filter" on page 887. For more information about firewall filters, see the Routing Policy Configuration Guide.

Example: Applying the Inbound Traffic Filter to the Encryption Interface

Apply the inbound firewall filter (**ipsec-decrypt-policy-filter**) to the decrypted packet to perform the final policy check. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and decrypts the incoming packet.

The Packet Forwarding Engine directs IPsec packets to the ES PIC. It uses the packet's security parameter index (SPI), protocol, and destination address to look up the SA configured on one of the ES interfaces. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and is used to decrypt the incoming packet. When the packets are processed (decrypted, authenticated, or both), the input firewall filter (**ipsec-decrypt-policy-filter**) is applied on the decrypted packet to perform the final policy check. **term1** defines the decrypted (and verified) traffic and performs the required policy check. For information about **term1**, see "Example: Configuring an Inbound Traffic Filter" on page 887.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The inbound traffic filter is applied after the ES PIC has processed the packet, so the decrypted traffic is defined as any traffic that the remote gateway is encrypting and sending to this router. IKE uses this filter to determine the policy required for a tunnel. This policy is used during the negotiation with the remote gateway to find the matching SA configuration.

```
[edit interfaces]
es-1/2/0 {
  unit 0 {
    tunnel {
      source 10.5.5.5; # tunnel source address
      destination 10.6.6.6; # tunnel destination address
    }
    family inet {
      filter {
        input ipsec-decrypt-policy-filter;
      }
      ipsec-sa manual-sa1; # SA name applied to packet
      address 10.1.1.8/32 { # local interface address inside local VPN
      destination 10.2.2.254; # destination address inside remote VPN
    }
  }
}
```

Configuring an ES Tunnel Interface for a Layer 3 VPN

To configure an ES tunnel interface for a Layer 3 VPN, you need to configure an ES tunnel interface on the provider edge (PE) router and on the customer edge (CE) router. You also need to configure IPsec on the PE and CE routers. For more information about configuring an ES tunnel for a Layer 3 VPN, see the Junos OS VPNs Configuration Guide.

Configuring ES PIC Redundancy

You can configure ES PIC redundancy on M Series and T Series routers that have multiple ES PICs. With ES PIC redundancy, one ES PIC is active and another ES PIC is on standby. When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and SAs, and acts as the new next hop for IPsec traffic. Reestablishment of tunnels on the backup ES PIC does not require new Internet Key Exchange (IKE) negotiations. If the primary ES PIC comes online, it remains in standby and does not preempt the backup. To determine which PIC is currently active, use the **show ipsec redundancy** command.

NOTE: ES PIC redundancy is supported on M Series and T Series routers.

To configure an ES PIC as the backup, include the **backup-interface** statement at the **[edit interfaces *fpc/pic/port* es-options]** hierarchy level:

backup-interface es-*fpc/pic/port*;

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring ES PIC Redundancy

After you create the inbound firewall filter, apply it to the master ES PIC. Here, the inbound firewall filter (**ipsec-decrypt-policy-filter**) is applied on the decrypted packet to perform the final policy check. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and decrypts the incoming packet. This example does not show SA and filter configuration. For information about SA and filter configuration, see the Junos OS System Basics Configuration Guide, the Routing Policy Configuration Guide, and "Example: Configuring an Inbound Traffic Filter" on page 887.

```
[edit interfaces]
es-1/2/0 {
   es-options {
      backup-interface es-1/0/0;
   }
   unit 0 {
      tunnel {
         source 10.5.5.5;
         destination 10.6.6.6;
      }
      family inet {
         ipsec-sa manual-sa1;
         filter {
            input ipsec-decrypt-policy-filter;
         }
         address 10.1.1.8/32 {
            destination 10.2.2.254;
         }
      }
   }
}
```

Configuring IPsec Tunnel Redundancy

You can configure IPsec tunnel redundancy by specifying a backup destination address. The local router sends keepalives to determine the remote site's reachability. When the peer is no longer reachable, a new tunnel is established. For up to 60 seconds during failover, traffic is dropped without notification being sent. Figure 8 on page 889 shows IPsec primary and backup tunnels.

Figure 8: IPsec Tunnel Redundancy

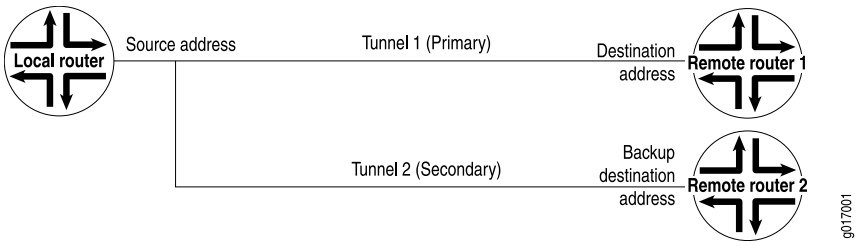

To configure IPsec tunnel redundancy, include the **backup-destination** statement at the **[edit interfaces unit *logical-unit-number* tunnel]** hierarchy level:

backup-destination*address*;
destination *address*;
source *address*;



NOTE: Tunnel redundancy is supported on M Series and T Series routers.

The primary and backup destinations must be on different routers.

The tunnels must be distinct from each other and policies must match.

For more information about tunnels, see Tunnel Properties.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 45

Summary of Encryption Configuration Statements

The following sections explain each of the encryption services statements. The statements are organized alphabetically.

address (Interfaces)

Syntax

```
address address {
  destination address;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the interface address.

Options ***address***—Address of the interface.

The remaining statement is explained separately.

Usage Guidelines See "Configuring Encryption Interfaces" on page 881.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

backup-destination

See backup-destination

Copyright © 2013, Juniper Networks, Inc.

backup-interface

Syntax backup-interface *interface-name*;

Hierarchy Level [edit interfaces *interface-name* es-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a backup ES Physical Interface Card (PIC). When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and security associations (SAs), and acts as the new next hop for IPsec traffic.

Options ***interface-name***—Name of ES interface to serve as the backup.

Usage Guidelines See "Configuring ES PIC Redundancy" on page 888.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination (Interfaces)

Syntax destination *address*;

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number* tunnel]
[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],
[edit interfaces *interface-name* unit *logical-unit-number* tunnel]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For CoS on ATM interfaces, specify the remote address of the connection.

For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.

For tunnel and encryption interfaces, specify the remote address of the tunnel.

Options ***address***—Address of the remote side of the connection.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Linear RED Profiles on ATM Interfaces
- Multilink and Link Services Logical Interface Configuration Overview on page 1136
- Configuring Encryption Interfaces on page 881
- Configuring Traffic Sampling on page 912
- Configuring Flow Monitoring on page 920
- Configuring Unicast Tunnels on page 1273

Copyright © 2013, Juniper Networks, Inc.

es-options

Syntax

```
es-options {
  backup-interface interface-name;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description On ES interfaces, configure ES interface-specific interface properties.

The **backup-interface** statement is explained separately.

Usage Guidelines See "Configuring ES PIC Redundancy" on page 888.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

family

Syntax

```
family inet {
  ipsec-sa sa-name;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure protocol family information for the logical interface.

Options *family*—Protocol family:

- **ccc**—Circuit cross-connect protocol suite
- **inet**—IP version 4 suite
- **inet6**—IP version 6 suite
- **iso**—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—MPLS
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpls**—Virtual private LAN service

The remaining statements are explained separately.

Usage Guidelines See "Configuring Encryption Interfaces" on page 881; for a general discussion of **family** statement options, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

filter

Syntax

```
filter {
  input filter-name;
  output filter-name;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the filters to be applied on an interface.

Options **input *filter-name***—Identifier for the input filter.

output *filter-name*—Identifier for the output filter.

Usage Guidelines See "Configuring Filters for Traffic Transiting the ES PIC" on page 883.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines See the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ipsec-sa

Syntax ipsec-sa *sa-name*;

Hierarchy Level [edit interfaces es-*fpc/pic/port* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the IP Security (IPsec) SA name associated with the interface.

Options ***sa-name***—IPsec SA name.

Usage Guidelines See "Configuring Encryption Interfaces" on page 881.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Junos OS System Basics Configuration Guide

source

Syntax source *source-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],
[edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.

Description For tunnel and encryption interfaces, specify the source address.

Options ***source-address***—Address of the source side of the connection.

Usage Guidelines See "Configuring Encryption Interfaces" on page 881, "Configuring Traffic Sampling" on page 912, and "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

tunnel

Syntax

```
tunnel {
  backup-destination destination-address;
  destination destination-address;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).

The statements are explained separately.

Usage Guidelines See "Configuring Encryption Interfaces" on page 881 and Tunnel Properties.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos OS VPNs Configuration Guide

Copyright © 2013, Juniper Networks, Inc.

unit (Interfaces)

Syntax

```
unit logical-unit-number {
  family inet {
    ipsec-sa sa-name;
  }
  tunnel {
    backup-destination destination-address;
    destination destination-address;
    routing-instance {
      destination routing-instance-name;
    }
    source source-address;
    ttl number;
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number***—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Usage Guidelines See "Configuring Encryption Interfaces" on page 881; for a general discussion of logical interface properties, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 5

Flow Monitoring and Discard Accounting Services

- Flow Monitoring and Discard Accounting Overview on page 903
- Flow Monitoring and Discard Accounting Configuration Guidelines on page 907
- Summary of Flow-Monitoring Configuration Statements on page 977
- Flow Collection Configuration Guidelines on page 1053
- Summary of Flow Collection Configuration Statements on page 1065
- Dynamic Flow Capture Configuration Guidelines on page 1083
- Flow-Tap Configuration Guidelines on page 1095
- Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements on page 1103

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 46

Flow Monitoring and Discard Accounting Overview

Using a Juniper Networks M Series Multiservice Edge or T Series Core Router, a selection of PICs (including the Monitoring Services PIC, Adaptive Services [AS] PIC, Multiservices PIC, or Multiservices DPC) and other networking hardware, you can monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to do the following:

- Gather and export detailed information about IP version 4 (IPv4) traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.
- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format (port mirror).



NOTE: Monitoring Services PICs, AS PICs, and Multiservices PICs must be mounted on an Enhanced Flexible PIC Concentrator (FPC) in an M Series or T Series router.

Multiservices DPCs installed in Juniper Networks MX Series 3D Universal Edge Routers support the same functionality, with the exception of the passive monitoring and flow-tap features.

This section provides general information on the following topics:

- Passive Flow Monitoring Overview on page 903
- Active Flow Monitoring Overview on page 904

Passive Flow Monitoring Overview

The router used for passive monitoring does not route packets from the monitored interface, nor does it run any routing protocols related to those interfaces; it only receives traffic flows, collects intercepted traffic, and exports it to cflowd servers and packet

Copyright © 2013, Juniper Networks, Inc.

analyzers. Figure 9 on page 904 shows a typical topology for the passive flow-monitoring application.

Figure 9: Passive Monitoring Application Topology

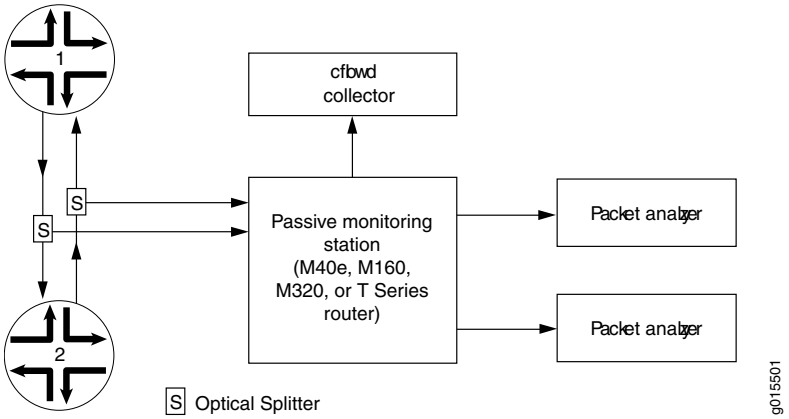

Traffic travels normally between Router 1 and Router 2. To redirect IPv4 traffic, you insert an optical splitter on the interface between these two routers. The optical splitter copies and redirects the traffic to the monitoring station, which is an M40e, M160, M320, or T Series router. The optical cable connects only the receive port on the monitoring station, never the transmit port. This configuration allows the monitoring station to receive traffic from the router being monitored but never to transmit it back.

If you are monitoring traffic flow, the Internet Processor II application-specific integrated circuit (ASIC) in the router forwards a copy of the traffic to the Monitoring Services, Adaptive Services, or Multiservices PIC in the monitoring station. If more than one monitoring PIC is installed, the monitoring station distributes the load of the incoming traffic across the multiple PICs. The monitoring PICs generate flow records in cflowd version 5 format, and the records are then exported to the cflowd collector.

If you are performing lawful interception of traffic between the two routers, the Internet Processor II ASIC filters the incoming traffic and forwards it to the Tunnel Services PIC. Filter-based forwarding is then applied to direct the traffic to the packet analyzers.

Optionally, the intercepted traffic or the cflowd records can be encrypted by the ES PIC or IP Security (IPsec) services and then sent to a cflowd server or packet analyzer.

Active Flow Monitoring Overview

Although the Monitoring Services PIC was designed initially for use as an offline passive flow monitoring tool, it can also be used in an active flow monitoring topology. In contrast, the AS or Multiservices PIC is designed exclusively for active flow monitoring. To use either the Monitoring Services PIC, AS PIC, or Multiservices PIC for active flow monitoring, you must install the PIC in an M Series or T Series router. The router participates in both the monitoring application and in the normal routing functionality of the network.

Starting with Junos OS Release 11.4, support for active monitoring is extended to logical systems running on T Series and MX Series routers. A logical system is a partition created from a physical router that performs independent routing tasks. Several logical systems

Copyright © 2013, Juniper Networks, Inc.

in a single router with their own interfaces, policies, instances, and routing tables can perform functions handled by several different routers. A shared services PIC handles flows from all the logical systems. Only version 9 flows, IPv4, and MPLS templates are supported. See "Example: Configuring Active Monitoring on Logical Systems" on page 925 for a sample configuration that enables active monitoring on a logical system.

Specified packets can be filtered and sent to the monitoring interface. For the Monitoring Services PIC, the interface name contains the **mo-** prefix. For the AS or Multiservices PIC, the interface name contains the **sp-** prefix.



NOTE: If you upgrade from the Monitoring Services PIC to the Adaptive Services or Multiservices PIC for active flow monitoring, you must change the name of your monitoring interface from mo-*fpc/pic/port* to sp-*fpc/pic/port*.

The major active flow monitoring actions you can configure at the **[edit forwarding-options]** hierarchy level are as follows:

- Sampling, with the **[edit forwarding-options sampling**] hierarchy. This option sends a copy of the traffic stream to an AS or Monitoring Services PIC, which extracts limited information (such as the source and destination IP address) from some of the packets in a flow. The original packets are forwarded to the intended destination as usual.
- Discard accounting, with the **[edit forwarding-options accounting]** hierarchy. This option quarantines unwanted packets, creates cflowd records that describe the packets, and discards the packets instead of forwarding them.
- Port mirroring, with the **[edit forwarding-options port-mirroring]** hierarchy. This option makes one full copy of all packets in a flow and delivers the copy to a single destination. The original packets are forwarded to the intended destination.
- Multiple port mirroring, with the **[edit forwarding-options next-hop-group]** hierarchy. This option allows multiple copies of selected traffic to be delivered to multiple destinations. (Multiple port mirroring requires a Tunnel Services PIC.)

Unlike passive flow monitoring, you do not need to configure a monitoring group. Instead, you can send filtered packets to a monitoring services or adaptive services interface (**mo-** or **sp-**) by using sampling or discard accounting. Optionally, you can configure port mirroring or multiple port mirroring to direct packets to additional interfaces.

These active flow monitoring options provide a wide variety of actions that can be performed on network traffic flows. However, the following restrictions apply:

- The router can perform sampling *or* port mirroring at any one time.
- The router can perform forwarding *or* discard accounting at any one time.

Because the Monitoring Services, AS, and Multiservices PICs allow only one action to be performed at any one time, the following configuration options are available:

- Sampling and forwarding
- Sampling and discard accounting

Copyright © 2013, Juniper Networks, Inc.

- Port mirroring and forwarding
- Port mirroring and discard accounting
- Sampling and port mirroring on different sets of traffic

Figure 10 on page 906 shows a sample topology.

Figure 10: Active Monitoring Configuration Topology

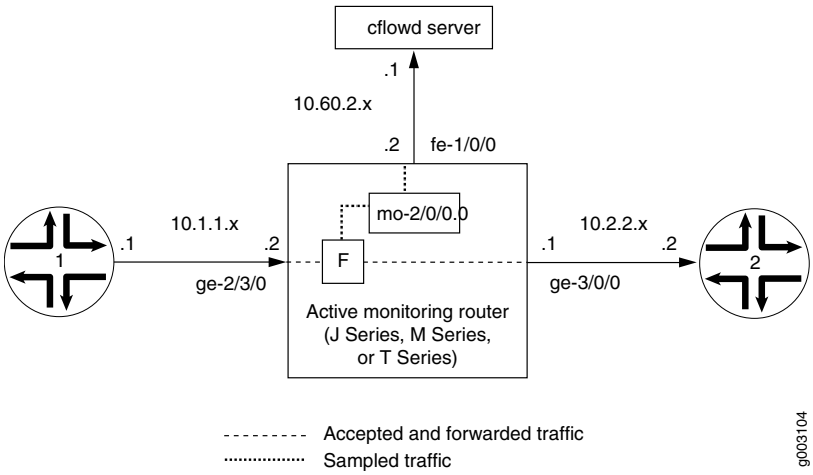

In Figure 10 on page 906, traffic from Router 1 arrives on the monitoring router's Gigabit Ethernet **ge-2/3/0** interface. The exit interface on the monitoring router leading to destination Router 2 is **ge-3/0/0**, but this could be any interface type (such as SONET, Gigabit Ethernet, and so on). The export interface leading to the cflowd server is **fe-1/0/0**.

To enable active monitoring, configure a firewall filter on the interface **ge-2/3/0** with the following match conditions:

- Traffic matching certain firewall conditions is sent to the Monitoring Services PIC using filter-based forwarding. This traffic is quarantined and not forwarded to other routers.
- All other traffic is port-mirrored to the Monitoring Services PIC. Port mirroring copies each packet and sends the copies to the port-mirroring next hop (in this case, a Monitoring Services PIC). The original packets are forwarded out of the router as usual.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 47

Flow Monitoring and Discard Accounting Configuration Guidelines

To configure flow monitoring and accounting interfaces, include the following statements at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
mo-fpc/pic/port {
    unit logical-unit-number {
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage direction;
            }
        }
        address address {
            destination address;
        }
        filter {
            group filter-group-number;
            input filter-name;
            output filter-name;
        }
        receive-options-packets;
        receive-ttl-exceeded;
        sampling direction;
    }
}
multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
    flow-control-options {
        down-on-flow-control;
        dump-on-flow-control;
        reset-on-flow-control;
    }
}
(at-fpc/pic/port | fe-fpc/pic/port | ge-fpc/pic/port) {
    passive-monitor-mode;
}
so-fpc/pic/port {
    unit logical-unit-number {
```

Copyright © 2013, Juniper Networks, Inc.

```
    passive-monitor-mode;
  }
}
```

To configure flow monitoring and accounting properties, include the following statements at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
accounting name {
  output {
    aggregate-export-interval seconds;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface interface-name {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
monitoring name {
  family family {
    output {
      cflowd hostname port port-number;
      export-format format;
      flow-active-timeout seconds;
      flow-export-destination {
        collector-pic;
      }
      flow-inactive-timeout seconds;
      interface interface-name {
        engine-id number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
      }
    }
  }
next-hop-group group-names {
  interface interface-name {
    next-hop address;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  port-mirroring {
    input {
      rate rate;
      run-length number;
      maximum-packet-length bytes
    }
    family (inet | inet6) {
      output {
        interface interface-name {
          next-hop address;
        }
        no-filter-check;
      }
    }
    traceoptions {
      file filename {
        files number;
        size bytes;
        (world-readable | no-world-readable);
      }
    }
  }
  sampling {
    disable;
    sample-once;
    input {
      rate number;
      run-length number;
      max-packets-per-second number;
      maximum-packet-length bytes;
    }
    traceoptions {
      no-remote-trace;
      file filename <files number> <size bytes> <match expression> <world-readable |
        no-world-readable>;
    }
    family (inet | inet6 | mpls) {
      disable;
      output {
        aggregate-export-interval seconds;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        extension-service service-name;
        flow-server hostname {
          aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
              caida-compliant;
            }
            source-prefix;
          }
          autonomous-system-type (origin | peer);
```

Copyright © 2013, Juniper Networks, Inc.

```
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
        version9 {
          template template-name;
        }
      }
      interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
      }
      file {
        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
      }
    }
  }
  instance instance-name {
    disable;
    input {
      rate number;
      run-length number;
      max-packets-per-second number;
      maximum-packet-length bytes;
    }
    family (inet | inet6 | mpls) {
      disable;
      output {
        aggregate-export-interval seconds;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        extension-service service-name;
        flow-server hostname {
          aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
              caida-compliant;
            }
            source-prefix;
          }
          autonomous-system-type (origin | peer);
          (local-dump | no-local-dump);
          port port-number;
          source-address address;
          version format;
          version9 {
            template template-name;
          }
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
        interface interface-name {
          engine-id number;
          engine-type number;
          source-address address;
        }
        inline-jflow {
          source-address address;
          flow-export-rate rate;
        }
      }
    }
  }
}
```



NOTE: For the complete **[edit forwarding-options]** hierarchy, see the Routing Policy Configuration Guide. This section documents only the statements used in flow monitoring and accounting services.

To configure flow monitoring that uses cflowd version 9, include the following statements at the **[edit services]** hierarchy level:

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

This chapter contains the following sections:

- Configuring Traffic Sampling on page 912
- Configuring Flow Monitoring on page 920
- Example: Configuring Active Monitoring on Logical Systems on page 925
- Enabling Flow Aggregation on page 928
- Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 929
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 932

Copyright © 2013, Juniper Networks, Inc.

- Configuring Sampling Instances on page 941
- Configuring Inline Sampling on page 942
- Configuring Inline Sampling on MX80 Routers on page 946
- Directing Replicated Flows to Multiple Flow Servers on page 947
- Logging cflowd Flows Before Export on page 949
- Configuring Port Mirroring on page 950
- Load Balancing Among Multiple Monitoring Interfaces on page 963
- Configuring Discard Accounting on page 967
- Enabling Passive Flow Monitoring on page 968
- Configuring Services Interface Redundancy with Flow Monitoring on page 975

Configuring Traffic Sampling

Traffic sampling enables you to copy traffic to a Physical Interface Card (PIC) that performs flow accounting while the router forwards the packet to its original destination. You can configure the router to perform sampling in either of two locations:

- On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the **then sample** statement.
- On the Monitoring Services, Adaptive Services, or Multiservices PIC.



NOTE: Routing Engine based sampling is not supported on VPN routing and forwarding (VRF) instances.

The following sections provide configuration instructions for traffic sampling:

- Minimum Configuration for Traffic Sampling on page 912
- Configuring Traffic Sampling on page 913
- Disabling Traffic Sampling on page 915
- Sampling Once on page 915
- Configuring Traffic Sampling Output on page 915
- Tracing Traffic Sampling Operations on page 917
- Traffic Sampling Examples on page 917

Minimum Configuration for Traffic Sampling

To configure traffic sampling on a logical interface, you must perform at least the following tasks:

- Create a firewall filter to apply to the logical interfaces being sampled by including the **filter** statement at the **[edit firewall family *family-name*]** hierarchy level. In the filter **then** statement, you must specify the action modifier **sample** and the action **accept**.

  filter *filter-name* {

Copyright © 2013, Juniper Networks, Inc.

```
    term term-name {
      then {
        sample;
        accept;
      }
    }
}
```

For more information about firewall filter actions and action modifiers, see the Routing Policy Configuration Guide.

- Apply the filter to the interfaces on which you want to sample traffic by including the **address** and **filter** statements at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family-name*]** hierarchy level:

```
address address {
  destination destination-address;
}
filter {
  input filter-name;
}
```

- Enable sampling and specify a nonzero sampling rate by including the **sampling** statement at the **[edit forwarding-options]** hierarchy level:

```
sampling {
  input {
    rate number;
    run-length number;
    max-packets-per-second number;
    maximum-packet-length bytes;
  }
}
```

Configuring Traffic Sampling

To configure traffic sampling on any logical interface, include the **input** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
input {
  rate number;
  run-length number;
  max-packets-per-second number;
  maximum-packet-length bytes;
}
```

When you use Routing Engine-based sampling, specify the threshold traffic value by including the **max-packets-per-second** statement. The value is the maximum number of packets to be sampled, beyond which the sampling mechanism begins dropping packets. The range is from 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.

For **maximum-packet-length** on MX-Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these

Copyright © 2013, Juniper Networks, Inc.

devices. For other devices, the range is from 0 to 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.

NOTE: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the **output** statement, the **max-packets-per-second** value is ignored.

Specify the sampling rate by setting the values for **rate** and **run-length** (see Figure 11 on page 914).

Figure 11: Configure Sampling Rate

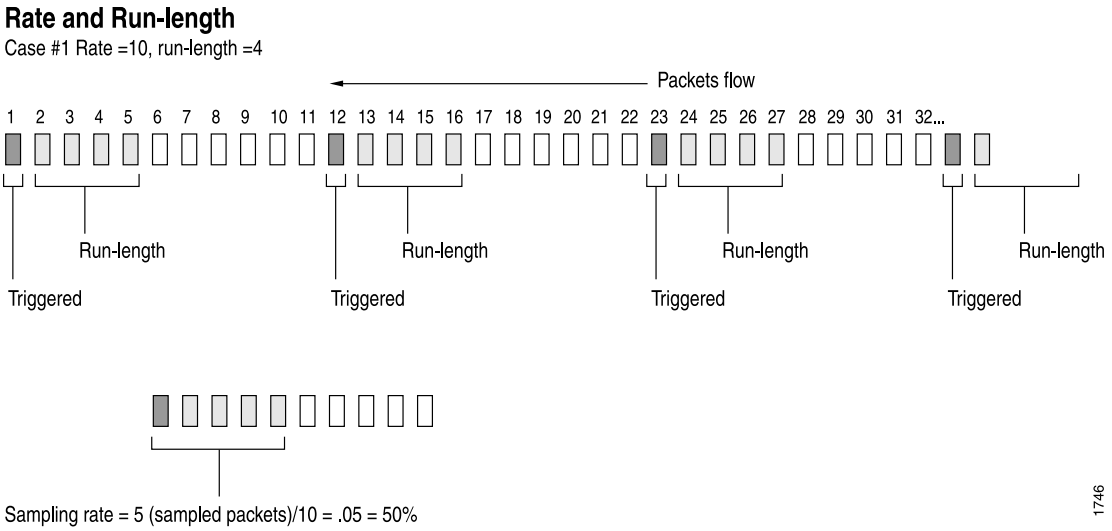

The **rate** statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, x number of packets out of every 10 is sampled, where x=run-length+1. By default, the rate is 0, which means that no traffic is sampled.

The **run-length** statement specifies the number of matching packets to sample following the initial one-packet trigger event. By default, the **run-length** is 0, which means that no more traffic is sampled after the trigger event. The range is from 0 through 20. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.

NOTE: The **run-length** and **maximum-packet-length** configuration statements are not supported on MX80 routers.

If you do not include the **input** statement, sampling is disabled.

To collect the sampled packets in a file, include the **file** statement at the **[edit forwarding-options sampling output]** hierarchy level. Output file formats are discussed later in the chapter.

Copyright © 2013, Juniper Networks, Inc.

Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the **disable** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
disable;
```

Sampling Once

To explicitly sample a packet for active monitoring only once, include the **sample-once** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
sample-once;
```

Setting this option avoids duplication of packets in cases where sampling is enabled at both the ingress and egress interfaces and simplifies analysis of the sampled traffic.

Configuring Traffic Sampling Output

To configure traffic sampling output, include the following statements at the **[edit forwarding-options sampling family *(inet | inet6 | mpls)* output]** hierarchy level:

```
aggregate-export-interval seconds;
flow-active-timeout seconds;
flow-inactive-timeout seconds;
extension-service service-name;
  flow-server hostname {
    aggregation {
      autonomous-system;
      destination-prefix;
      protocol-port;
      source-destination-prefix {
        caida-compliant;
      }
      source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
    version9 {
      template template-name;
    }
  }
  interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
  }
  file {
    disable;
    filename filename;
    files number;
    size bytes;
```

Copyright © 2013, Juniper Networks, Inc.

```
    (stamp | no-stamp);
    (world-readable | no-world-readable);
  }
```

To configure inline flow monitoring on MX Series routers, include the **inline-jflow** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output]** hierarchy level. Inline sampling exclusively supports a new format called IP_FIX that uses UDP as the transport protocol. When you configure inline sampling, you must include the **version-ipfix** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output flow-server *address*]** hierarchy level and also at the **[edit services flow-monitoring]** hierarchy level. For more information about configuring inline flow monitoring, see "Configuring Inline Sampling" on page 942.

To direct sampled traffic to a flow-monitoring interface, include the **interface** statement. The **engine-id** and **engine-type** statements specify the identity and type numbers of the interface; they are dynamically generated based on the Flexible PIC Concentrator (FPC), PIC, and slot numbers and the chassis type. The **source-address** statement specifies the traffic source.

To configure flow sampling version 9 output, you need to include the **template** statement at the **[edit forwarding-options sampling output version9]** hierarchy level. For information on cflowd, see "Enabling Flow Aggregation" on page 928.

The **aggregate-export-interval** statement is described in "Configuring Discard Accounting" on page 967, and the f**low-active-timeout** and **flow-inactive-timeout** statements are described in "Configuring Flow Monitoring" on page 920.

Traffic sampling results are automatically saved to a file in the **/var/tmp** directory. To collect the sampled packets in a file, include the **file** statement at the **[edit forwarding-options sampling family inet output]** hierarchy level:

```
file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}
```

Traffic Sampling Output Format

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the **/var/tmp** directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
# Apr  7 15:48:50
Time                    Dest           Src Dest Src Proto TOS Pkt Intf  IP   TCP

                        addr          addr port port          len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
```


Copyright © 2013, Juniper Networks, Inc.

```
Apr 7 15:48:56 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
```

To set the timestamp option for the file **my-sample**, enter the following:

```
[edit forwarding-options sampling output file]
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the **stamp** option, the **Time** field is displayed.

```
# Apr  7 15:48:50
# Time            Dest        Src  Dest   Src Proto  TOS   Pkt  Intf    IP   TCP
#                 addr       addr  port  port              len   num  frag flags
# Feb  1 20:31:21
#                 Dest        Src  Dest   Src Proto  TOS   Pkt  Intf    IP   TCP
#                 addr       addr  port  port              len   num  frag flags
```

Tracing Traffic Sampling Operations

Tracing operations track all traffic sampling operations and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/sampled**. The default file size is 128K, and 10 files are created before the first one gets overwritten.

To trace traffic sampling operations, include the **traceoptions** statement at the **[edit forwarding-options sampling]** hierarchy level:

```
traceoptions {
  no-remote-trace;
  file filename <files number> <size bytes> <match expression> <world-readable |
    no-world-readable>;
}
```

Traffic Sampling Examples

The following sections provide examples of configuring traffic sampling:

- Example: Sampling a Single SONET/SDH Interface on page 917
- Example: Sampling All Traffic from a Single IP Address on page 918
- Example: Sampling All FTP Traffic on page 919

Example: Sampling a Single SONET/SDH Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET/SDH interface and collects it in a file named **sonet-samples.txt**.

Create the filter:

```
[edit firewall family inet]
filter {
  input sample-sonet {
    then {
      sample;
      accept;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
```

Apply the filter to the SONET/SDH interface:

```
[edit interfaces]
so-0/0/1 {
  unit 0 {
    family inet {
      filter {
        input sample-sonet;
      }
      address 10.127.68.254/32 {
        destination 172.16.74.7;
      }
    }
  }
}
```

Finally, configure traffic sampling:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 100;
      run-length 2;
    }
  }
  family inet {
    output {
      file {
        filename sonet-samples.txt;
        files 40;
        size 5m;
      }
    }
  }
}
```

Example: Sampling All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of **172.16.92.31**, and collects it in a file named **samples-172-16-92-31.txt**.

Create the filter:

```
[edit firewall family inet]
filter one-ip {
  term get-ip {
    from {
      source-address 172.16.92.31;
    }
    then {
      sample;
```


Copyright © 2013, Juniper Networks, Inc.

```
        accept;
      }
    }
  }
```

Apply the filter to the Gigabit Ethernet interface:

```
[edit interfaces]
ge-4/1/1 {
  unit 0 {
    family inet {
      filter {
        input one-ip;
      }
      address 10.45.92.254;
    }
  }
}
```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 1;
    }
  }
  family inet {
    output {
      file {
        filename samples-172-16-92-31.txt;
        files 100;
        size 100k;
      }
    }
  }
}
```

Example: Sampling All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using the FTP data transfer protocol in the output path of a specific T3 interface, and collects the information in a file named **t3-ftp-traffic.txt**.

Create a filter:

```
[edit firewall family inet]
filter ftp-stats {
  term ftp-usage {
    from {
      destination-port [ftp ftp-data];
    }
    then {
      sample;
      accept;
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
```

Apply the filter to the T3 interface:

```
[edit interfaces]
t3-7/0/2 {
  unit 0 {
    family inet {
      filter {
        input ftp-stats;
      }
      address 10.35.78.254/32 {
        destination 10.35.78.4;
      }
    }
  }
}
```

Finally, gather statistics on 10 percent of the candidate samples:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 10;
    }
  }
  family inet {
    output {
      file {
        filename t3-ftp-traffic.txt;
        files 50;
        size 1m;
      }
    }
  }
}
```

Configuring Flow Monitoring

The flow-monitoring application performs traffic flow monitoring and enables lawful interception of traffic between two routers. Traffic flows can either be passively monitored by an offline router or actively monitored by a router participating in the network.

To configure flow monitoring you need to do the following:

- Configuring Flow-Monitoring Interfaces on page 920
- Configuring Flow-Monitoring Properties on page 922
- Example: Configuring Flow Monitoring on page 924

Configuring Flow-Monitoring Interfaces

To enable flow monitoring on the Monitoring Services PIC, include the **mo-*fpc/pic/port*** statement at the **[edit interfaces]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
mo-fpc/pic/port {
  unit logical-unit-number {
    family inet {
      address address {
        destination address;
      }
      filter {
        group filter-group-number;
        input filter-name;
        output filter-name;
      }
      sampling {
        [ input output ];
      }
    }
  }
  multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
    flow-control-options {
      down-on-flow-control;
      dump-on-flow-control;
      reset-on-flow-control;
    }
  }
}
```

Specify the physical and logical location of the flow-monitoring interface. You cannot use **unit 0**, because it is already used by internal processes. Specify the source and destination addresses. The **filter** statement allows you to associate an input or output filter or a filter group that you have already configured for this purpose. The **sampling** statement specifies the traffic direction: **input**, **output**, or both.

The **multiservice-options** statement allows you to configure properties related to flow-monitoring interfaces:

- Include the **core-dump** statement to enable storage of core files in **/var/tmp**.
- Include the **syslog** statement to enable storage of system logging information in **/var/log**.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Boot images for monitoring services interfaces are specified at the **[edit chassis images pic]** hierarchy level. You must include the following configuration to make the flow monitoring feature operable:

```
[edit system]
ntp {
   boot-server ntp.juniper.net;
   server 172.17.28.5;
}
processes {
   ntp enable;
}
```

For more information, see the Junos OS System Basics Configuration Guide.

- Include the **flow-control-options** statement to configure flow control.

Configuring Flow-Monitoring Properties

To configure flow-monitoring properties, include the **monitoring** statement at the **[edit forwarding-options]** hierarchy level:

```
monitoring name {
   family inet {
      output {
         cflowd hostname port port-number;
         export-format format;
         flow-active-timeout seconds;
         flow-export-destination {
            collector-pic;
         }
         flow-inactive-timeout seconds;
         interface interface-name {
            engine-id number;
            engine-type number;
            input-interface-index number;
            output-interface-index number;
            source-address address;
         }
      }
   }
```

A monitoring instance is a named entity that specifies collector information under the **monitoring *name*** statement. The following sections describe the properties you can configure:

- Directing Traffic to Flow-Monitoring Interfaces on page 923
- Exporting Flows on page 923
- Configuring Time Periods when Flow Monitoring is Active and Inactive on page 923

Copyright © 2013, Juniper Networks, Inc.

Directing Traffic to Flow-Monitoring Interfaces

To direct traffic to a flow-monitoring interface, include the **interface** statement at the **[edit forwarding-options monitoring *name* output]** hierarchy level. By default, the Junos OS automatically assigns values for the **engine-id** and **engine-type** statements:

- **engine-id**—Monitoring interface location.
- **engine-type**—Platform-specific monitoring interface type.

The **source-address** statement specifies the traffic source for transmission of cflowd information; you must configure it manually. If you provide a different **source-address** statement for each monitoring services output interface, you can track which interface processes a particular cflowd record.

By default, the **input-interface-index** value is the SNMP index of the input interface. You can override the default by including a specific value. The **input-interface-index** and **output-interface-index** values are exported in fields present in the cflowd version 5 flow format.



NOTE: On J Series Services Routers, cflowd sampling in the input direction of an interface reports the output interface index as 0.

Exporting Flows

To direct traffic to a flow collection interface, include the **flow-export-destination** statement. For more information about flow collection, see Flow Collection.

To configure the cflowd version number, include the **export-format** statement at the **[edit forwarding-options monitoring *name* output]** hierarchy level. By default, version 5 is used. Version 8 enables the router software to aggregate the flow information using broader criteria and reduce cflowd traffic. Version 8 aggregation is performed periodically (every few seconds) on active flows and when flows are allowed to expire. Because the aggregation is performed periodically, active timeout events are ignored.

For more information on cflowd properties, see "Enabling Flow Aggregation" on page 928.

Configuring Time Periods when Flow Monitoring is Active and Inactive

To configure time periods for active flow monitoring and intervals of inactivity, include the **flow-active-timeout** and **flow-inactive-timeout** statements at the **[edit forwarding-options monitoring *name* output]** hierarchy level:

- The **flow-active-timeout** statement specifies the time interval between flow exports for active flows. If the interval between the time the last packet was received and the time the flow was last exported exceeds the configured value, the flow is exported.

  This timer is needed to provide periodic updates when a flow has a long duration. The active timeout setting enables the router to retain the start time for the flow as a constant and send out periodic cflowd reports. This in turn allows the collector to

Copyright © 2013, Juniper Networks, Inc.

register the start time and determine that a flow has survived for a duration longer than the configured active timeout.



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

- The **flow-inactive-timeout** statement specifies the interval of inactivity for a flow that triggers the flow export. If the interval between the current time and the time that the last packet for this flow was received exceeds the configured inactive timeout value, the flow is allowed to expire.

  If the flow stops transmitting for longer than the configured inactive timeout value, the router purges it from the flow table and exports the cflowd record. As a result, the flow is forgotten as far as the PIC is concerned and if the same 5-tuple appears again, it is assigned a new start time and considered a new flow.

Both timers are necessary. The active timeout setting is needed to provide information for flows that constantly transmit packets for a long duration. The inactive timeout setting enables the router to purge flows that have become inactive and would waste tracking resources.



NOTE: The router must contain an Adaptive Services, Multiservices, or Monitoring Services PIC for the **flow-active-timeout** and **flow-inactive-timeout** statements to take effect.

Example: Configuring Flow Monitoring

The following is an example of flow-monitoring properties configured to support input SONET/SDH interfaces, output monitoring services interfaces, and export to cflowd for flow analysis. To complete the configuration, you also need to configure the interfaces and set up a virtual private network (VPN) routing and forwarding (VRF) instance. For a complete example, see the Junos OS Feature Guides. For information on cflowd, see "Enabling Flow Aggregation" on page 928.

```
[edit forwarding-options]
monitoring group1 {
  family inet {
    output {
      cflowd 192.168.245.2 port 2055;
      export-format cflowd-version-5;
      flow-active-timeout 60;
      flow-inactive-timeout 30;
      interface mo-4/0/0.1 {
        engine-id 1;
        engine-type 1;
```


Copyright © 2013, Juniper Networks, Inc.

```
            input-interface-index 44;
            output-interface-index 54;
            source-address 192.168.245.1;
          }
          interface mo-4/1/0.1 {
            engine-id 2;
            engine-type 1;
            input-interface-index 45;
            output-interface-index 55;
            source-address 192.168.245.1;
          }
          interface mo-4/2/0.1 {
            engine-id 3;
            engine-type 1;
            input-interface-index 46;
            output-interface-index 56;
            source-address 192.168.245.1;
          }
          interface mo-4/3/0.1 {
            engine-id 4;
            engine-type 1;
            input-interface-index 47;
            output-interface-index 57;
            source-address 192.168.245.1;
          }
        }
      }
    }
```

Example: Configuring Active Monitoring on Logical Systems

This example shows a sample configuration that allows you to configure active monitoring on a logical system. The following section shows the configuration on the master router:

```
[edit forwarding-options]
sampling {
  instance inst1 {
    input {
      rate 1;
    }
    family inet;
      output {
        flow-server 2.2.2.2 {
          port 2055;
          version9 {
            template {
              ipv4;
            }
          }
        }
      }
      interface sp-0/1/0 {
        source-address 10.11.12.13;
      }
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    family mpls;
      output {
        flow-server 2.2.2.2 {
          port 2055;
          version9 {
            template {
              mpls;
            }
          }
        }
      }
      interface sp-0/1/0 {
        source-address 10.11.12.13;
      }
    }
  }
  services {
    flow-monitoring {
      version9 {
        template ipv4 {
          flow-active-timeout 60;
          flow-inactive-timeout 60;
          ipv4-template;
          template-refresh-rate {
            packets 1000;
            seconds 10;
          }
          option-refresh-rate {
            packets 1000;
            seconds 10;
          }
        }
        template mpls {
          mpls-template;
        }
      }
    }
  }
```

The configuration for the logical router uses the input parameters and the output interface for sampling from the master router. Each logical router should have separate template definitions for the flow-server configuration. The following section shows the configuration on the logical router:

```
  logical-systems {
    ls-1 {
      firewall {
        family inet {
          filter test-sample {
            term term-1 {
              then {
                sample;
                accept;
              }
            }
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        filter {
          input test-sample;
          output test-sample;
        }
      }
    }
  }
}
forwarding-options {
  sampling {
    instance sample-inst1 {
      family inet;
        output {
          flow-server 2.2.2.2 {
            port 2055;
            version9 {
              template {
                ipv4-ls1;
              }
            }
          }
        }
      }
    }
    family mpls;
      output {
        flow-server 2.2.2.2 {
          port 2055;
          version9 {
            template {
              mpls-ls1;
            }
          }
        }
      }
    }
  }
}
services {
  flow-monitoring {
    version9 {
      template ipv4-ls1 {
        flow-active-timeout 60;
        flow-inactive-timeout 60;
        ipv4-template;
        template-refresh-rate {
          packets 1000;
          seconds 10;
        }
```

```
            option-refresh-rate {
              packets 1000;
              seconds 10;
            }
          }
          template mpls-ls1 {
            mpls-template;
          }
        }
      }
    }
  }
}
```

Enabling Flow Aggregation

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs either the cflowd application available from CAIDA (**http://www.caida.org**) or the newer version 9 format defined in RFC 3954, *Cisco Systems NetFlow Services Export Version 9*. Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process. To do this, include the **route-record** statement at the **[edit routing-options]** hierarchy level (for routing instances, include the statement at the **[edit routing-instances *routing-instance-name* routing-options]** hierarchy level):

```
[edit routing-instances routing-instance-name routing-options]
route-record;
```

By default, flow aggregation is disabled.

By using flow aggregation, you can obtain various types of byte and packet counts of flows through a router. The application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

You configure flow aggregation in different ways, depending on whether you want to export flow records in cflowd version 5 or 8 format, or the separate version 9 format. The latter allows you to sample MPLS, IPv4, IPv6, and peer AS billing traffic. You can also combine configuration statements between the MPLS and IPv4 formats. Version 5 Routing Engine-based sampling in the **family inet** statement is allowed when the **family inet6** statement has version 9 PIC-based sampling enabled.



NOTE: When PIC-based sampling is enabled, collection of flow statistics for sampled packets on flows in virtual private networks (VPNs) is also supported. No additional CLI configuration is required.

For configuration instructions for flow aggregation, see the following sections:

- Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 929
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 932

Copyright © 2013, Juniper Networks, Inc.

- Directing Replicated Flows to Multiple Flow Servers on page 947
- Logging cflowd Flows Before Export on page 949

Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd

To enable the collection of cflowd version 5 or version 8 flow formats, include the **flow-server** statement:

```
flow-server hostname {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  (local-dump | no-local-dump);
  port port-number;
  version format;
}
```

You can include this statement at the following hierarchy levels:

- **[edit forwarding-options sampling family (inet | inet6 | mpls) output]**
- **[edit forwarding-options sampling instance *instance-name* output]**
- **[edit forwarding-options accounting *name* output cflowd *hostname*]**

You must configure the **family inet** statement on logical interface **unit 0** on the monitoring interface, as in the following example:

```
[edit interfaces]
sp-3/0/0 {
  unit 0 {
    family inet {
      ...
    }
  }
}
```

Copyright © 2013, Juniper Networks, Inc.



NOTE: Boot images for monitoring services interfaces are specified at the **[edit chassis images pic]** hierarchy level. You must enable the NTP client to make the cflowd feature operable, by including the following configuration:

```
[edit system]
ntp {
   boot-server ntp.juniper.net;
   server 172.17.28.5;
}
processes {
   ntp enable;
}
```

For more information, see the Junos OS System Basics Configuration Guide.

You can also configure cflowd version 5 for flow-monitoring applications by including the **cflowd** statement at the **[edit forwarding-options monitoring *name* family inet output]** hierarchy level:

```
cflowd hostname {
   port port-number;
}
```

The following restrictions apply to cflowd flow formats:

- You can configure up to one version 5 and one version 8 flow format at the **[edit forwarding-options accounting *name* output]** hierarchy level.
- You can configure up to eight version 5 or one version 8 flow format at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output]** hierarchy level for Routing Engine-based sampling by including the **flow-server** statement. In contrast, PIC-based sampling allows you to specify one cflowd version 5 server and one version 8 server simultaneously. However, the two cflowd servers must have different IP addresses.
- You can configure up to eight version 5 flow formats at the **[edit forwarding-options monitoring *name* output]** hierarchy level. Version 8 flow formats and aggregation are not supported for flow-monitoring applications.
- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.
- The configuration includes a proprietary v5 extension template for supporting 4-byte AS information in flow records. Its template version is set to 500, indicating it to be proprietary. All other fields remain the same; the source AS and destination AS are each 4 bytes long, rather than 2 bytes as in the traditional v5 template. This option is

Copyright © 2013, Juniper Networks, Inc.

available at the **[edit forwarding-options sampling family inet output flow-server *server-name* version]** hierarchy level.

In the **cflowd** statement, specify the name or identifier of the host that collects the flow aggregates. You must also include the User Datagram Protocol (UDP) port number on the host and the version, which gives the format of the exported cflowd aggregates. To collect cflowd records in a log file before exporting, include the **local-dump** statement.



NOTE: You can specify both host (cflowd) sampling and port mirroring in the same configuration; however, only one action takes effect at any one time. Port mirroring takes precedence. For more information, see "Configuring Port Mirroring" on page 950.

For cflowd version 8 only, you can specify aggregation of specific types of traffic by including the **aggregation** statement. This conserves memory and bandwidth by enabling cflowd to export targeted flows rather than all aggregated traffic. To specify a flow type, include the **aggregation** statement:

```
aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
        caida-compliant;
    }
    source-prefix;
}
```

You can include this statement at the following hierarchy levels:

- **[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server *hostname*]**
- **[edit forwarding-options accounting *name* output cflowd *hostname*]**

The **autonomous-system** statement configures aggregation by the AS number; this statement might require setting the separate cflowd **autonomous-system-type** statement to include either **origin** or **peer** AS numbers. The **origin** option specifies to use the origin AS of the packet source address in the Source Autonomous System cflowd field. The **peer** option specifies to use the peer AS through which the packet passed in the Source Autonomous System cflowd field. By default, cflowd exports the origin AS number.

The **destination-prefix** statement configures aggregation by the destination prefix only.

The **protocol-port** statement configures aggregation by the protocol and port number; requires setting the separate **cflowd port** statement.

The **source-destination-prefix** statement configures aggregation by the source and destination prefix. Version 2.1b1 of CAIDA's cflowd application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the **caida-compliant** statement, the Junos

Copyright © 2013, Juniper Networks, Inc.

OS complies with Version 2.1b1 of cflowd. If you do not include the **caida-compliant** statement in the configuration, the Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*.

The **source-prefix** statement configures aggregation by the source prefix only.

Collection of sampled packets in a local ASCII file is not affected by the **cflowd** statement.

The following commands enable RE and PIC based sampling at the **set forwarding options sampling** hierarchy level:

- **set input rate** *rate*
- **set input run-length** *length*
- **set family inet output flow-server** *flowcollector* **port** *udp port*
- **set family inet output flow-server** *flowcollector***no-local-dump**
- **set family inet output flow-server** *flowcollector* **version <5/8>**

The following commands enable RE and PIC based sampling at the **set interfaces** hierarchy level:

- *interface to be sampled* unit *unit* family inet filter *input/output filtername*

The following commands enable RE and PIC based sampling at the **set firewall family** hierarchy level:

- **set inet filter** *filtername* **term 1 then count** *filtername***ing**
- **set inet filter** *filtername* **term 1 then sample**
- **set inet filter** *filtername* **term 1 then accept**

The following command enables PIC based sampling at the **set forwarding options sampling** hierarchy level:

- **set family inet output interface** *sp-*/*/** **source address** *source address*

Configuring Flow Aggregation to Use Version 9 Flow Templates

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



NOTE: Version 9 requires that you install a services PIC, such as the Adaptive Services PIC or Multiservices PIC in the router. On MX Series routers, the Multiservices DPC fulfills this requirement. For more information on determining which services PIC is suitable for your router, see "Enabling Service Packages" on page 39 or the appropriate hardware documentation.

Copyright © 2013, Juniper Networks, Inc.

The following sections contain additional information:

- Configuring the Traffic to Be Sampled on page 933
- Configuring the Version 9 Template Properties on page 933
- Restrictions on page 935
- Fields Included in Each Template Type on page 935
- MPLS Sampling Behavior on page 936
- Verification on page 937
- Examples: Configuring Version 9 Flow Templates on page 937

Configuring the Traffic to Be Sampled

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

```
[edit forwarding-options sampling input]
family (inet | inet6 | mpls) {
maximum-packet-length bytes
max-packets-per-second number;
rate number;
run-length number;
```

You can include **family inet** ,**family inet6**, or **family mpls**.



NOTE: If you specify sampling for peer AS billing traffic, the **family** statement supports only IPv4 and IPv6 traffic (inet or inet6). Peer AS billing traffic is enabled only at the global instance hierarchy level and is not available for per Packet Forwarding Engine instances.

Configuring the Version 9 Template Properties

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

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

The following details apply to the configuration statements:

Copyright © 2013, Juniper Networks, Inc.

- You assign each template a unique name by including the **template *name*** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template**, **ipv6–template**, **mpls-ipv4-template**, **mpls-template**, or **peer-as-billing-template**.
- If the template is used for MPLS traffic, you can also specify up to three label positions for the MPLS header label data by including the **label-position** statement; the default values are **[1 2 3]**.
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds. Values you specify in template definitions override the global timeout values configured at the **[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server]** hierarchy level.



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

- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 60 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

```
interfaces interface-name {
   unit 0 {
      family inet6 {
         sampling {
            input;
            output;
         }
      }
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

Restrictions

The following restrictions apply to version 9 templates:

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

  MPLS | Layer 2 Header | IPv4

  In this case, **mpls-ipv4** flows are not created on the PIC, because the IPv4 header does not directly follow the MPLS header. Packets are dropped on the PIC and are accounted as parser errors.
- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.

Fields Included in Each Template Type

The following fields are common to all template types:

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

The IPv4 template includes the following specific fields:

- IPv4 Source Address
- IPv4 Destination Address
- L4 Source Port
- L4 Destination Port
- IPv4 TOS
- IPv4 Protocol
- ICMP type and code

Copyright © 2013, Juniper Networks, Inc.

- TCP Flags
- IPv4 Next Hop Address

The IPv6 template includes the following specific fields:

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

The MPLS template includes the following specific fields:

- MPLS Label #1
- MPLS Label #2
- MPLS Label #3
- MPLS EXP Information
- FEC IP Address

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

The peer AS billing template includes the following specific fields:

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

MPLS Sampling Behavior

This section describes the behavior when MPLS sampling is used on egress interfaces in various scenarios (label pop or swap) on provider routers (P routers). For more information on configuration and background specific to MPLS applications, see the Junos OS MPLS Applications Configuration Guide.

Copyright © 2013, Juniper Networks, Inc.

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

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

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

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

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

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

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

Verification

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

All other **show services accounting** commands also support version 9 templates, except for **show services accounting flow-detail** and **show services accounting aggregation aggregation-type**. For more information about operational mode commands, see the Junos OS Operational Mode Commands.

Examples: Configuring Version 9 Flow Templates

The following is a sample version 9 template configuration:

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

Copyright © 2013, Juniper Networks, Inc.

```
        }
        template peer-as-billing-template-1 {
          peer-as-billing-template;
          }
        }
      }
    }
  }
```

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

```
  firewall {
    family mpls {
      filter mpls_sample {
        term default {
          then {
            accept;
            sample;
          }
        }
      }
    }
  }
```

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

```
  interfaces {
    at-0/1/1 {
      unit 0 {
        family mpls {
          filter {
            input mpls_sample;
          }
        }
      }
    }
    sp-7/0/0 {
      unit 0 {
        family inet;
        family mpls;
      }
    }
  }
```

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

```
  forwarding-options {
    sampling {
      input {
        family mpls {
          rate 1;
        }
      }
      family mpls {
        output {
```


Copyright © 2013, Juniper Networks, Inc.

```
        flow-active-timeout 60;
        flow-inactive-timeout 30;
        flow-server 1.2.3.4 {
          port 2055;
          version9 {
            template mpls-ipv4-template-1;
          }
        }
        interface sp-7/0/0 {
          source-address 1.1.1.1;
        }
        }
      }
    }
  }
```

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

```
  firewall {
    family inet {
      filter peer-as-filter {
        term 0 {
            from {
              destination-class dcu-1;
              interface ge-2/1/0;
              forwarding-class class-1;
            }
            then count count_team_0;
        }
        }
        term 1 {
            from {
              destination-class dcu-2;
              interface ge-2/1/0;
              forwarding-class class-1;
            }
            then count count_team_1;
        }
        term 2 {
            from {
              destination-class dcu-3;
              interface ge-2/1/0;
              forwarding-class class-1;
            }
            then count count_team_2;
        }
      }
    }
  }
```

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

```
  forwarding-options {
    family inet {
```

```
        filter output peer-as-filter;
    }
}
```

The following sample configuration applies the peer AS DCU policy options to collect usage statistics for the traffic stream for as-path ingressing at a specific input interface with the firewall configuration hierarchy applied as Forwarding Table Filters (FTFs). The configuration functionality with COS capability can be achieved through FTFs for destination-class usage with forwarding-class for specific input interfaces:

```
policy-options {
    policy-statement P1 {
        from {
            protocol bgp;
            neighbor 10.2.25.5; #BGP router configuration;
            as-path AS-1; #AS path configuration;
        }
        then destination-class dcu-1; #Destination class configuration;
    }
    policy-statement P2 {
        from {
            neighbor 1.2.25.5;
            as-path AS-2;
        }
        then destination-class dcu2;
    }
    policy-statement P3 {
        from {
            protocol bgp;
            neighbor 192.2.1.1;
            as-path AS-3;
        }
        then destination-class dcu3;
    }
    as-path AS-1 3131:1111:1123;
    as-path AS-2 100000;
    as-path AS-3 192:29283:2;
}
```

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

```
forwarding-options {
    sampling {
        }
        input {
            rate 1;
        }
        family inet {
            output {
                flow-server 10.209.15.58 {
                    port 300;
                    version9 {
                        template {
                            peer-as;
                        }
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
          interface sp-5/2/0 {
            source-address 2.3.4.5;
          }
        }
      }
    }
  }
  family inet {
    filter {
      output peer-as-filter;
    }
  }
```

Configuring Sampling Instances

You can configure active sampling by defining a sampling instance that specifies a name for the sampling parameters and bind the instance name to a Packet Forwarding Engine. This configuration enables you to define multiple named sampling parameter sets associated with multiple destinations (as many as the number of Packet Forwarding Engines in the chassis), with multiple protocol families per sampling destination. With the cflowd version 5 and version 8 and flow aggregation version 9, you can use templates to organize the data gathered from sampling.

To implement this feature, you include the **instance** statement at the [edit forwarding-options sampling] hierarchy level.

The following considerations apply to the sampling instance configuration:

- This configuration is supported on the IP version 4 (**inet**), IP version 6 (**ipv6**), and MPLS protocol families.
- You can configure the **rate** and **run-length** options at the **[edit forwarding-options sampling input]** hierarchy level to apply common values for all families on a global basis. Alternatively, you can configure these options at the **[edit forwarding-options sampling instance *instance-name* input]** hierarchy level to apply specific values for each instance or at the **[edit forwarding-options sampling instance *instance-name* family *family* input]** hierarchy level to apply specific values for each protocol family you configure.
- For MX-Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these devices. For other devices, the range is from 0 to 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.



NOTE: **The run-length and maximum-packet-length configuration statements are not supported on MX80 routers.**

Copyright © 2013, Juniper Networks, Inc.

To associate the defined instance with a particular Packet Forwarding Engine, you include the **sampling-instances** statement at the **[edit chassis fpc *number*]** hierarchy level, as in the following example:

```
chassis {
  fpc 2 {
    sampling-instances samp1;
  }
}
```

Related Documentation

- Flow Monitoring
- More Information About Passive and Active Flow Monitoring
- Example: Sampling Instance Configuration
- Configuring Flow Aggregation (cflowd)
- [edit forwarding-options sampling] Hierarchy Level

Configuring Inline Sampling

On MX Series routers only, you can configure active sampling to be performed on an inline data path without the need for a services Dense Port Concentrator (DPC). To do this, you define a sampling instance with specific properties. One Flexible PIC Concentrator (FPC) can support only one instance; for each instance, either services PIC-based sampling or inline sampling is supported per family. As a result, a particular instance can define PIC-based sampling for one family and inline sampling for a different family. Both IPv4 and IPv6 are supported for inline sampling.

Inline sampling supports a specified sampling output format designated IP_FIX. It uses UDP as the transport protocol.

Before you configure inline sampling, you should ensure that you have adequately sized hash tables for IPv4 and IPv6 flow sampling. These tables can use 1 to 15 256k areas, and each is assigned a default value of one such area. When anticipated traffic volume requires larger tables, allocate larger tables.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The following points are applicable for supporting backward compatibility with Junos OS releases earlier than Release 12.1. when you configure th Ipv4 and Ipv6 flow table sizes for inline sampling:

- If you do not configure the flow table size for inline services by excluding the flow-table-size statement at the **[edit chassis fpc slot-number inline-services]** hierarchy level, fifteen 256K entries are allocated by default for the IPv4 flow table and one 1K entry is allocated by default for the IPv6 flow table on the Packet Forwarding Engine.
- If you configure the **ipv4-flow-table-size** units statement at the **[edit chassis fpc slot-number inline-services flow-table-size]** hierarchy level and if you do not configure the **ipv6-flow-table-size** units statement at the **[edit chassis fpc slot-number inline-services flow-table-size]** hierarchy level, the number of units of 256K entries that you configure for the IPv4 flow table is allocated. For the IPv6 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you do not configure the **ipv4-flow-table-size** units statement at the **[edit chassis fpc slot-number inline-services flow-table-size]** hierarchy level and if you configure the **ipv6-flow-table-size** units statement at the **[edit chassis fpc slot-number inline-services flow-table-size]** hierarchy level, the number of units of 256K entries that you configure for the IPv6 flow table is allocated. For the IPv4 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you configure the sizes of both the IPv4 and IPv6 flow tables, the flow tables are created on the Packet Forwarding Engine based on the size that you specified.

To allocate IPv4 and IPv6 flow hash tables:

1. Go to the flow-table-size hierarchy level for inline services on the FPC that processes the monitored flows.

   ```
   [edit]
   user@host# edit chassis fpc 0 inline-services flow-table-size
   ```

2. Specify the required sizes for the sampling hash tables.

   ```
   [edit chassis fpc 0 inline-services flow-table-size]
   user@host# set ipv4-flow-table-size  5
   user@host# set ipv6-flow-table-size 5
   ```

   NOTE: When you set the flow hash table sizes, remember:

   - Any change in the configured size of flow hash table sizes initiates an automatic reboot of the FPC.
   - The total number of units used for both IPv4 and IPv6 cannot exceed 15.

Copyright © 2013, Juniper Networks, Inc.

The configuration for inline sampling on MX80 routers is slightly different.

To configure inline sampling on all other MX Series routers:

1. Enable inline sampling and specify the source address for the traffic.

   [edit forwarding-options sampling instance *instance-name* family inet output]
   user@host# **set inline-jflow source address** *address*

2. Specify the IP_FIX output format.

   [edit forwarding-options sampling instance *instance-name* family inet output flow-server *address*]
   user@host# **set version-ipfix template ipv4**

3. Specify the output properties.

   [edit services flow-monitoring]
   user@host# **set version-ipfix**

   The output format properties are common to other output formats and are described in "Configuring Flow Aggregation to Use IPFIX Flow Templates".

The following is an example of the sampling configuration for an instance that supports inline sampling on **family inet** and PIC-based sampling on **family inet6**:

```
[edit forwarding-options]
sampling {
   instance {
      sample-ins1 {
         input {
            rate 1;
         }
         family inet {
            output {
               flow-server 2.2.2.2 {
                  port 2055;
                  version-ipfix {
                     template {
                        ipv4;
                     }
                  }
               }
               inline-jflow {
                  source-address 10.11.12.13;
               }
            }
         }
         family inet6 {
            output {
               flow-server 2.2.2.2 {
                  port 2055;
                  version-ipfix {
                     template {
                        ipv6;
                     }
                  }
               }
```


Copyright © 2013, Juniper Networks, Inc.

```
            interface sp-0/1/0 {
              source-address 10.11.12.13;
            }
          }
        }
      }
    }
  }
```

The following example shows the output format configuration:

```
services {
  flow-monitoring {
    version-ipfix {
      template ipv4 {
        flow-active-timeout 60;
        flow-inactive-timeout 60;
        ipv4-template;
        template-refresh-rate {
          packets 1000;
          seconds 10;
        }
        option-refresh-rate {
          packets 1000;
          seconds 10;
        }
      }
    }
  }
}
```

The following considerations apply to the inline flow-monitoring instance configuration:

- Sampling run-length and clip-size are not supported.
- For inline configurations, each family can support only one collector.

NOTE: Inline sampling instances can handle only up to 65536 AS paths. If the total number of AS paths exceed the maximum limit, the AS paths that have AS index greater than 65536 are discarded and counted as error. Flow records associated with such AS paths show the AS value as *0xFFFFFFFF* . However, this limitation does not impact normal forwarding operations.

NOTE: On routers with Multiservices PICs or Multiservices DPCs, all fragments of a fragmented IPv4 packet other than the first fragment of the packet are processed accurately by the flow monitoring application running on MS-PIC or MS-DPC. The flow monitoring mechanism handles such fragments accurately by setting the layer 4 related fields in the associated flows to zero.

Related Documentation

- Configuring Inline Sampling on MX80 Routers on page 946

Copyright © 2013, Juniper Networks, Inc.

- inline-jflow on page 1005

Configuring Inline Sampling on MX80 Routers

To configure inline sampling on MX80 routers:

1. Associate a sampling instance with the Forwarding Engine Processor.

   ```
   [edit]
   user@host# set chassis tfeb slot sampling-instance sampling-instance
   ```

   The Forwarding Engine Processor slot is always **0** because MX80 routers have only one Packet Forwarding Engine. In this configuration, the sampling instance is **sample-ins1**.

   ```
   [edit]
   user@host# set chassis tfeb 0 sampling-instance sample-ins1
   ```

   NOTE: MX80 routers support only one sampling instance.

2. Configure the rate at the **[edit forwarding-options sampling instance instance-name input]** hierarchy level to apply specific values for the sampling instance **sample-ins1**.

   ```
   [edit forwarding-options sampling instance sample-ins1 input]
   user@host# set rate number
   ```

   In this configuration, the rate is **1000**.

   ```
   [edit forwarding-options sampling instance sample-ins1 input]
   user@host# set rate 1000
   ```

3. Enable inline sampling and specify the source address for the traffic.

   ```
   [edit forwarding-options sampling instance sample-ins1 family inet output]
   user@host# set inline-jflow source-address address
   ```

   In this configuration, the source address is **10.11.12.13**.

   ```
   [edit forwarding-options sampling instance sample-ins1 family inet output]
   user@host# set inline-jflow source-address 10.11.12.13
   ```

4. Specify the output properties.

   ```
   [edit services flow-monitoring]
   user@host# set version-ipfix
   ```

   The output format properties are common to other output formats and are described in “Configuring Flow Aggregation to Use IPFIX Flow Templates”.

The following is an example of the sampling configuration for an instance that supports inline sampling on MX80 routers:

```
[edit forwarding-options]
sampling {
  instance {
    sample-ins1 {
```


Copyright © 2013, Juniper Networks, Inc.

```
            input {
              rate 1000;
            }
            family inet {
              output {
                inline-jflow {
                  source-address 10.11.12.13;
                }
              }
            }
          }
        }
      }
```



NOTE: You need not configure a Flexible PIC Concentrator (FPC) slot because MX80 routers have only one Packet Forwarding Engine.

The following considerations apply to the inline flow-monitoring instance configuration:

- This configuration does not support MPLS-IPv6.
- Clip-size is not supported.

Related Documentation

- Configuring Flow Aggregation to Use IPFIX Flow Templates
- Configuring Inline Sampling on page 942
- inline-jflow on page 1005

Directing Replicated Flows to Multiple Flow Servers

You can configure replication of the sampled flow records for use by multiple flow servers. You can use either sampling based on the Routing Engine, using cflowd version 5 or version 8, or sampling based on the services PIC, using flow aggregation version 9, as described in the following sections:

- Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers on page 947
- Directing Replicated Version 9 Flow Aggregates to Multiple Servers on page 948

Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers

Routing Engine–based sampling supports up to eight flow servers for both cflowd version 5 and version 8 configurations. The total number of servers is limited to eight regardless of how many are configured for cflowd v5 or v8.

When you configure cflowd-based sampling, the export packets are replicated to all flow servers configured to receive them. If two servers are configured to receive v5 records, both the servers will receive records for a specified flow.

Copyright © 2013, Juniper Networks, Inc.

NOTE: With Routing Engine–based sampling, if multiple flow servers are configured with version 8 export format, all of them must use the same aggregation type. For example, all servers receiving version 8 export could be configured for **source-destination** aggregation type.

The following configuration example allows replication of export packets to two flow servers.

```
forwarding-options {
  sampling {
    instance inst1 {
      input {
        rate 1;
      }
      family inet;
        output {
          flow-server 10.10.3.2 {
            port 2055;
            version 5;
            source-address 192.168.164.119;
          }
          flow-server 172.17.20.62 {
            port 2055;
            version 5;
            source-address 192.168.164.119;
          }
        }
      }
    }
  }
}
```

Directing Replicated Version 9 Flow Aggregates to Multiple Servers

The export packets generated for a template are replicated to all the flow servers that are configured to receive information for that template. The maximum number of servers supported is eight.

This also implies that periodic updates required by version 9 (RFC 3954) are sent to each configured collector. The following updates are sent periodically as part of this requirement:

- Options data
- Template definition

The refresh period for options data and template definition is configured on a per-template basis at the **[edit services flow-monitoring]** hierarchy level.

The following configuration example allows replication of version 9 export packets to two flow servers.

```
forwarding-options {
```


Copyright © 2013, Juniper Networks, Inc.

```
sampling {
  instance inst1 {
    input {
      rate 1;
    }
    family inet;
      output {
        flow-server 10.10.3.2 {
          port 2055;
          version9 {
            template {
              ipv4;
            }
          }
        }
        flow-server 172.17.20.62 {
          port 2055;
          version9 {
            template {
              ipv4;
            }
          }
        }
      }
      flow-inactive-timeout 30;
      flow-active-timeout 60;
      interface sp-4/0/0 {
        source-address 10.10.3.4;
      }
    }
  }
}
```

Logging cflowd Flows Before Export

To collect the cflowd flows in a log file before they are exported, include the **local-dump** statement at the **[edit forwarding-options sampling output flow-server *hostname*]** hierarchy level:

```
[edit forwarding-options sampling output flow-server hostname]
local-dump;
```

By default, the flows are collected in **/var/log/sampled**; to change the filename, include the **filename** statement at the **[edit forwarding-options sampling traceoptions]** hierarchy level. For more information about changing the filename, see "Configuring Traffic Sampling Output" on page 915.



NOTE: Because the **local-dump** statement adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

Copyright © 2013, Juniper Networks, Inc.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a cflowd header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43    Src addr: 192.53.127.1
Jun 27 18:35:43    Dst addr: 192.6.255.15
Jun 27 18:35:43    Nhop addr: 192.6.255.240
Jun 27 18:35:43    Input interface: 5
Jun 27 18:35:43    Output interface: 3
Jun 27 18:35:43    Pkts in flow: 15
Jun 27 18:35:43    Bytes in flow: 600
Jun 27 18:35:43    Start time of flow: 7230
Jun 27 18:35:43    End time of flow: 7271
Jun 27 18:35:43    Src port: 26629
Jun 27 18:35:43    Dst port: 179
Jun 27 18:35:43    TCP flags: 0x10
Jun 27 18:35:43    IP proto num: 6
Jun 27 18:35:43    TOS: 0xc0
Jun 27 18:35:43    Src AS: 7018
Jun 27 18:35:43    Dst AS: 11111
Jun 27 18:35:43    Src netmask len: 16
Jun 27 18:35:43    Dst netmask len: 0
```

[... 41 more version 5 flow entries; then the following header:]

```
Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43   Num-records: 42
Jun 27 18:35:43   Version: 5
Jun 27 18:35:43   low seq num: 118
Jun 27 18:35:43   Engine id: 0
Jun 27 18:35:43   Engine type: 3
```

Configuring Port Mirroring

On routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

You can configure simultaneous use of sampling and port mirroring, and set an independent sampling rate and run-length for port-mirrored packets. However, if a packet is selected for both sampling and port mirroring, only one action can be performed and port mirroring takes precedence. For example, if you configure an interface to sample every packet input to the interface and a filter also selects the packet to be port mirrored to another interface, only the port mirroring would take effect. All other packets not matching the explicit filter port-mirroring criteria continue to be sampled when forwarded to their final destination.

Copyright © 2013, Juniper Networks, Inc.



NOTE: Configuration for both port mirroring and traffic sampling are handled by the same daemon, so in order to view a trace log file for port mirroring, you must configure the **traceoptions** option under traffic sampling.

To prepare traffic for port mirroring, include the **filter** statement at the **[edit firewall family inet]** hierarchy level:

```
filter filter-name;
```

This filter at the **[edit firewall family (inet | inet6)]** hierarchy level selects traffic to be port-mirrored:

```
filter filter-name {
    term term-name {
        then {
            port-mirror;
            accept;
        }
    }
}
```

To configure port mirroring on a logical interface, configure the following statements at the **[edit forwarding-options port-mirroring]** hierarchy level:

```
[edit forwarding-options port-mirroring family (inet|inet6)]
input {
    maximum-packet-length bytes
    rate rate;
    run-length number;
}
output {
    interface interface-name {
        next-hop address;
    }
    no-filter-check;
}
```



NOTE: The **input** statement can also be configured at the **[edit forwarding-options port-mirroring]** hierarchy level. This is only maintained for backward compatibility. However, the configuration of the **output** statement is deprecated at the **[edit forwarding-options port-mirroring]** hierarchy level.

Specify the port-mirroring destination by including the **next-hop** statement at the **[edit forwarding-options port-mirroring output interface *interface-name*]** hierarchy level:

```
next-hop address;
```



NOTE: For IPv4 port mirroring to reach a next-hop destination, you must manually include a static Address Resolution Protocol (ARP) entry in the router configuration.

Copyright © 2013, Juniper Networks, Inc.

The **no-filter-check** statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it. en

The interface used to send the packets to the analyzer is the output interface configured above at the **[edit forwarding-options port-mirroring family (inet | inet6) output]** hierarchy level. You can use any physical interface type, including generic routing encapsulation (GRE) tunnel interfaces. The next-hop address specifies the destination address; this statement is mandatory for non point-to-point interfaces, such as Ethernet interfaces.

To configure the sampling rate or duration, include the **rate** or **run-length** statement at the **[edit forwarding-options port-mirroring input]** hierarchy level.

You can trace port-mirroring operations the same way you trace sampling operations. For more information, see "Tracing Traffic Sampling Operations" on page 917.

For more information about port mirroring, see the following sections:

- Configuring Tunnels on page 952
- Port Mirroring with Next-Hop Groups on page 953
- Configuring Inline Port Mirroring on page 954
- Filter-Based Forwarding with Multiple Monitoring Interfaces on page 955
- Restrictions on page 955
- Configuring Port Mirroring on Services Interfaces on page 956
- Examples: Configuring Port Mirroring on page 957

Configuring Tunnels

In typical applications, you send the sampled packets to an analyzer or a workstation for analysis, rather than another router. If you must send this traffic over a network, you should use tunnels. For more information about tunnel interfaces, see Tunnel Properties.

If your router is equipped with a Tunnel PIC, you can forward duplicate packets to multiple interfaces by configuring a next-hop group. To configure a next-hop group, include the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level:

```
[edit forwarding-options]
next-hop-group group-names {
    interface interface-name {
        next-hop address;
    }
}
```

The **interface** statement specifies the interface that sends out sampled information. The **next-hop** statement specifies the next-hop addresses to which to send the sampled information.

Next-hop groups have the following restrictions:

- Next-hop groups are supported for IPv4 addresses only.
- Next-hop groups are supported on M Series routers only, except the M120 and the M320.

Copyright © 2013, Juniper Networks, Inc.

- Next-hop groups support up to 16 next-hop addresses.
- Up to 30 next-hop groups are supported.
- Each next-hop group must have at least two next-hop addresses.

Port Mirroring with Next-Hop Groups

You can configure next-hop groups for MX, TX, and T Series routers using either IP addresses or Layer 2 addresses for the next hops. Use the **group-type [inet | layer-2]** statement at **[edit forwarding-options next-hop-group *next-hop-group-name*]** hierarchy level to establish the next-hop groups. You can reference more than one port mirroring instance in a filter on MX Series routers. Use the **port-mirror-instance *instance-name*** statement at the **[edit firewall family *family-name* filter *filter-name* term *term-name*]** to refer to one of several port mirroring instances. For more information about this configuration, see the JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 13.1.



NOTE: On the Trio chipset for MX series routers, port mirroring instances can only be bound to the FPC level and not up to the PIC level. For MX series routers with a DPC card, both levels are supported.

On MX, TX, and T Series routers only, you can configure port mirroring using next-hop groups, also known as *multipacket port mirroring*, without the presence of a Tunnel PIC. To configure this functionality, include the **next-hop-group** statement at the **[edit forwarding-options port-mirror family inet output]** or **[edit forwarding-options port-mirror instance *instance-name* family inet output]** hierarchy level:

```
[edit forwarding-options]
port-mirror {
    family inet {
        output {
            next-hop-group group-name;
        }
    }
}
```

or

```
[edit forwarding-options]
port-mirror {
    instance instance-name {
        family (inet | vpls) {
            output {
                next-hop-group group-name;
            }
        }
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

You define the next-hop group by including the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level. For an example, see "Examples: Configuring Port Mirroring" on page 957.This configuration is supported only with IPv4 addresses.

You can disable this configuration by including a **disable** or **disable-all-instances** statement at the **[edit forwarding-options port-mirror]** hierarchy level or by including a **disable** statement at the **[edit forwarding-options port-mirror instance *instance-name*]** hierarchy level. You can display the settings and network status by issuing the **show forwarding-options next-hop-group** and **show forwarding-options port-mirroring** operational commands.



NOTE: If you try to bind any derived instance to the FPC, a commit error will occur.

Configuring Inline Port Mirroring

Inline port mirroring provides you with the ability to specify instances that are not bound to the flexible PIC concentrator (FPC) in the firewall filter's **then port-mirror-instance** action. This way, you are not limited to only two port-mirror instances per FPC. Inline port mirroring decouples the port-mirror destination from the input parameters like **rate**. While the input parameters are programmed in the switch interface board, the next-hop destination of the mirrored packet is available in the packet itself. Inline port mirroring is supported only on Trio-based modular port concentrators (MPCs).

Using inline port mirroring, a port-mirror instance will have an option to inherit input parameters from another instance that specifies it, as shown in the following CLI configuration example:

```
instance pm2 {
   + input-parameters-instance pm1;
   family inet {
      output {
         interface ge-1/2/3.0 {
            next-hop 50.0.0.3;
         }
      }
   }
}
```

Multiple levels of inheritance are not allowed. One instance can be referred by multiple instances. An instance can refer to another instance that is defined before it. Forward references are not allowed and an instance cannot refer to itself, doing so will cause an error during configuration parsing.

The user can specify an instance that is not bound to the FPC in the firewall filter. The specified filter should inherit one of the two instances that have been bound to the FPC. If it does not, the packet is not marked for port-mirroring. If it does, then the packet will be sampled using the input parameters specified by the referred instance but the copy will be sent to the its own destination.

Copyright © 2013, Juniper Networks, Inc.

Filter-Based Forwarding with Multiple Monitoring Interfaces

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. Obviously, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from those from the previous tables the packet has passed through, to avoid packet looping inside the Packet Forwarding Engine.

For more information about FBF configuration, see the Junos OS Routing Protocols Configuration Guide. For an example of FBF applied to an output interface, see "Examples: Configuring Port Mirroring" on page 957.

Restrictions

The following restrictions apply to port-mirroring configurations:

- The interface you configure for port mirroring should not participate in any kind of routing activity.
- The destination address you specify should not have a route to the ultimate traffic destination. For example, if the sampled IPv4 packets have a destination address of **10.68.9.10** and the port-mirrored traffic is sent to **10.68.20.15** for analysis, the device associated with the latter address should not know a route to **10.68.9.10**. Also, it should not send the sampled packets back to the source address.
- IPv4 and IPv6 traffic is supported. For IPv6 port mirroring, you must configure the next-hop router with an IPv6 neighbor before mirroring the traffic, similar to an ARP request for IPv4 traffic. All the restrictions applied to IPv4 configurations should also apply to IPv6.
- On M120 and M320 routers, multiple next-hop mirroring is not supported.
- Because M320 routers do not support multiple bindings of port-mirror instances per FPC, the **port-mirror-instance** action is not supported in firewall filters for these routers.
- On M Series routers other than the M120 and M320 routers, only one family protocol (either IPv4 or IPv6) is supported at a time.
- Port mirroring supports up to 16 next hops, but there is no next-hop group support for **inet6**.
- Only transit data is supported.
- You can configure multiple port-mirroring interfaces per router.
- On routers containing an Internet Processor II application-specific integrated circuit (ASIC), you must include a firewall filter with both the **accept** action and the **port-mirror** action modifier on the inbound interface. Do not include the **discard** action, or port mirroring will not work.

Copyright © 2013, Juniper Networks, Inc.

- If the port-mirroring interface is a non-point-to-point interface, you must include an IP address under the **port-mirroring** statement to identify the other end of the link. This IP address must be reachable for you to see the sampled traffic. If the port-mirroring interface is an Ethernet interface, the router should have an Address Resolution Protocol (ARP) entry for it. The following sample configuration sets up a static ARP entry.
- You do not need to configure firewall filters on both inbound and outbound interfaces, but at least one is necessary on the inbound interface to provide the copies of the packets to send to an analyzer.
- Inline port mirroring is supported only on Trio-based MPCs.

Configuring Port Mirroring on Services Interfaces

A special situation arises when you configure unit **0** of a services interface (AS or Multiservices PIC) to be the port-mirroring logical interface, as in the following example:

```
[edit forwarding-options]
port-mirroring {
   input {
      rate 1;
   }
   family inet {
      output {
         interface sp-1/0/0.0;
      }
   }
}
```

Since any traffic directed to unit **0** on a services interface is targeted for monitoring (cflowd packets are generated for it), the sample port-mirroring configuration indicates that the customer would like to have cflowd records generated for the port-mirrored traffic.

However, generation of cflowd records requires the following additional configuration; if it is missing, the port-mirrored traffic is simply dropped by the services interface without generating any cflowd packets.

```
[edit forwarding-options]
sampling {
   instance instance1 { # named instances of sampling parameters
      input {
         rate 1;
      }
      family inet {
         output {
            flow-server 172.16.28.65 {
               port 1230;
            }
            interface sp-1/0/0 {  # If the port-mirrored traffic requires monitoring, this
                                  # interface must be same as that specified in the
                                  # port-mirroring configuration.
               source-address 3.1.2.3;
            }
         }
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
```



NOTE: Another way to configure sp-1/0/0 to generate cflowd records is to use only the sampling configuration, but include a firewall filter **sample** action instead of a **port-mirror** action.

Examples: Configuring Port Mirroring

The following example sends port-mirrored traffic to multiple cflowd servers or packet analyzers:

```
[edit interfaces]
ge-1/0/0 { # This is the input interface where packets enter the router.
  unit 0 {
    family inet {
      filter {
        input mirror_pkts; # Here is where you apply the first filter.
      }
      address 10.11.0.1/24;
    }
  }
}
ge-1/1/0 { # This is an exit interface for HTTP packets.
  unit 0 {
    family inet {
      address 10.12.0.1/24;
    }
  }
}
ge-1/2/0 { # This is an exit interface for HTTP packets.
  unit 0 {
    family inet {
      address 10.13.0.1/24;
    }
  }
}
so-0/3/0 { # This is an exit interface for FTP packets.
  unit 0 {
    family inet {
      address 10.1.1.1/30;
    }
  }
}
so-4/3/0 { # This is an exit interface for FTP packets.
  unit 0 {
    family inet {
      address 10.2.2.2/30;
    }
  }
}
so-7/0/0 { # This is an exit interface for all remaining packets.
```

Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family inet {
        address 10.5.5.5/30;
      }
    }
}
so-7/0/1 { # This is an exit interface for all remaining packets.
    unit 0 {
      family inet {
        address 10.6.6.6/30;
      }
    }
}
vt-3/3/0 { # The tunnel interface is where you send the port mirrored traffic.
    unit 0 {
      family inet;
    }
    unit 1 {
      family inet {
        filter {
          input collect_pkts; # This is where you apply the second firewall filter.
        }
      }
    }
}
[edit forwarding-options]
port-mirroring { # This is required when you configure next-hop groups.
    input {
      rate 1; # This rate port mirrors one packet for every one received (1:1 = all
        # packets).
    }
    family inet {
      output { # This sends traffic to a tunnel interface to prepare for multiport mirroring.
        interface vt-3/3/0.1;
        no-filter-check;
      }
    }
}
next-hop-group ftp-traffic { # Point-to-point interfaces require you to specify the interface
     # name only.
    interface so-4/3/0.0;
    interface so-0/3/0.0;
}
next-hop-group http-traffic { # You need to configure a next hop for multipoint interfaces
     # (Ethernet).
    interface ge-1/1/0.0 {
      next-hop 10.12.0.2;
    }
    interface ge-1/2/0.0 {
      next-hop 10.13.0.2;
    }
}
next-hop-group default-collect {
    interface so-7/0/0.0;
    interface so-7/0/1.0;
}
```


Copyright © 2013, Juniper Networks, Inc.

```
[edit firewall]
family inet {
  filter mirror_pkts { # Apply this filter to the input interface.
    term catch_all {
      then {
        count input_mirror_pkts;
        port-mirror; # This action sends traffic to be copied and port mirrored.
        accept;
      }
    }
  }
  filter collect_pkts { # Apply this filter to the tunnel interface.
    term ftp-term { # This term sends FTP traffic to an FTP next-hop group.
      from {
        protocol ftp;
      }
      then next-hop-group ftp-traffic;
    }
    term http-term {# This term sends HTTP traffic to an HTTP next-hop group.
      from {
        protocol http;
      }
      then next-hop-group http-traffic;
    }
    term default {# This term sends all remaining traffic to a final next-hop group.
      then next-hop-group default-collectors;
    }
  }
}
```

The following example demonstrates configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1**, respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25**, and finally leaves the router from interface **so-2/0/0.0**.

```
[edit interfaces]
so-0/0/3 {
  unit 0 {
    family inet {
      filter {
        output fbf;
      }
      address 10.50.10.2/25;
    }
  }
}
fe-1/2/0 {
  unit 0 {
    family inet {
```

```
            address 10.50.50.2/25;
          }
        }
      }
      so-2/0/0 {
        unit 0 {
          family inet {
            address 10.50.20.2/25;
          }
        }
      }
      [edit firewall]
      filter fbf {
        term 0 {
          from {
            source-address {
              10.50.200.0/25;
            }
          }
          then routing-instance fbf;
        }
        term d {
          then count d;
        }
      }
      [edit routing-instances]
      fbf {
        instance-type forwarding;
        routing-options {
          static {
            route 10.50.100.0/25 next-hop so-2/0/0.0;
          }
        }
      }
      [edit routing-options]
      interface-routes {
        rib-group inet fbf-group;
      }
      static {
        route 10.50.100.0/25 next-hop 10.50.10.1;
      }
      rib-groups {
        fbf-group {
          import-rib [ inet.0 fbf.inet.0 ];
        }
      }
```

The following example shows configuration of port mirroring using next-hops groups or multipacket port mirroring:

```
  forwarding-options {
    next-hop-group inet_nhg {
      group-type inet;
      interface ge-2/0/2.101 {
        next-hop 10.2.0.2;
      }
```


Copyright © 2013, Juniper Networks, Inc.

```
    interface ge-2/2/8.2 {
      next-hop 10.8.0.2;
    }
  }
  next-hop-group vpls_nhg {
    group-type layer-2;
    interface ge-2/0/1.100;
    interface ge-2/2/9.0;
    inactive: next-hop-subgroup vpls_subg {
      interface ge-2/0/1.101;
      interface ge-2/2/9.1;
    }
  }
  next-hop-group vpls_nhg_2 {
    group-type layer-2;
    interface ge-2/2/1.100;
    interface ge-2/3/9.0;
  }
  port-mirror {
    disable-all-instances; /* Disable all port-mirroring instances */
    disable; /* Disable the global instance */
    input {
      rate 10; # start mirroring every 10th packet
      run-length 4; # mirror 4 additional packets
    }
    family inet {
      output {
        next-hop-group inet_nhg;
      }
    }
    family vpls {
      output {
        next-hop-group vpls_nhg;
      }
    }
    instance {
      inst1 {
        disable; /* Disable this instance */
        input {
          rate 1;
          maximum-packet-length 200;
        }
        family inet {
          output {
            next-hop-group inet_nhg;
          }
        }
        family vpls {
          output {
            next-hop-group vpls_nhg_2;
          }
        }
      }
    }
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

The following example shows configuration of port mirroring using next-hops groups or multipacket port mirroring on a T series router:

```
forwarding-options {
  next-hop-group inet_nhg {
    group-type inet;
    interface so-0/0/0.0; # There is no need for the nexthop address on T series routers
    interface ge-2/0/2/.0 {
      next-hop 1.2.3.4
    }
    next-hop-subgroup sub_inet {
      interface so-1/2/0.0;
      interface ge-6/1/2.0 {
        next-hop 6.7.8.9;
      }
    }
    next-hop-group vpls_nhg_2 {
      group-type layer-2;
      interface ge-2/2/1.100;
      interface ge-2/3/9.0;
    }
  }
  port-mirroring {
    disable-all-instances; /*Disable all port-mirroring instances */
    disable; /* Disable the global instance */
    input {
      rate 10;
      run-length 4;
    }
    family inet {
      output {
        next-hop-group inet_nhg;
      }
    }
    family vpls {
      output {
        next-hop-group vpls_nhg;
      }
    }
    instance {
      inst1 {
        disable; /* Disable this instance */
        input {
          rate 1;
          maximum-packet-length 200;
        }
        family inet {
          output {
            next-hop-group inet_nhg;
          }
        }
        family vpls {
          output {
            next-hop-group vpls_nhg_2;
          }
        }
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

The following example shows configuration of inline port mirroring using PM1 and PM2 as our port mirror instances.

```
instance {
  pm1 {
    input {
      rate 3;
    }
    family inet {
      output {
        interface ge-1/2/2.0 {
          next-hop 40.0.0.2;
        }
      }
    }
  }
  pm2 {
    input-parameters-instance pm1;
    family inet {
      output {
        interface ge-1/2/3.0 {
          next-hop 50.0.0.3;
        }
      }
    }
  }
  }
  firewall {
    filter pm_filter {
      term t1 {
        then port-mirror-instance pm2;
      }
    }
  }
  chassis {
    fpc 1 {
      port-mirror-instance pm1;
    }
  }
```

The packets will be sampled at a rate of 3 and the copy is sent to 50.0.0.3.

Load Balancing Among Multiple Monitoring Interfaces

The active monitoring application was initially intended for port-mirroring packets on an interface on a normal network router to single or multiple destinations. By port-mirroring these packets to a tunnel interface and using filter-based forwarding on the tunnel interface, port-mirrored packets can be load-balanced across set of interfaces. This method employs existing configuration statements for passive monitoring.

Copyright © 2013, Juniper Networks, Inc.

The configuration consists of the following parts; sample values are included for illustration only.

- Firewall filter configuration—Firewall filter PORT-MIRROR-TO-VT is used to port-mirror the packet to a Tunnel PIC, and filter **catch**, applied on the virtual tunnel (**vt**) interface, is used to send traffic to a filter-based routing instance.

```
[edit firewall]
filter PORT-MIRROR-TO-VT {
  term a {
    then {
      port-mirror;
      accept;
    }
  }
}
filter catch {
  term def {
    then {
      count counter;
      routing-instance fbf_instance;
    }
  }
}
```

For more information about firewall filters, see the Routing Policy Configuration Guide.

- Interface configuration—Apply filter PORT-MIRROR-TO-VT to the interface on which traffic is to be monitored actively.

```
[edit interfaces]
ge-1/3/0 {
  unit 0 {
    family inet {
      filter {
        input PORT-MIRROR-TO-VT;
      }
      address 10.38.0.2/30;
    }
  }
}
vt-3/2/0 {
  unit 0 {
    family inet {
      filter {
        input catch;
      }
    }
  }
}
mo-6/1/0 {
  unit 0 {
    family inet;
  }
}
mo-6/2/0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family inet;
    }
  }
  mo-6/3/0 {
    unit 0 {
      family inet;
    }
  }
  mo-7/1/0 {
    unit 0 {
      family inet;
    }
  }
  mo-7/2/0 {
    unit 0 {
      family inet;
    }
  }
  mo-7/3/0 {
    unit 0 {
      family inet;
    }
  }
```

For more information on configuring interface properties, see the Junos® OS Network Interfaces.

- Routing instance configuration for filter-based forwarding:

```
[edit routing-instances fbf_instance]
instance-type forwarding;
routing-options {
  static {
    route 0.0.0.0/0 next-hop [ mo-7/1/0.0 mo-7/2/0.0 mo-7/3/0.0 mo-6/3/0.0
      mo-6/2/0.0 mo-6/1/0.0 ];
  }
}
```

For more information on routing instance configuration, see the Junos OS Routing Protocols Configuration Guide.

- Routing table groups—Configure the routing table group to resolve the routes installed in the routing instances to directly connected next hops on the interface:

```
[edit routing-options]
interface-routes {
  rib-group inet common;
}
rib-groups {
  common {
    import-rib [ inet.0 fbf_instance.inet.0 ];
  }
}
forwarding-table {
  export pplb;
```

Copyright © 2013, Juniper Networks, Inc.

```
}
```

For more information on routing table groups, see the Junos OS Routing Protocols Configuration Guide.

- Policy for per-packet load balancing:

```
[edit policy-options]
policy-statement pplb {
  then {
    load-balance per-packet;
  }
}
```

For more information on routing policy groups, see the Routing Policy Configuration Guide.

- Port mirroring and monitoring groups—Configure the monitoring services options, and also define hash-based load balancing:

```
[edit forwarding-options]
port-mirroring {
  input {
    rate 1;
  }
  family inet {
    output {
      interface vt-3/2/0.0;
      no-filter-check;
    }
  }
}
monitoring group1 {
  family inet {
    output {
      export-format cflowd-version-5;
      flow-active-timeout 60;
      flow-inactive-timeout 15;
      cflowd 10.36.252.1 port 2055;
      interface mo-6/1/0.0 {
        source-address 10.36.252.2;
      }
      interface mo-6/2/0.0 {
        source-address 10.36.252.2;
      }
      interface mo-6/3/0.0 {
        source-address 10.36.252.2;
      }
      interface mo-7/1/0.0 {
        source-address 10.36.252.2;
      }
      interface mo-7/2/0.0 {
        source-address 10.36.252.2;
      }
      interface mo-7/3/0.0 {
        source-address 10.36.252.2;
      }
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
    hash-key {
      family inet {
        layer-3;
      }
    }
```

For more information on hash keys, see the Routing Policy Configuration Guide.

Configuring Discard Accounting

Discard accounting is similar to traffic sampling, but varies from it in two ways:

- In discard accounting, the packet is intercepted by the monitoring PIC and is not forwarded to its destination.
- Traffic sampling allows you to limit the number of packets sampled by configuring the **max-packets-per-second**, **rate**, and **run-length** statements. Discard accounting does not provide these options, and a high packet count can potentially overwhelm the monitoring PIC.

To configure discard accounting, include the **accounting** statement at the **[edit forwarding-options]** hierarchy level:

```
accounting name {
  output {
    aggregate-export-interval seconds;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface interface-name {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

A discard instance is a named entity that specifies collector information under the **accounting *name*** statement. Discard instances are referenced in firewall filter **term** statements by including the **then discard accounting *name*** statement.

Most of the other statements are also found at the **[edit forwarding-options sampling]** hierarchy level. For information on cflowd, see "Enabling Flow Aggregation" on page 928. The **flow-active-timeout** and **flow-inactive-timeout** statements are described in "Configuring Flow Monitoring" on page 920.

To direct sampled traffic to a flow-monitoring interface, include the **interface** statement. The **engine-id** and **engine-type** statements specify the accounting interface used on the traffic, and the **source-address** statement specifies the traffic source.

You cannot use rate-limiting with discard accounting; however, you can specify the duration of the interval for exporting aggregated accounting information by including the **aggregate-export-interval** statement in the configuration. This enables you to put a boundary on the amount of traffic exported to a flow-monitoring interface.

Enabling Passive Flow Monitoring

You can monitor IPv4 traffic from another router if you have the following components installed in an M Series, MX Series, or T Series router:

- Monitoring Services, Adaptive Services, or Multiservices PICs to perform the service processing
- SONET/SDH, Fast Ethernet, or Gigabit Ethernet PICs as transit interface

On SONET/SDH interfaces, you enable passive flow monitoring by including the **passive-monitor-mode** statement at the **[edit interfaces so-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces so-fpc/pic/port unit logical-unit-number]
passive-monitor-mode;
```

On Asynchronous Transfer Mode (ATM), Fast Ethernet, or Gigabit Ethernet interfaces, you enable passive flow monitoring by including the **passive-monitor-mode** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
passive-monitor-mode;
```

IPv6 passive monitoring is not supported on Monitoring Services PICs. You must configure port mirroring to forward the packets from the passive monitored ports to other interfaces. Interfaces configured on the following FPCs and PIC support IPv6 passive monitoring on the T640 and T1600 routers:

- Enhanced Scaling FPC2
- Enhanced Scaling FPC3
- Enhanced II FPC1
- Enhanced II FPC2

Copyright © 2013, Juniper Networks, Inc.

- Enhanced II FPC3
- Enhanced Scaling FPC4
- Enhanced Scaling FPC4.1
- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (supported on both WAN-PHY and LAN-PHY mode for both IPv4 and IPv6 addresses)
- Gigabit Ethernet PIC with SFP
- 10-Gigabit Ethernet PIC with XENPAK (T1600 router)
- SONET/SDH OC192/STM64 PIC (T1600 router)
- SONET/SDH OC192/STM64 PICs with XFP (T1600 router)
- SONET/SDH OC48c/STM16 PIC with SFP (T1600 router)
- SONET/SDH OC48/STM16 (Multi-Rate)
- SONET/SDH OC12/STM4 (Multi–Rate) PIC with SFP
- Type 1 SONET/SDH OC3/STM1 (Multi–Rate) PIC with SFP

To configure port mirroring, include the **port-mirroring** statement at the **[edit forwarding-options]** hierarchy level.

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you include the **passive-monitor-mode** statement in the configuration:

- The ATM interface is always up, and the interface does not receive or transmit incoming control packets, such as Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) cells.
- The SONET/SDH interface does not send keepalives or alarms and does not participate actively on the network.
- Gigabit and Fast Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.
- Ethernet interfaces do not support the **stacked-vlan-tagging** statement for both IPv4 and IPv6 packets in passive monitoring mode.

On monitoring services interfaces, you enable passive flow monitoring by including the **family** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, specifying the **inet** option:

```
[edit interfaces interface-name unit logical-unit-number]
family inet;
```

Copyright © 2013, Juniper Networks, Inc.

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see "Configuring Flow-Monitoring Interfaces" on page 920.

For conformity with the cflowd record structure, you must include the **receive-options-packets** and **receive-ttl-exceeded** statements at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet]
receive-options-packets;
receive-ttl-exceeded;
```

For more information, see the following sections:

- Passive Flow Monitoring for MPLS Encapsulated Packets on page 970
- Example: Enabling IPv4 Passive Flow Monitoring on page 972
- Example: Enabling IPv6 Passive Flow Monitoring on page 974

Passive Flow Monitoring for MPLS Encapsulated Packets

On monitoring services interfaces, you can process MPLS packets that have not been assigned label values and have no corresponding entry in the **mpls.0** routing table. This allows you to assign a default route to unlabeled MPLS packets.

To configure a default label value for MPLS packets, include the **default-route** statement at the **[edit protocols mpls interface *interface-name* label-map]** hierarchy level:

```
[edit protocols mpls interface interface-name label-map]
default-route {
   (next-hop (address | interface-name | address/interface-name)) | (reject | discard);
   (pop | (swap <out-label>);
   class-of-service value;
   preference preference;
   type type;
}
```

For more information about static labels, see the Junos OS MPLS Applications Configuration Guide.

Removing MPLS Labels from Incoming Packets

The Junos OS can forward only IPv4 packets to a Monitoring Services, Adaptive Services, or Multiservices PIC. IPv4 and IPv6 packets with MPLS labels cannot be forwarded to a monitoring PIC. By default, if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded. To monitor IPv4 and IPv6 packets with MPLS labels, you must remove the MPLS labels as the packets arrive on the interface.

You can remove up to two MPLS labels from an incoming packet by including the **pop-all-labels** statement at the **[edit interfaces *interface-name* (atm-options | fastether-options | gigether-options | sonet-options) mpls]** hierarchy level:

```
[edit interfaces interface-name (atm-options | fastether-options | gigether-options |
   sonet-options) mpls]
pop-all-labels {
   required-depth [ numbers ];
```


Copyright © 2013, Juniper Networks, Inc.

```
}
```

By default, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. You can specify the number of MPLS labels that an incoming packet must have for the **pop-all-labels** statement to take effect by including the **required-depth** statement at the **[edit interfaces *interface-name* (atm-options | fastether-options | gigether-options | sonet-options) mpls pop-all-labels]** hierarchy level:

```
[edit interfaces interface-name (atm-options | fastether-options | gigether-options |
   sonet-options) mpls pop-all-labels]
required-depth [ numbers ];
```

The required depth can be **1**, **2**, or **[1 2]**. If you include the **required-depth 1** statement, the **pop-all-labels** statement takes effect for incoming packets with one label only. If you include the **required-depth 2** statement, the **pop-all-labels** statement takes effect for incoming packets with two labels only. If you include the **required-depth [1 2]** statement, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. A required depth of **[1 2]** is equivalent to the default behavior of the **pop-all-labels** statement.

When you remove MPLS labels from incoming packets, note the following:

- The **pop-all-labels** statement has no effect on IP packets with three or more MPLS labels.
- When you enable MPLS label removal, you must configure all ports on a PIC with the same label popping mode and required depth.
- You use the **pop-all-labels** statement to enable passive monitoring applications, not active monitoring applications.
- You cannot apply MPLS filters or accounting to the MPLS labels because the labels are removed as soon as the packet arrives on the interface.
- On ATM2 interfaces, you must use a label value greater than 4095 because the lower range of MPLS labels is reserved for label-switched interface (LSI) and virtual private LAN service (VPLS) support. For more information, see the Junos OS VPNs Configuration Guide.
- The following ATM encapsulation types are not supported on interfaces with MPLS label removal:
  - **atm-ccc-cell-relay**
  - **atm-ccc-vc-mux**
  - **atm-mlppp-llc**
  - **atm-tcc-snap**
  - **atm-tcc-vc-mux**
  - **ether-over-atm-llc**
  - **ether-vpls-over-atm-llc**

Copyright © 2013, Juniper Networks, Inc.

Example: Enabling IPv4 Passive Flow Monitoring

The following example shows a complete configuration for enabling passive flow monitoring on an Ethernet interface.

In this example, the Gigabit Ethernet interface can accept all Ethernet packets. It strips VLAN tags (if there are any) and up to two MPLS labels blindly, and passes IPv4 packets to the monitoring interface. With this configuration, it can monitor IPv4, VLAN+IPv4, VLAN+MPLS+IPv4, and VLAN+MPLS+MPLS+IPv4 labeled packets.

The Fast Ethernet interface can accept only packets with VLAN ID 100. All other packets are dropped. With this configuration, it can monitor VLAN (ID=100)+IPv4, VLAN (ID=100)+MPLS+IPv4, and VLAN (ID=100)+MPLS+MPLS+IPv4 labeled packets.

```
[edit firewall]
family inet {
   filter input-monitoring-filter {
      term def {
         then {
            count counter;
            accept;
         }
      }
   }
}
[edit interfaces]
ge-0/0/0 {
   passive-monitor-mode;
   gigether-options {
      mpls {
         pop-all-labels;
      }
   }
   unit 0 {
      family inet {
         filter {
            input input-monitoring-filter;
         }
      }
   }
}
fe-0/1/0 {
   passive-monitor-mode;
   vlan-tagging;
   fastether-options {
      mpls {
         pop-all-labels required-depth [ 1 2 ];
      }
   }
   unit 0 {
      vlan-id 100;
      family inet {
         filter {
            input input-monitoring-filter;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
}
mo-1/0/0 {
    unit 0 {
        family inet {
            receive-options-packets;
            receive-ttl-exceeded;
        }
    }
    unit 1 {
        family inet;
    }
}
[edit forwarding-options]
monitoring mon1 {
    family inet {
        output {
            export-format cflowd-version-5;
            cflowd 50.0.0.2 port 2055;
            interface mo-1/0/0.0 {
                source-address 50.0.0.1;
            }
        }
    }
}
[edit routing-instances]
monitoring-vrf {
    instance-type vrf;
    interface ge-0/0/0.0;
    interface fe-0/1/0.0;
    interface mo-1/0/0.1;
    route-distinguisher 68:1;
    vrf-import monitoring-vrf-import;
    vrf-export monitoring-vrf-export;
    routing-options {
        static {
            route 0.0.0.0/0 next-hop mo-1/0/0.1;
        }
    }
}
[edit policy-options]
policy-statement monitoring-vrf-import {
    then {
        reject;
    }
}
policy-statement monitoring-vrf-export {
    then {
        reject;
    }
}
```

Copyright © 2013, Juniper Networks, Inc.

Example: Enabling IPv6 Passive Flow Monitoring

The following example shows a complete configuration for enabling IPv6 passive flow monitoring on an Ethernet interface.

In this example, the Gigabit Ethernet interface can accept all Ethernet packets. It strips VLAN tags (if there are any) and up to two MPLS labels blindly, and passes IPv6 packets to the monitoring interface. With this configuration, the Gigabit Ethernet interface can monitor IPv6, VLAN+IPv6, VLAN+MPLS+IPv6, and VLAN+MPLS+MPLS+IPv6 labeled packets.

The vlan-tagged Gigabit Ethernet interface can accept only packets with VLAN ID 100. All other packets are dropped. With this configuration, it can monitor VLAN (ID=100)+IPv6, VLAN (ID=100)+MPLS+IPv6, and VLAN (ID=100)+MPLS+MPLS+IPv6 labeled packets.

```
[edit interfaces]
xe-0/1/0 {
   passive-monitor-mode;
   unit 0 {
      family inet6 {
         filter {
            input port-mirror6;
         }
         address 2001::1/128;
      }
   }
}
   xe-0/1/2 {
      passive-monitor-mode;
      vlan-tagging;
      unit 0 {
         vlan-id 100;
         family inet6 {
            filter {
               input port-mirror6;
            }
         }
      }
   }
   xe-0/1/1 {
      unit 0 {
         family inet6 {
            address 2000::1/128;
         }
      }
   }
[edit firewall]
family inet6 {
   filter port-mirror6 {
      term term2 {
         then {
            count count_pm;
            port-mirror;
            accept;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
  }
  [edit forwarding options]
  port-mirroring {
    input {
      rate 1;
    }
    family inet6 {
      output {
        interface xe-0/1/1.0 {
          next-hop 2000::3;
        }
        no-filter-check;
      }
    }
  }
```

Configuring Services Interface Redundancy with Flow Monitoring

Active monitoring services configurations on AS, Multiservices PICs, and Multiservices DPCs support redundancy. To configure redundancy, you specify a redundancy services PIC (**rsp**) interface in which the primary AS or Multiservices PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary PIC is restored, it remains on standby and does not preempt the secondary PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the **show interfaces redundancy** command.



NOTE: On flow-monitoring configurations, the only service option supported is *warm standby*, in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected. However, configuration is preserved and available on the new active PIC.

As with the other services that support warm standby, you can issue the **request interfaces (revert | switchover)** command to switch manually between the primary and secondary flow monitoring interfaces.

For more information, see "Configuring AS or Multiservices PIC Redundancy" on page 522. For information on operational mode commands, see the Junos OS Operational Mode Commands.

A sample configuration follows.

```
interface {
  rsp0 {
    redundancy-options {
      primary sp-0/0/0;
      secondary sp-1/3/0;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    unit 0 {
      family inet;
    }
  }
}
interface {
  ge-0/2/0 {
    unit 0 {
      family inet {
        filter {
          input as_sample;
        }
      }
      address 10.58.255.49/28;
    }
  }
}
forwarding-options {
  sampling {
    instance instance1 { # named instances of sampling parameters
      input {
        rate 1;
        run-length 0;
        max-packets-per-second 65535;
      }
      family inet {
        output {
          flow-server 10.10.10.2 {
            port 5000;
            version 5;
          }
          flow-active-timeout 60;
          interface rsp0 {
            source-address 10.10.10.1;
          }
        }
      }
    }
  }
}
firewall {
  filter as_sample {
    term t1 {
      then {
        sample;
        accept;
      }
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 48

Summary of Flow-Monitoring Configuration Statements

The following sections explain each of the flow-monitoring configuration statements. The statements are organized alphabetically.

Copyright © 2013, Juniper Networks, Inc.

accounting

Syntax

```
accounting name {
  output {
    aggregate-export-interval seconds;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface interface-name {
      engine-id (Forwarding Options) number;
      engine-type number;
      source-address (Forwarding Options) address;
    }
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the discard accounting instance name and options.

The statements are explained separately.

Usage Guidelines See "Configuring Discard Accounting" on page 967.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

address (Interfaces)

Syntax
```
address address {
  destination address;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*family *family*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the interface address.

Options ***address***—Address of the interface.

The remaining statement is explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920 or "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Junos® OS Network Interfaces for other options not associated with flow monitoring.

aggregate-export-interval

Syntax aggregate-export-interval *seconds*;

Hierarchy Level [edit forwarding-options accounting *name* output],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output],
[edit forwarding-options sampling family (inet |inet6 |mpls) output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the duration, in seconds, of the interval for exporting aggregate accounting information.

Options ***seconds***—Duration.

Usage Guidelines See "Configuring Discard Accounting" on page 967.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

aggregation

Syntax

```
aggregation {
  autonomous-system;
  destination-prefix;
  protocol-port;
  source-destination-prefix {
    caida-compliant;
  }
  source-prefix;
}
```

Hierarchy Level

[edit forwarding-options accounting output cflowd *hostname*],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For cflowd version 8 only, specify the type of data to be aggregated; cflowd records and sends only those flows that match the specified criteria.

Options

autonomous-system—Aggregate by autonomous system (AS) number.

caida-compliant—Record source and destination mask-length values in compliance with the Version 2.1b1 release of CAIDA's cflowd application. If this statement is not configured, the Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*, dated August 30, 1999.

destination-prefix—Aggregate by destination prefix.

protocol-port—Aggregate by protocol and port number.

source-destination-prefix—Aggregate by source and destination prefix.

source-prefix—Aggregate by source prefix.

Usage Guidelines

See "Enabling Flow Aggregation" on page 928.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

autonomous-system-type

Syntax autonomous-system-type (origin | peer);

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the type of AS numbers that cflowd exports.

Default origin

Options **origin**—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.

peer—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.

Usage Guidelines See "Enabling Flow Aggregation" on page 928.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

cflowd

See the following sections:

- cflowd (Discard Accounting) on page 982
- cflowd (Flow Monitoring) on page 983

cflowd (Discard Accounting)

Syntax

```
cflowd hostname {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  label-position {
    template template-name;
  }
  (local-dump | no-local-dump);
  port port-number;
  source-address (Forwarding Options) address;
  version format;
}
```

Hierarchy Level [edit forwarding-options accounting *name* output],

Release Information Statement introduced before Junos OS Release 7.4.

Description Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.

You can configure up to one version 5 and one version 8 flow format at the **[edit forwarding-options accounting *name* output]** hierarchy level.

Options ***hostname***—The IP address or identifier of the host system (the workstation running the cflowd utility).

The remaining statements are explained separately.

Usage Guidelines See "Enabling Flow Aggregation" on page 928.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

cflowd (Flow Monitoring)

Syntax

```
cflowd hostname {
  port port-number;
}
```

Hierarchy Level [edit forwarding-options monitoring *name* inet output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.

You can configure up to eight version 5 flow formats at the **[edit forwarding-options monitoring *name* output]** hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.

Options ***hostname***—The IP address or identifier of the host system (the workstation running the cflowd utility).

The remaining statement is explained separately.

Usage Guidelines See "Enabling Flow Aggregation" on page 928.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

core-dump

Syntax (core-dump | no-core-dump);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory **/var/tmp** contains core files. The Junos OS saves the current core file (0) and the four previous core files, which are numbered from 1 through 4 (from newest to oldest):



NOTE: By default, all members of a configured user group (with read-only permissions) can access the core dump files and attach them to cases associated with JTAC.

- **core-dump**—Enable the core dumping operation.
- **no-core-dump**—Disable the core dumping operation.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

destination (Tunnel Remote End)

Syntax destination *address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description For tunnel interfaces, specify the remote address of the tunnel.

Options ***destination-address***—Address of the remote side of the connection.

Usage Guidelines See "Configuring Unicast Tunnels" on page 1273, "Configuring Traffic Sampling" on page 912, and "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

disable (Forwarding Options)

Syntax disable;

Hierarchy Level [edit forwarding-options port-mirror],
[edit forwarding-options port-mirror instance *instance-name*],
[edit forwarding-options sampling],
[edit forwarding-options sampling instance *instance-name*],
[edit forwarding-options sampling family (inet |inet6 |mpls)],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file]

Release Information Statement introduced before Junos OS Release 7.4.
Statement added to **port-mirror** hierarchy in Junos OS Release 9.6.

Description Disable traffic accounting, port mirroring, or sampling.

Usage Guidelines See "Configuring Traffic Sampling" on page 912 or "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

disable-all-instances

Syntax disable-all-instances;

Hierarchy Level [edit forwarding-options port-mirror]

Release Information Statement introduced in Junos OS Release 9.6.

Description Disable all port mirroring instances globally.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

engine-id (Forwarding Options)

Syntax engine-id *number*;

Hierarchy Level [edit forwarding-options accounting *name* output interface *interface-name*],
[edit forwarding-options monitoring *name* output interface *interface-name*],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output interface *interface-name*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the engine ID number for flow monitoring and accounting services.

Options ***number***—Identity of accounting interface.

Usage Guidelines See "Configuring Traffic Sampling" on page 912, "Configuring Flow Monitoring" on page 920, or "Configuring Discard Accounting" on page 967.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

engine-type

Syntax engine-type *number*;

Hierarchy Level
[edit forwarding-options accounting *name* output interface *interface-name*],
[edit forwarding-options monitoring *name* output interface *interface-name*],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output interface *interface-name*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the engine type number for flow monitoring and accounting services. The engine type attribute refers to the type of the flow switching engine, such as the route processor or a line module. The configured engine type is inserted in output **cflowd** packets. The **Source ID**, a 32-bit value to ensure uniqueness for all flows exported from a particular device, is the equivalent of the engine type and the engine ID fields.

NOTE: You must configure a source address in the output interface statements. The interface-level statement of engine-type is added automatically but you may override this value with manually configured statements to track different flows with a single **cflowd** collector.

Options *number*—Platform-specific accounting interface type.

Usage Guidelines See "Configuring Traffic Sampling" on page 912, "Configuring Flow Monitoring" on page 920, or "Configuring Discard Accounting" on page 967.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

extension-service

Syntax

```
extension-service service-name {
  provider-specific rules;
}
```

Hierarchy Level

[edit forwarding-options sampling instance *instance-name* family (inet |inet6) output]
[edit forwarding-options sampling family (inet |inet6) output]
[edit services service-set *service-set-name*]

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Define a customer specific sampling configuration.

Define a service set or traffic monitoring for applications using application-specific configuration guidelines.



NOTE: If the extension-service statement is specified while configuring a service set, the service-order statement is mandatory.

Options

provider-specific rules—Provider-specific subhierarchy for services and service sets. See the application-specific documentation for details.

service-name—Name of the service.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- service-order
- sampling on page 1037

Copyright © 2013, Juniper Networks, Inc.

export-format

Syntax export-format *format*;

Hierarchy Level [edit forwarding-options monitoring *name* output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Flow monitoring export format.

Options ***format***—Format of the flows.
Values: 5 or 8
Default: 5

Usage Guidelines See "Exporting Flows" on page 923.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- version on page 1047

Copyright © 2013, Juniper Networks, Inc.

family

See the following sections:

- family (Interfaces) on page 990
- family (Monitoring) on page 991
- family (Port Mirroring) on page 992
- family (Sampling) on page 993

family (Interfaces)

Syntax

```
family family {
  address address {
    destination destination-address;
  }
  filter {
    group filter-group-number;
    input filter-name;
    output filter-name;
  }
  sampling direction;
  receive-options-packets;
  receive-ttl-exceeded;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure protocol family information for the logical interface.

Options ***family***—Protocol family; for flow monitoring and accounting services, only the IP version 4 (IPv4) protocol (**inet**) is supported.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other options not used with services interfaces.

Copyright © 2013, Juniper Networks, Inc.

family (Monitoring)

Syntax

```
family inet {
  output {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    export-format format;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      port port-number;
    }
    interface interface-name {
      engine-id number;
      engine-type number;
      input-interface-index number;
      output-interface-index number;
      source-address address;
    }
  }
}
```

Hierarchy Level [edit forwarding-options monitoring *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify input and output interfaces and properties for flow monitoring. Only IPv4 (**inet**) is supported.

The statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

family (Port Mirroring)

Syntax

```
family (inet | inet6) {
  output {
    interface interface-name {
      next-hop address;
    }
    no-filter-check;
  }
}
```

Hierarchy Level [edit forwarding-options port-mirroring]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the protocol family to be sampled. Only IPv4 (**inet**) and IPv6 (**inet6**) are supported.

The statements are explained separately.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

family (Sampling)

Syntax

```
family (inet | inet6 | mpls) {
  disable;
  output {
    aggregate-export-interval seconds;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    extension-service service-name;
    flow-server hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      (local-dump | no-local-dump);
      port port-number;
      source-address address;
      version format;
      version9 {
        template template-name;
      }
    }
    interface interface-name {
      engine-id number;
      engine-type number;
      source-address address;
    }
    file {
      disable;
      filename filename;
      files number;
      size bytes;
      (stamp | no-stamp);
      (world-readable | no-world-readable);
    }
    inline-jflow {
      source-address address;
      flow-export-rate rate;
    }
  }
}
```

Hierarchy Level

[edit forwarding-options sampling],
[edit forwarding-options sampling instance *instance-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
mpls option introduced in Release 8.3.
inet6 option introduced in Release 9.4.

Copyright © 2013, Juniper Networks, Inc.

Description Configure the protocol family to be sampled. IPv4 (**inet**) is supported for most purposes, but you can configure **family mpls** to collect and export MPLS label information or **family inet6** to collect and export IPv6 traffic using flow aggregation version 9.

The remaining statements are explained separately.



NOTE: **The inline-jflow statement is valid only under the [edit forwarding-options sampling instance *instance-name* family inet output] hierarchy level. The file statement is valid only under the [edit forwarding-options sampling family inet output] hierarchy level.**

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

file

See the following sections:

- file (Sampling) on page 995
- file (Trace Options) on page 995

file (Sampling)

Syntax

```
file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}
```

Hierarchy Level [edit forwarding-options sampling family inet output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Collect the traffic samples in a file.

The statements are explained separately.

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

file (Trace Options)

Syntax file *filename* <files *number* <size *bytes*> <world-readable | no-world-readable>;

Hierarchy Level [edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure information about the files that contain trace logging information.

Options *filename*—The name of the file containing the trace information.

Default: /var/log/sampled

The remaining statements are explained separately.

Usage Guidelines See "Tracing Traffic Sampling Operations" on page 917.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

filename

Syntax filename *filename*;

Hierarchy Level [edit forwarding-options sampling family (inet |inet6 |mpls) output file]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the name of the output file.

Options ***filename***—Name of the file in which to place the traffic samples. All files are placed in the directory **/var/tmp**.

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

files

Syntax files *number*;

Hierarchy Level [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptions file]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the total number of files to be saved with samples or trace data.

Options ***number***—Maximum number of traffic sampling or trace log files. When a file named ***sampling-file*** reaches its maximum size, it is renamed ***sampling-file*.0**, then ***sampling-file*.1**, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.

Range: 1 through 100 files

Default: 5 files for sampling output; 10 files for trace log information

Usage Guidelines See "Configuring Port Mirroring" on page 950 or "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

filter

Syntax

```
filter {
  input filter-name;
  output filter-name;
  group filter-group-number;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply a firewall filter to an interface. You can also use filters for encrypted traffic.

Options

group *filter-group-number*—Define an interface to be part of a filter group. The default filter group number is 0.

input *filter-name*—Name of one filter to evaluate when packets are received on the interface.

output *filter-name*—Name of one filter to evaluate when packets are transmitted on the interface.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Routing Policy Configuration Guide or the Junos OS System Basics Configuration Guide

Copyright © 2013, Juniper Networks, Inc.

flow-active-timeout

Syntax flow-active-timeout *seconds*;

Hierarchy Level
[edit forwarding-options accounting *name* output],
[edit forwarding-options monitoring *name* output],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output],
[edit forwarding-options sampling family (inet |inet6 |mpls) output],
[edit services flow-monitoring version9]

Release Information Statement introduced before Junos OS Release 7.4.

Description Interval after which an active flow is exported.

NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.

Options *seconds*—Duration of the timeout period.

Range: 60 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)

Default: 1800 seconds (for forwarding-options configurations); 60 seconds (for services configurations)

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

Usage Guidelines See "Configuring Time Periods when Flow Monitoring is Active and Inactive" on page 923 or "Configuring the Version 9 Template Properties" on page 933.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-export-rate

Syntax flow-export-rate *rate*;

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family inet output inline-jflow]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the flow export rate of monitored packets in kpps.

Options ***rate***—Flow export rate of monitored packets in kpps (from 1 to 400).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Discard Accounting on page 967
- Configuring Flow Monitoring on page 920
- Configuring Traffic Sampling on page 912

flow-control-options

Syntax

```
flow-control-options {
  down-on-flow-control;
  dump-on-flow-control;
  reset-on-flow-control;
}
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 8.4.

Description Configure the flow control options for application recovery in case of a prolonged flow control failure.

- **down-on-flow-control**—Bring interface down during prolonged flow control.
- **dump-on-flow-control**—Cause core dump during prolonged flow control.
- **reset-on-flow-control**—Reset interface during prolonged flow control.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-export-destination

Syntax

```
flow-export-destination {
    (cflowd-collector | collector-pic);
}
```

Hierarchy Level [edit forwarding-options monitoring *group-name* family inet output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure flow collection.

Options **cflowd-collector**—cflowd collector.

collector-pic—Collector PIC.

Usage Guidelines See "Exporting Flows" on page 923.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-inactive-timeout

Syntax flow-inactive-timeout *seconds*;

Hierarchy Level [edit forwarding-options accounting *name* output],
[edit forwarding-options monitoring *name* output],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output],
[edit forwarding-options sampling family (inet |inet6 |mpls) output],
[edit services flow-monitoring version9]

Release Information Statement introduced before Junos OS Release 7.4.

Description Interval of inactivity that marks a flow inactive.



NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.

Options *seconds*—Duration of the timeout period.

Range: 60 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)

Default: 1800 seconds (for forwarding-options configurations); 60 seconds (for services configurations)

Usage Guidelines See "Configuring Time Periods when Flow Monitoring is Active and Inactive" on page 923 or "Configuring the Version 9 Template Properties" on page 933.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-monitoring

Syntax

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

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the active monitoring properties for flow aggregation version 9.

The statements are explained separately.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-server

Syntax

```
flow-server hostname {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  (local-dump | no-local-dump);
  port port-number;
  source-address address;
  version format;
  version9 {
    template template-name;
  }
}
```

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output], [edit forwarding-options sampling family (inet |inet6 |mpls) output]

Release Information Statement introduced before Junos OS Release 7.4.
version9 statement introduced in Junos OS Release 8.3.

Description Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect. Specify a host system to collect sampled flows using the version 9 format.

You can configure up to one version 5 and one version 8 flow format at the **[edit forwarding-options sampling family (inet | inet6| mpls) output flow-server *hostname*]** hierarchy level. For the same configuration, you can specify only either version 9 flow record formats or formats using versions 5 and 8, not both types of formats.

Options ***hostname***—The IP address–IPv4 or IPv6–or identifier of the host system (the workstationeither running the cflowd utility or collecting traffic flows using version 9).

You can configure only one host system for version 9.



NOTE: IPv6 coniguration for **flow-server** is supported only in Junos OS Release 12.3 and later.

Note that when you configure an IPv6 address for the **flow-server** statement, you must also configure an IPv6 address for the **inline-jflow source-address** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

The remaining statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Traffic Sampling on page 912

forwarding-options

Syntax
forwarding-options { ... }

Hierarchy Level
[edit]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Configure traffic forwarding.

The statements that apply to services interfaces are explained separately. For other statements, see the Routing Policy Configuration Guide.

Usage Guidelines
See "Configuring Flow Monitoring" on page 920 and "Configuring Traffic Sampling" on page 912.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Flow Monitoring on page 920
- Configuring Traffic Sampling on page 912

Copyright © 2013, Juniper Networks, Inc.

inline-jflow

Syntax

```
inline-jflow {
  source-address address;
  flow-export-rate rate;
}
```

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family inet output]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify inline flow monitoring for traffic from the designated address.

Options ***address***—Source IP address.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Inline Sampling on page 942

Copyright © 2013, Juniper Networks, Inc.

input

See the following sections:

- input (Port Mirroring) on page 1006
- input (Sampling) on page 1006

input (Port Mirroring)

Syntax

```
input {
  maximum-packet-lengthbytes
  rate number;
  run-length number;
}
```

Hierarchy Level

[edit forwarding-options port-mirroring],
[edit forwarding-options port-mirroring instance *instance-name*]
[edit forwarding-options port-mirroring family (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure port mirroring on a logical interface.

The statements are explained separately.

Usage Guidelines

See "Configuring Port Mirroring" on page 950.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

input (Sampling)

Syntax

```
input {
  max-packets-per-second number;
  rate number;
  run-length number;
  maximum-packet-length bytes;
}
```

Hierarchy Level

[edit forwarding-options sampling],
[edit forwarding-options sampling instance *instance-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure traffic sampling on a logical interface.

The statements are explained separately.

Usage Guidelines

See "Configuring Traffic Sampling" on page 912.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

input-interface-index

Syntax input-interface-index *number*;

Hierarchy Level [edit forwarding-options monitoring *name* output interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify a value for the input interface index that overrides the default supplied by SNMP.

Options ***number***—Input interface index value.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

instance

See the following sections:

- instance (Port Mirroring) on page 1008
- instance (Sampling) on page 1009

instance (Port Mirroring)

Syntax

```
instance instance-name {
  disable;
  input {
    rate number;
    maximum-packet-length bytes
  }
  family (inet | inet6 | vpls) {
    output {
      next-hop-group group-name;
    }
  }
}
```

Hierarchy Level [edit forwarding-options port-mirroring]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a port-mirroring instance.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Sampling Instances" on page 941.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

instance (Sampling)

Syntax

```
instance instance-name {
  disable;
  input {
    rate number;
    run-length number;
    max-packets-per-second number;
    maximum-packet-length bytes;
  }
  family (inet | inet6 | mpls) {
    disable;
    output {
      aggregate-export-interval seconds;
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      extension-service service-name;
      flow-server hostname {
        aggregation {
          autonomous-system;
          destination-prefix;
          protocol-port;
          source-destination-prefix {
            caida-compliant;
          }
          source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
        version9 {
          template template-name;
        }
      }
      interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
      }
      inline-jflow {
        source-address address;
        flow-export-rate rate;
      }
    }
  }
}
```

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a sampling instance.

Copyright © 2013, Juniper Networks, Inc.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Sampling Instances" on page 941.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

interface

See the following sections:

- interface (Accounting or Sampling) on page 1011
- interface (Monitoring) on page 1012
- interface (Port Mirroring) on page 1012

interface (Accounting or Sampling)

Syntax

```
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
```

Hierarchy Level

[edit forwarding-options accounting *name* output],
[edit forwarding-options sampling family (inet |inet6 |mpls) output],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the output interface for monitored traffic.

Options

interface-name—Name of the interface.

The remaining statements are explained separately.

Usage Guidelines

See "Configuring Discard Accounting" on page 967 or "Configuring Traffic Sampling" on page 912.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

interface (Monitoring)

Syntax

```
interface interface-name {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
}
```

Hierarchy Level [edit forwarding-options monitoring *name* family inet output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the output interface for monitored traffic.

Options ***interface-name***—Name of the interface.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interface (Port Mirroring)

Syntax

```
interface interface-name {
    next-hop address;
}
```

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6) output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options ***interface-name***—Name of the interface.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines See the Junos® OS Network Interfaces for general information.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

ipv4-template

Syntax ipv4-template;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify that the flow aggregation version 9 template is used only for IPv4 records.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

ipv6-template

Syntax ipv6-template;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify that the flow aggregation version 9 template is used only for IPv6 records.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

label-position

Syntax label-position [*positions*];

Hierarchy Level [edit services flow-monitoring version9 template *template-name* mpls-ipv4-template],
[edit services flow-monitoring version9 template *template-name* mpls-template]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify positions for up to three labels in the template.

Default [1 2 3]

Options ***positions***—Numbered positions for the labels.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

local-dump

Syntax (local-dump | no-local-dump);

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable collection of cflowd records in a log file.

Options **no-local-dump**—Do not dump cflowd records to a log file before exporting.

local-dump—Dump cflowd records to a log file before exporting.

Usage Guidelines See "Enabling Flow Aggregation" on page 928.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

match

Syntax match *expression*;

Hierarchy Level [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling traceoptions file]

Release Information Statement introduced before Junos OS Release 7.4.

Description Regular expression for lines to be logged for tracing.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Port Mirroring on page 950
- Configuring Traffic Sampling on page 912

Copyright © 2013, Juniper Networks, Inc.

maximum-packet-length

Syntax maximum-packet-length *bytes*;

Hierarchy Level [edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance *instance-name* input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance *instance-name* input]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.

Description Set the maximum length of the packet used for port mirroring or traffic sampling. Packets with lengths greater than the specified maximum are truncated.



NOTE: **The maximum-packet-length statement is not supported on MX80 routers.**

Options *bytes*—Maximum length (in bytes) of the mirrored packet or the sampled packet.

Range: 0 through 9216

Default: 0

For MX-Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these devices. For other devices, the range is from 0 to 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Port Mirroring
- Configuring Traffic Sampling

Copyright © 2013, Juniper Networks, Inc.

max-packets-per-second

Syntax max-packets-per-second *number*;

Hierarchy Level [edit forwarding-options sampling input],
[edit forwarding-options sampling instance *instance-name* input]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the traffic threshold that must be exceeded before packets are dropped. A value of **0** instructs the Packet Forwarding Engine not to sample any traffic.



NOTE: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the **output** statement, the **max-packets-per-second** value is ignored.

Options *number*—Maximum number of packets per second.

Range: 0 through 65,535

Default: 1000

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

monitoring

Syntax

```
monitoring name {
  family inet {
    output {
      cflowd hostname port-number;
      export-format cflowd-version-5;
      flow-active-timeout seconds;
      flow-export-destination {
        (cflowd-collector | collector-pic);
      }
      flow-inactive-timeout seconds;
      interface interface-name {
        number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
      }
    }
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the flow monitoring instance name and properties.

The statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

mpls-ipv4-template

Syntax

```
mpls-ipv4-template {
  label-position [ positions ];
}
```

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the flow aggregation version 9 properties for templates that combine IPv4 and MPLS records. The remaining statement is explained separately.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

mpls-template

Syntax

```
mpls-template {
  label-position [ positions ];
}
```

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the flow aggregation version 9 properties for templates used only for MPLS records. The remaining statement is explained separately.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

multiservice-options

Syntax

```
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
  flow-control-options {
    down-on-flow-control;
    dump-on-flow-control;
    reset-on-flow-control;
  }
}
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For flow-monitoring interfaces only, configure multiservice-specific interface properties.

The statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

next-hop (Forwarding Options)

Syntax next-hop *address*;

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6) output interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the next-hop address for sending copies of packets to an analyzer.

Options ***address***—IP address of the next-hop router.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

next-hop-group

See the following sections:

- next-hop-group (Forwarding Options) on page 1021
- next-hop-group (Port Mirroring) on page 1022

next-hop-group (Forwarding Options)

Syntax

```
next-hop-group group-name {
  interface interface-name {
    next-hop address;
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the next-hop address for sending copies of packets to an analyzer.

The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options

address—IP address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses.

group-name—Name of next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group must have at least two next-hop addresses.

interface-name—Name of interface used to reach the next-hop destination.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

next-hop-group (Port Mirroring)

Syntax next-hop-group *group-name*;

Hierarchy Level [edit forwarding-options port-mirroring family (inet | vpls) output],
[edit forwarding-options port-mirroring instance *instance-name* family (inet | vpls) output]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the next-hop address for sending copies of packets to an analyzer. This configuration enables multipacket port mirroring on MX Series routers without the use of a Tunnel PIC.

The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options ***group-name***—Name of next-hop group.

Usage Guidelines See "Port Mirroring with Next-Hop Groups" on page 953.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

no-core-dump

See core-dump

no-filter-check

Syntax no-filter-check;

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6) output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Disable filter checking on the port-mirroring interface.

This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

no-local-dump

See local-dump

 Copyright © 2013, Juniper Networks, Inc.

no-remote-trace (Trace Options)

Syntax no-remote-trace;

Hierarchy Level [edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions]

Release Information Statement introduced before Junos OS Release 7.4.

Description Disable remote tracing.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Tracing Traffic Sampling Operations on page 917

no-stamp

See stamp

no-syslog

See syslog

no-world-readable

See world-readable

Copyright © 2013, Juniper Networks, Inc.

option-refresh-rate

Syntax option-refresh-rate packets *packets* seconds *seconds*;

Hierarchy Level [edit services flow-monitoring version9],
[edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the refresh rate, in either packets or seconds.

Options ***packets***—Refresh rate, in number of packets.

Range: 1 through 480,000

Default: 4800

seconds—Refresh rate, in number of seconds.

Range: 10 through 600

Default: 60

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

output

See the following sections:

- output (Accounting) on page 1025
- output (Monitoring) on page 1026
- output (Port Mirroring) on page 1026
- output (Sampling) on page 1027

output (Accounting)

Syntax

```
output {
  aggregate-export-interval seconds;
  cflowd hostname {
    aggregation {
      autonomous-system;
      destination-prefix;
      protocol-port;
      source-destination-prefix {
        caida-compliant;
      }
      source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
  }
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
  }
}
```

Hierarchy Level [edit forwarding-options accounting *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Usage Guidelines See "Configuring Discard Accounting" on page 967.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

output (Monitoring)

Syntax

```
output {
  cflowd hostname port port-number;
  export-format format;
  flow-active-timeout seconds;
  flow-export-destination {
    (cflowd-collector | collector-pic);
  }
  flow-inactive-timeout seconds;
  interface interface-name {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
  }
}
```

Hierarchy Level [edit forwarding-options monitoring *name* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

output (Port Mirroring)

Syntax

```
output {
  interface interface-name {
    next-hop address;
  }
  no-filter-check;
}
```

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure output interfaces and flow properties.

The statements are explained separately.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

output (Sampling)

Syntax

```
output {
  aggregate-export-interval seconds;
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  extension-service service-name;
  flow-server hostname {
    aggregation {
      autonomous-system;
      destination-prefix;
      protocol-port;
      source-destination-prefix {
        caida-compliant;
      }
      source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
    version9 {
      template template-name;
    }
  }
  interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
  }
  file {
    disable;
    filename filename;
    files number;
    size bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
  }
  inline-jflow {
    source-address address;
    flow-export-rate rate;
  }
}
```

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls)],
[edit forwarding-options sampling family (inet |inet6 |mpls)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output files and interfaces, and flow properties.

The statements are explained separately.

Copyright © 2013, Juniper Networks, Inc.



NOTE: The inline-jflow statement is valid only under the [edit forwarding-options sampling instance *instance-name* family inet output] hierarchy level. The file statement is valid only under the [edit forwarding-options sampling family inet output] hierarchy level.

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

output-interface-index

Syntax output-interface-index *number*;

Hierarchy Level [edit forwarding-options monitoring *name* output interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify a value for the output interface index that overrides the default supplied by SNMP.



NOTE: On J Series routers, cflowd sampling in the input direction of an interface reports the output interface index as 0.

Options *number*—Output interface index value.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

passive-monitor-mode

Syntax passive-monitor-mode;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Asynchronous Transfer Mode (ATM), SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, monitor packet flows from another router. If you include this statement in the configuration, the SONET/SDH interface does not send keepalives or alarms, and does not participate actively on the network.

Usage Guidelines See "Enabling Passive Flow Monitoring" on page 968.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- multiservice-options on page 1020

pop-all-labels

Syntax

```
pop-all-labels {
  required-depth number;
}
```

Hierarchy Level

[edit interfaces *interface-name* atm-options mpls],
[edit interfaces *interface-name* fastether-options mpls],
[edit interfaces *interface-name* gigether-options mpls],
[edit interfaces *interface-name* sonet-options mpls]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, removes up to two MPLS labels from incoming IP packets. For passive monitoring on T Series devices, removes up to five MPLS labels from incoming IP packets.

This statement has no effect on IP packets with more than two MPLS labels, or IP packets with more than five MPLS labels on T Series devices. Packets with MPLS labels cannot be processed by the monitoring PIC; if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded.

The remaining statement is explained separately.

Default

If you omit this statement, the MPLS labels are not removed, and the packet is not processed by the monitoring PIC.

Usage Guidelines

See "Passive Flow Monitoring for MPLS Encapsulated Packets" on page 970.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces

Copyright © 2013, Juniper Networks, Inc.

port

Syntax port *port-number*;

Hierarchy Level [edit forwarding-options accounting *name* output cflowd *hostname*],
[edit forwarding-options monitoring *name* family inet output cflowd *hostname*],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the User Datagram Protocol (UDP) port number on the cflowd host system.

Options ***port-number***—Any valid UDP port number on the host system.

Usage Guidelines See "Enabling Flow Aggregation" on page 928.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

port-mirroring

Syntax

```
port-mirroring {
  input {
    maximum-packet-lengthbytes
    rate rate;
    run-length number;
  }
  family inet {
    output {
      interface interface-name {
        next-hop address;
      }
      no-filter-check;
    }
  }
  instance instance-name {
    disable;
    input {
      rate rate;
      maximum-packet-length number;
    }
    family inet {
      output {
        next-hop-group group-name;
      }
    }
  }
  traceoptions {
    file filename <files number> <size bytes> <world-readable | no-world-readable>;
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the input, output, and traceoptions properties for sending copies of packets to an analyzer.

The statements are explained separately.

Usage Guidelines See "Configuring Port Mirroring" on page 950.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rate

Syntax rate *number*;

Hierarchy Level [edit forwarding-options port-mirroring input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance *instance-name* input],
[edit forwarding-options port-mirroring family (inet|inet6) input]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

Options *number*—Denominator of the ratio.

Range: 1 through 65,535

Usage Guidelines See "Configuring Port Mirroring" on page 950 or "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

receive-options-packets

Syntax receive-options-packets;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description When you enable passive monitoring, this statement is required for conformity with cflowd records structure.

Usage Guidelines See "Enabling Passive Flow Monitoring" on page 968.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

receive-ttl-exceeded

Syntax receive-ttl-exceeded;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description When you enable passive monitoring, this statement is required for conformity with cflowd records structure.

Usage Guidelines See "Enabling Passive Flow Monitoring" on page 968.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

required-depth

Syntax required-depth *number*;

Hierarchy Level [edit interfaces *interface-name* atm-options mpls pop-all-labels],
[edit interfaces *interface-name* fastether-options mpls pop-all-labels],
[edit interfaces *interface-name* gigether-options mpls pop-all-labels],
[edit interfaces *interface-name* sonet-options mpls pop-all-labels]

Release Information Statement introduced before Junos OS Release 7.4.

Description For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, specify the number of MPLS labels an incoming packet must have for the **pop-all-labels** statement to take effect.

If you include the **required-depth 1** statement, the **pop-all-labels** statement takes effect for incoming packets with one label only. If you include the **required-depth 2** statement, the **pop-all-labels** statement takes effect for incoming packets with two labels only.

Options *number*—Number of MPLS labels on incoming IP packets.

Range: 1 through 2 labels.

Default: If you omit this statement, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. The default is equivalent to including the **required-depth [1 2]** statement.

Usage Guidelines See "Passive Flow Monitoring for MPLS Encapsulated Packets" on page 970.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces

 Copyright © 2013, Juniper Networks, Inc.

run-length

Syntax run-length *number*;

Hierarchy Level [edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance *port-mirroring-instance-name* input],
[edit forwarding-options port-mirroring family (inet|inet6) input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance *instance-name* input]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.

Description Set the number of samples following the initial trigger event. This allows you to sample packets following those already being sampled.

Options *number*—Number of samples.

Range: 0 through 20

Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Applying Filters to Forwarding Tables
- Configuring Port Mirroring on page 950
- Configuring Traffic Sampling on page 912

sample-once

Syntax sample-once;

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced in Junos OS Release 9.6.

Description Sample traffic for active monitoring only once.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Traffic Sampling on page 912

Copyright © 2013, Juniper Networks, Inc.

sampling

See the following sections:

- sampling (Forwarding Options) on page 1037
- sampling (Interfaces) on page 1039

Copyright © 2013, Juniper Networks, Inc.

sampling (Forwarding Options)

Syntax

```
sampling {
    disable;
    sample-once;
    family (inet | inet6 | mpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            extension-service service-name;
            file {
                disable;
                filename filename;
                files number;
                size bytes;
                (stamp | no-stamp);
                (world-readable | no-world-readable);
            }
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    instance instance-name {
        disable;
        family (inet | inet6 | mpls) {
```

Copyright © 2013, Juniper Networks, Inc.

```
        disable;
        output {
          aggregate-export-interval seconds;
          extension-service service-name;
          flow-active-timeout seconds;
          flow-inactive-timeout seconds;
          flow-server hostname {
            aggregation {
              autonomous-system;
              destination-prefix;
              protocol-port;
              source-destination-prefix {
                caida-compliant;
              }
              source-prefix;
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address address;
            version format;
            version-ipfix {
              template template-name;
            }
            version9 {
              template template-name;
            }
          }
          inline-jflow {
            source-address address;
            flow-export-rate rate;
          }
          interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
          }
        }
      }
      input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
      }
    }
    traceoptions {
      no-remote-trace;
      file filename <files number> <size bytes> <match expression> <world-readable |
        no-world-readable>;
    }
  }
```

Hierarchy Level [edit forwarding-options]

Copyright © 2013, Juniper Networks, Inc.

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure traffic sampling.

The statements are explained separately.

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Applying Filters to Forwarding Tables
- Configuring Active Flow Monitoring Using Version 9
- Configuring Flow Aggregation (cflowd)
- Configuring Port Mirroring
- Tracing Traffic-Sampling Operations

sampling (Interfaces)

Syntax sampling *direction*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the direction of traffic to be sampled.

Options ***direction*** can be one of the following:

input—Configure at least one expected ingress point.

output—Configure at least one expected egress point.

input output—On a single interface, configure at least one expected ingress point and one expect egress point.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos Services Interfaces Configuration Release 12.3
- Configuring Flow Monitoring on page 920

Copyright © 2013, Juniper Networks, Inc.

services

Syntax services { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure router services.

The underlying statements are explained separately.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

size

Syntax size *bytes*;

Hierarchy Level [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptions file]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space.

When a traffic sampling file named **sampling-file** reaches the maximum size, it is renamed **sampling-file.0**. When the **sampling-file** again reaches its maximum size, **sampling-file.0** is renamed **sampling-file.1** and **sampling-file** is renamed **sampling-file.0**. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.

Options ***bytes***—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB for sampling data; 128 KB for log information

Usage Guidelines See "Configuring Port Mirroring" on page 950 or "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

source-address (Forwarding Options)

Syntax source-address *address*;

Hierarchy Level [edit forwarding-options accounting *name* outputinterface *interface-name*],
[edit forwarding-options monitoring *name*familyfamily inet output interface *interface-name*],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output interface *interface-name*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output interface *interface-name*],
[edit forwarding-options sampling instance *instance-name* family inet output inline-jflow]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the source address for monitored packets.

Options ***address***—Interface source address.

Usage Guidelines See "Configuring Discard Accounting" on page 967, "Configuring Flow Monitoring" on page 920, or "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

stamp

Syntax (stamp | no-stamp);

Hierarchy Level [edit forwarding-options sampling family (inet |inet6 |mpls) output file]

Release Information Statement introduced before Junos OS Release 7.4.

Description Include a timestamp with each line in the output file.

Options **no-stamp**—Do not include timestamps. This is the default.

stamp—Include a timestamp with each line of packet sampling information.
Default: No timestamp is included.

Usage Guidelines See "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

syslog

Syntax (syslog | no-syslog);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the **/var/log** directory.

- **syslog**—Enable PIC system logging.
- **no-syslog**—Disable PIC system logging.

Usage Guidelines See "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

template

See the following sections:

- template (Forwarding Options) on page 1043
- template (Services) on page 1044

template (Forwarding Options)

Syntax template *template-name*;

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname* version9],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname* version9]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify flow aggregation version 9 template to be used for output of sampling records.

Options ***template-name***—Name of version 9 template.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

template (Services)

Syntax

```
template template-name {
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  ipv4-template;
  ipv6-template;
  mpls-template {
    label-position [ positions ];
  }
  mpls-ipv4-template {
    label-position [ positions ];
  }
  peer-as-billing-template;
  option-refresh-rate packets packets seconds seconds;
  template-refresh-rate packets packets seconds seconds;
}
```

Hierarchy Level [edit services flow-monitoring version9]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the flow aggregation version 9 template properties. The remaining statements are explained separately.

Options *template-name*—Name of the version 9 template.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

template-refresh-rate

Syntax template-refresh-rate packets *packets* seconds *seconds*;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the refresh rate, in either packets or seconds.

Options *packets*—Refresh rate, in number of packets.
Range: 1 through 480,000
Default: 4800

seconds—Refresh rate, in number of seconds.
Range: 10 through 600
Default: 60

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

traceoptions (Forwarding Options)

Syntax
```
traceoptions {
  no-remote-trace;
  file filename <files number> <size bytes> <match expression> <world-readable |
    no-world-readable>;
}
```

Hierarchy Level [edit forwarding-options port-mirroring],
[edit forwarding-options sampling]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure traffic sampling tracing operations.

The statements are explained separately.

Usage Guidelines See "Tracing Traffic Sampling Operations" on page 917.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

unit

Syntax

```
unit logical-unit-number {
  family inet {
    address address {
      destination destination-address;
    }
    filter {
      group filter-group-number;
      input filter-name;
      output filter-name;
    }
    sampling direction;
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number***—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Usage Guidelines For general information, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do no affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

version

Syntax version *format*;

Hierarchy Level [edit forwarding-options accounting *name* output flow-server *hostname*],
[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the version format of the aggregated flows exported to a cflowd server.

Options *format*—Format of the flows.

Values: 5 or 8

Default: 5

Usage Guidelines See "Enabling Flow Aggregation" on page 928.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- export-format on page 989

Copyright © 2013, Juniper Networks, Inc.

version9

See the following sections:

- version9 (Forwarding Options) on page 1048
- version9 (Services) on page 1049

version9 (Forwarding Options)

Syntax

```
version9 {
  template template-name;
}
```

Hierarchy Level

[edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output flow-server *hostname*],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server *hostname*]

Release Information

Statement introduced in Junos OS Release 8.3.

Description

Specify flow aggregation version 9 properties to apply to output sampling records. The remaining statements are explained separately.

Usage Guidelines

See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

version9 (Services)

Syntax

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

Hierarchy Level [edit services flow-monitoring]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify flow aggregation version 9 template properties. The remaining statements are explained separately.

Usage Guidelines See "Configuring Flow Aggregation to Use Version 9 Flow Templates" on page 932.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

version-ipfix

See the following sections:

- version-ipfix (Forwarding Options) on page 1050
- version-ipfix (Services) on page 1051

version-ipfix (Forwarding Options)

Syntax

```
version-ipfix {
  template template-name;
}
```

Hierarchy Level [edit forwarding-options sampling instance *instance-name* family inet output flow-server *address*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the output format to support inline flow monitoring.

Options ***template-name***—Currently **ipv4** is the only output template format supported.

Usage Guidelines See "Configuring Inline Sampling" on page 942.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- export-format on page 989

Copyright © 2013, Juniper Networks, Inc.

version-ipfix (Services)

Syntax

```
version-ipfix {
  template template-name {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    ipv4-template;
    option-refresh-rate packets packets seconds seconds;
    template-refresh-rate packets packets seconds seconds;
  }
}
```

Hierarchy Level [edit services flow-monitoring]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the output template properties to support inline flow monitoring. The remaining statements are explained separately.

Usage Guidelines See "Configuring Inline Sampling" on page 942.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

world-readable

Syntax (world-readable | no-world-readable);

Hierarchy Level [edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptionsfile]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable unrestricted file access.

Options **no-world-readable**—Restrict file access to owner. This is the default.

world-readable—Enable unrestricted file access.

Default: no-world-readable

Usage Guidelines See "Configuring Port Mirroring" on page 950 or "Configuring Traffic Sampling" on page 912.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 49

Flow Collection Configuration Guidelines

You can process and export multiple cflowd records with a flow collector interface. You create a flow collector interface on a Monitoring Services II or Multiservices 400 PIC. The flow collector interface combines multiple cflowd records into a compressed ASCII data file and exports the file to an FTP server. To convert a services PIC into a flow collector interface, include the **flow-collector** statement at the **[edit chassis fpc** *fpc-slot* **pic** *pic-slot* **monitoring-services application]** hierarchy level.

You can use the services PIC for either flow collection or monitoring, but not for both types of service simultaneously. When converting the PIC between service types, you must configure the **flow-collector** statement, take the PIC offline, and then bring the PIC back online. Restarting the router does not enable the new service type.

A flow collector interface, designated by the **cp-***fpc/pic/port* interface name, requires three logical interfaces for correct operation. Units 0 and 1 are used to send the compressed ASCII data files to an FTP server, while Unit 2 is used to receive cflowd records from a monitoring services interface.



NOTE: Unlike conventional interfaces, the **address** statement at the **[edit interfaces cp-***fpc/pic/port* **unit** *unit-number* **family inet]** hierarchy level corresponds to the IP address of the Routing Engine. Likewise, the **destination** statement at the **[edit interfaces cp-***fpc/pic/port* **unit** *unit-number* **family inet address** *ip-address***]** hierarchy level corresponds to the IP address of the flow collector interface. As a result, you must configure the **destination** statement for Unit 0 and 1 with *local* addresses that can reach the FTP server. Similarly, configure the **destination** statement for Unit 2 with a *local* IP address so it can reach the monitoring services interface that sends cflowd records.

To activate flow collector services after the services PIC is converted into a flow collector, include the **flow-collector** statement at the **[edit services]** hierarchy level. After you activate the flow collector, you need to configure the following components:

- Destination of the FTP server
- File specifications
- Input interface-to-flow collector interface mappings
- Transfer log settings

Copyright © 2013, Juniper Networks, Inc.

To configure flow collection, include the **flow-collector** statement at the **[edit services]** hierarchy level:

```
flow-collector {
    analyzer-address address;
    analyzer-id name;
    destinations {
        ftp:url {
            password "password";
        }
        file-specification {
            variant variant-number {
                data-format format;
                name-format format;
                transfer {
                    record-level number;
                    timeout seconds;
                }
            }
        }
        interface-map {
            collector interface-name;
            file-specification variant-number;
            interface-name {
                collector interface-name;
                file-specification variant-number;
            }
        }
        retry number;
        retry-delay seconds;
        transfer-log-archive {
            archive-sites {
                ftp:url {
                    password "password";
                    username username;
                }
            }
            filename-prefix prefix;
            maximum-age minutes;
        }
    }
}
```

This chapter contains the following sections:

- Configuring Flow Collection on page 1055
- Sending cflowd Records to Flow Collector Interfaces on page 1058
- Configuring Flow Collection Mode and Interfaces on Services PICs on page 1058
- Example: Configuring Flow Collection on page 1058

Copyright © 2013, Juniper Networks, Inc.

Configuring Flow Collection

This section describes the following tasks for configuring flow collection:

- Configuring Destination FTP Servers for Flow Records on page 1055
- Configuring a Packet Analyzer on page 1055
- Configuring File Formats on page 1056
- Configuring Interface Mappings on page 1056
- Configuring Transfer Logs on page 1057
- Configuring Retry Attempts on page 1057

Configuring Destination FTP Servers for Flow Records

Flow collection destinations are where the compressed ASCII data files are sent after the cflowd records are collected and processed. To specify the destination FTP server, include the **destinations** statement at the **[edit services flow-collector]** hierarchy level. You can specify up to two FTP server destinations and include the password for each configured server. If two FTP servers are configured, the first server in the configuration is the primary server and the second is a backup server.

To configure a destination for flow collection files, include the **destinations** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
destinations {
  ftp:url {
    password "password";
  }
}
```

To specify the destination FTP server, include the **ftp:*url*** statement. The value ***url*** is the FTP server address for the primary flow collection destination and can include macros.

When you include macros in the **ftp:*url*** statement, a directory can be created only for a single level. For example, the path **ftp://10.2.2.2/%m/%Y** expands to **ftp://10.2.2.2/01/2005**, and the software attempts to create the directory **01/2005** on the destination FTP server. If the **01/** directory already exists on the destination FTP server, the software creates the **/2005/ directory** one level down. If the **01/** directory does not exist on the destination FTP server, the software cannot create the **/2005/ directory**, and the FTP server destination will fail. For more information about macros, see **ftp**.

To specify the FTP server password, include the **password "*password*"** statement. The password must be enclosed in quotation marks. You can specify up to two destination FTP servers. The first destination specified is considered the primary destination.

Configuring a Packet Analyzer

You can specify values for the IP address and identifier of a packet analyzer to which the flow collector interface sends traffic for analysis. The values you specify here override any default values configured elsewhere.

Copyright © 2013, Juniper Networks, Inc.

To configure an IP address and identifier for the packet analyzer, include the **analyzer-address** and **analyzer-id** statements at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
analyzer-address address;
analyzer-id name;
```

Configuring File Formats

You configure data file formats, name formats, and transfer characteristics for the flow collection files. File records are sent to the destination FTP server when the timer expires or when a preset number of records are received, whichever comes first.

To configure the flow collection file format, include the **file-specification** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
file-specification {
  variant variant-number {
    data-format format;
    name-format format;
    transfer {
      record-level number;
      timeout seconds;
    }
  }
}
```

To set the data file format, include the **data-format** statement. To set the file name format, include the **name-format** statement. To set the export timer and file size thresholds, include the **transfer** statement and specify values for the **timeout** and **record-level** options.

For example, you can specify the name format as follows:

```
[edit services flow-collector file-specification variant variant-number]
name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
```

In this example, **cFlowd-py69Ni69-0** is the static portion used verbatim, **%D** is the date in YYYYMMDD format, **%T** is the time in HHMMSS format, **%I** is the value of **ifAlias**, **%N** is the generation number, and **bcp.bi.gz** is a user-configured string. A number of macros are supported for expressing the date and time information in different ways; for a complete list, see the summary section for name-format.

Configuring Interface Mappings

You can match an input interface with a flow collector interface and apply the preset file specifications to the input interface.

To configure an interface mapping, include the **interface-map** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
interface-map {
```


Copyright © 2013, Juniper Networks, Inc.

```
    collector interface-name;
    file-specification variant-number;
    interface-name {
      collector interface-name;
      file-specification variant-number;
    }
  }
```

To configure the default flow collector and file specifications for all input interfaces, include the **file-specification** and **collector** statements at the **[edit services flow-collector interface-map]** hierarchy level. To override the default settings and apply flow collector and file specifications to a specific input interface, include the **file-specification** and **collector** statements at the **[edit services flow-collector interface-map *interface-name*]** hierarchy level.

Configuring Transfer Logs

You can configure the filename, export interval, maximum size, and destination FTP server for log files containing the transfer activity history for a flow collector interface.

To configure a transfer log, include the **transfer-log-archive** statement at the **[edit services flow-collector]** hierarchy level:

```
[edit services flow-collector]
transfer-log-archive {
  archive-sites {
    ftp:url {
      password "password";
      username username;
    }
  }
  filename-prefix prefix;
  maximum-age minutes;
}
```

To configure the destination for archiving files, include the **archive-sites** statement. Specify the filename as follows:

```
[edit services flow-collector transfer-log]
filename "cFlowd-py69Ni69-0-%D_%T";
```

where **cFlowd-py69Ni69-0** is the static portion used verbatim, **%D** is the date in YYYYMMDD format, and **%T** is the time in HHMMSS format.

You can optionally include the following statements:

- **filename-prefix**—Sets a standard prefix for all the logged files.
- **maximum-age**—Specifies the duration a file remains on the server. The range is 1 through 360 minutes.

Configuring Retry Attempts

You can specify values for situations in which the flow collector interface needs more than one attempt to transfer log files to the FTP server:

Copyright © 2013, Juniper Networks, Inc.

- Maximum number of retry attempts
- Amount of time the flow collector interface waits between successive retries

To configure retry settings, include the **retry** and **retry-delay** statements at the **[edit services flow-collector]** hierarchy level:

```
retry number;
retry-delay seconds;
```

The **retry** value can be from 0 through 10. The **retry-delay** value can be from 0 through 60 seconds.

Sending cflowd Records to Flow Collector Interfaces

To specify a flow collector interface as the destination for cflowd records coming from a services PIC, include the **collector-pic** statement at the **[edit forwarding-options monitoring *group-name* family inet output flow-export-destination]** hierarchy level:

```
[edit forwarding-options monitoring group-name family inet output flow-export-destination]
collector-pic;
```

You can select either the flow collector interface or a cflowd server as the destination for cflowd records, but not both at the same time.

Configuring Flow Collection Mode and Interfaces on Services PICs

You can select the services PIC to run in either flow collection mode or monitoring mode, but not both.

To set the services PIC to run in flow collection mode, include the **flow-collector** statement at the **[edit chassis fpc *slot-number* pic *pic-number* monitoring-services application]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number monitoring-services application]
flow-collector;
```

For further information on configuring chassis properties, see the Junos OS System Basics Configuration Guide.

To specify flow collection interfaces, you configure the **cp** interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
cp-fpc/pic/port {
    ...
}
```

Example: Configuring Flow Collection

Figure 12 on page 1059 shows the path traveled by monitored traffic as it passes through the router. Packets arrive at input interfaces **so-0/1/0**, **so-3/0/0**, and **so-3/1/0**. The raw packets are directed into a filter-based forwarding routing instance and processed into cflowd records by the monitoring services interfaces **mo-7/1/0**, **mo-7/2/0**, and **mo-7/3/0**.

Copyright © 2013, Juniper Networks, Inc.

The cflowd records are compressed into files at the flow collector interfaces **cp-6/0/0** and **cp-7/0/0** and sent to the FTP server for analysis. Finally, a mandatory class-of-service (CoS) configuration is applied to export channels 0 and 1 on the flow collector interfaces to manage the outgoing processed files.

Figure 12: Flow Collector Interface Topology Diagram

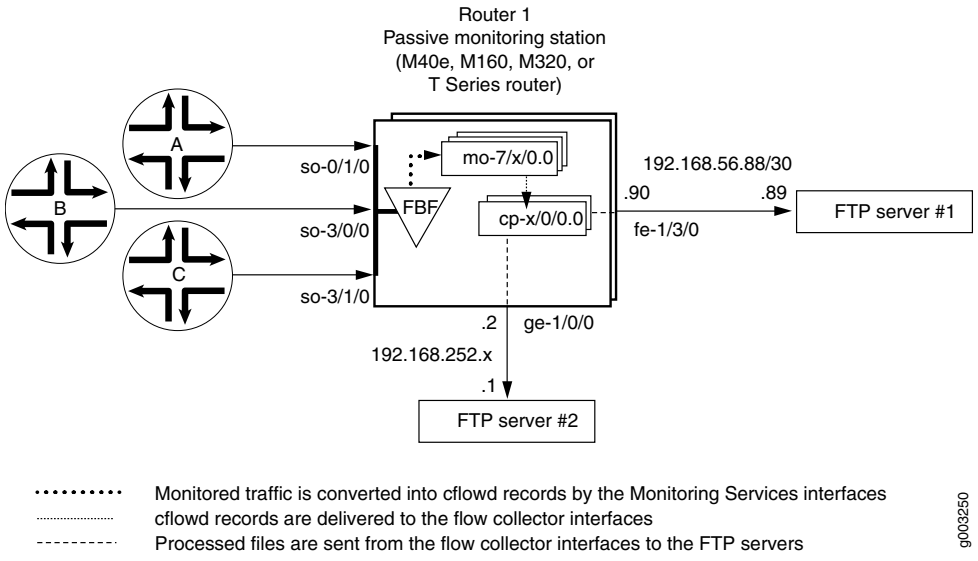


```
[edit]
chassis {
  fpc 6 {
    pic 0 {
      monitoring-services {
        application flow-collector; # This converts a Monitoring Services II or
                    # Multiservices 400 PIC into a flow collector interface.
      }
    }
  }
  fpc 7 {
    pic 0 {
      monitoring-services {
        application flow-collector; # This converts a Monitoring Services II or
                    # Multiservices 400 PIC into a flow collector interface.
      }
    }
  }
}
interfaces {
  cp-6/0/0 {
    unit 0 { # Logical interface .0 on a flow collector interface is export
      family inet { # channel 0 and sends records to the FTP server.
        filter {
          output cp-ftp; # Apply the CoS filter here.
        }
        address 10.0.0.1/32 {
          destination 10.0.0.2;
        }
```

```
    }
  }
  unit 1 { # Logical interface .1 on a flow collector interface is export
    family inet {# channel 1 and sends records to the FTP server.
      filter {
        output cp-ftp; # Apply the CoS filter here.
      }
      address 10.1.1.1/32 {
        destination 10.1.1.2;
      }
    }
  }
  unit 2 { # Logical interface .2 on a flow collector interface is the flow
    family inet { # receive channel that communicates with the Routing Engine.
      address 10.2.2.1/32 { # Do not apply a CoS filter on logical interface .2.
        destination 10.2.2.2;
      }
    }
  }
}
cp-7/0/0 {
  unit 0 {# Logical interface .0 on a flow collector interface is export
    family inet {# channel 0 and sends records to the FTP server.
      filter {
        output cp-ftp;# Apply the CoS filter here.
      }
      address 10.3.3.1/32 {
        destination 10.3.3.2;
      }
    }
  }
  unit 1 {# Logical interface .1 on a flow collector interface is export
    family inet {# channel 1 and sends records to the FTP server.
      filter {
        output cp-ftp;# Apply the CoS filter here.
      }
      address 10.4.4.1/32 {
        destination 10.4.4.2;
      }
    }
  }
  unit 2 {# Logical interface .2 on a flow collector interface is the flow
    family inet {# receive channel that communicates with the Routing Engine.
      address 10.5.5.1/32 {# Do not apply a CoS filter on logical interface .2.
        destination 10.5.5.2;
      }
    }
  }
}
fe-1/3/0 { # This is the exit interface leading to the first FTP server.
  unit 0 {
    family inet {
      address 192.168.56.90/30;
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

```
    ge-1/0/0 { # This is the exit interface leading to the second FTP server.
      unit 0 {
        family inet {
          address 192.168.252.2/24;
        }
      }
    }
    mo-7/1/0 { # This is the first interface that creates cflowd records.
      unit 0 {
        family inet;
      }
    }
    mo-7/2/0 { # This is the second interface that creates cflowd records.
      unit 0 {
        family inet;
      }
    }
    mo-7/3/0 { # This is the third interface that creates cflowd records.
      unit 0 {
        family inet;
      }
    }
    so-0/1/0 { # This is the first input interface that receives traffic to be monitored.
      encapsulation ppp;
      unit 0 {
        passive-monitor-mode; # This allows the interface to be passively monitored.
        family inet {
            filter {
              input catch; # The filter-based forwarding filter is applied here.
            }
        }
      }
    }
    so-3/0/0 { # This is the second interface that receives traffic to be monitored.
      encapsulation ppp;
      unit 0 {
        passive-monitor-mode; # This allows the interface to be passively monitored.
        family inet {
            filter {
              input catch; # The filter-based forwarding filter is applied here.
            }
        }
      }
    }
  so-3/1/0 { # This is the third interface that receives traffic to be monitored.
    encapsulation ppp;
    unit 0 {
      passive-monitor-mode; # This allows the interface to be passively monitored.
      family inet {
        filter {
          input catch; # The filter-based forwarding filter is applied here.
        }
      }
    }
  }
  forwarding-options {
```

```
    monitoring group1 {# Always define your monitoring group here.
      family inet {
        output {
          export-format cflowd-version-5;
          flow-active-timeout 60;
          flow-inactive-timeout 15;
          flow-export-destination collector-pic;  # Sends records to the flow collector.
          interface mo-7/1/0.0 {
            source-address 192.168.252.2;
          }
          interface mo-7/2/0.0 {
            source-address 192.168.252.2;
          }
          interface mo-7/3/0.0 {
            source-address 192.168.252.2;
          }
        }
      }
  }
  firewall {
    family inet {
      filter cp-ftp { # This filter provides CoS for flow collector interface traffic.
        term t1 {
          then forwarding-class expedited-forwarding;
        }
      }
    }
    filter catch { # This firewall filter sends incoming traffic into the
      interface-specific;# filter-based forwarding routing instance.
        term def {
          then {
            count counter;
            routing-instance fbf_instance;
          }
        }
      }
    }
  routing-options {
    interface-routes {
      rib-group inet common;
    }
    rib-groups {
      common {
        import-rib [inet.0 fbf_instance.inet.0];
      }
    }
    forwarding-table {
      export pplb;
    }
  }
  policy-options {
    policy-statement pplb {
      then {
        load-balance per-packet;
      }
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
}
routing-instances {
   fbf_instance { # This instance sends traffic to the monitoring services interface.
      instance-type forwarding;
      routing-options {
         static {
            route 0.0.0.0/0 next-hop mo-7/1/0.0;
         }
      }
   }
}
class-of-service { # A class-of-service configuration for the flow collector interface
   interfaces { # is required for flow collector services.
      cp-6/0/0 {
         scheduler-map cp-map;
      }
      cp-7/0/0 {
         scheduler-map cp-map;
      }
   }
}
scheduler-maps {
   cp-map {
      forwarding-class best-effort scheduler Q0;
      forwarding-class expedited-forwarding scheduler Q1;
      forwarding-class network-control scheduler Q3;
   }
}
schedulers {
   Q0 {
      transmit-rate remainder;
      buffer-size percent 90;
   }
   Q1 {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority strict-high;
   }
   Q3 {
      transmit-rate percent 5;
      buffer-size percent 5;
   }
}
services {
   flow-collector { # Define properties for flow collector interfaces here.
      analyzer-address 10.10.10.1; # This is the IP address of the analyzer.
      analyzer-id server1; # This helps to identify the analyzer.
      retry 3; # Maximum number of attempts by the PIC to send a file transfer log.
      retry-delay 30; # The time interval between attempts to send a file transfer log.
      destinations { # This defines the FTP servers that receive flow collector output.
         "ftp://user@192.168.56.89//tmp/collect1/" { # The primary FTP server.
            password "$9$lXJK8xN-w2oZdbZDHmF30O1"; # SECRET-DATA
         }
         "ftp://user@192.168.252.1//tmp/collect2/" { # The secondary FTP server.
            password "$9$elbvL7-dsgaGVwGjkP3nOBI"; # SECRET-DATA
         }
```

```
    }
    file-specification { # Define sets of flow collector characteristics here.
      def-spec {
        name-format "default-allInt-0-%D_%T-%I_%N.bcp.bi.gz";
        data-format flow-compressed; # The default compressed output format.
      } # When no overrides are specified, a collector uses default transfer values.
      f1 {
        name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
        data-format flow-compressed; # The default compressed output format.
        transfer timeout 1800 record-level 1000000; # Here are configured values.
      }
    }
    interface-map { # Allows you to map interfaces to flow collector interfaces.
      file-specification def-spec; # Flows generated for default traffic are sent to the
      collector cp-7/0/0; # default flow collector interface "cp-7/0/0".
      so-0/1/0.0 { # Flows generated for the so-0/1/0 interface are sent
        collector cp-6/0/0; # to cp-6/0/0, and the file-specification used is
      } # "default."
      so-3/0/0.0 { # Flows generated for the so-3/0/0 interface are sent
        file-specification f1; # to cp-6/0/0, and the file-specification used is "f1."
        collector cp-6/0/0;
      }
      so-3/1/0.0; # Because no settings are defined, flows generated for this
    } # interface use interface cp-7/0/0 and the default file specification.
    transfer-log-archive { # Sends flow collector interface log files to an FTP server.
      filename-prefix so_3_0_0_log;
      maximum-age 15;
      archive-sites {
        "ftp://user@192.168.56.89//tmp/transfers/" {
          password "$9$IFaEyevMXNVsWLsgaU.m6/C";
        }
      }
    ]
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 50

Summary of Flow Collection Configuration Statements

The following sections explain each of the flow collection configuration statements. The statements are organized alphabetically.

analyzer-address

Syntax analyzer-address *address*;

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an IP address for the packet analyzer that overrides the default value.

Options ***address***—IP address for packet analyzer.

Usage Guidelines See "Configuring a Packet Analyzer" on page 1055.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

analyzer-id

Syntax analyzer-id *name*;

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an identifier for the packet analyzer that overrides the default value.

Options *name*—Identifier for packet analyzer.

Usage Guidelines See "Configuring a Packet Analyzer" on page 1055.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

archive-sites

Syntax

```
archive-sites {
  ftp:url {
    password "password";
    username username;
  }
}
```

Hierarchy Level [edit services flow-collector transfer-log-archive]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the destination for transfer logs.

Options The statements are explained separately.

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.

Copyright © 2013, Juniper Networks, Inc.

collector

Syntax collector *interface-name*;

Hierarchy Level [edit services flow-collector interface-map]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the default flow collector interface for interface mapping.

Options **collector *interface-name***—Default flow collector interface.

Usage Guidelines See "Configuring Interface Mappings" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

data-format

Syntax data-format *format*;

Hierarchy Level [edit services flow-collector file-specification variant *variant-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the data format for a specific file format variant.

Options ***format***—Data format. Specify **flow-compressed** as the data format.

Usage Guidelines See "Configuring File Formats" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destinations

Syntax

```
destinations {
  ftp:url {
    password "password";
  }
}
```

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary destination FTP servers.

Options The statements are explained separately.

Usage Guidelines See "Configuring Destination FTP Servers for Flow Records" on page 1055.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

filename-prefix

Syntax filename-prefix *prefix*;

Hierarchy Level [edit services flow-collector transfer-log-archive]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the filename prefix for log files.

Options ***prefix***—Filename identifier.

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

file-specification

See the following sections:

- file-specification (File Format) on page 1069
- file-specification (Interface Mapping) on page 1069

file-specification (File Format)

Syntax

```
file-specification {
  variant variant-number {
    data-format format;
    name-format format;
    transfer {
      record-level number;
      timeout seconds;
    }
  }
}
```

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the file format for the flow collection files.

Options The statements are explained separately.

Usage Guidelines See "Configuring File Formats" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

file-specification (Interface Mapping)

Syntax

```
file-specification {
  variant variant-number;
}
```

Hierarchy Level [edit services flow-collector interface-map]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the default file specification for interface mapping.

Options **variant *variant-number***—Default file format variant.

Usage Guidelines See "Configuring Interface Mappings" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-collector

Syntax

```
flow-collector {
  analyzer-address address;
  analyzer-id name;
  destinations {
    ftp:url {
      password "password";
    }
  }
  file-specification {
    variant variant-number {
      data-format format;
      name-format format;
      transfer {
        record-level number;
        timeout seconds;
      }
    }
  }
  interface-map {
    collector interface-name;
    file-specification variant-number;
    interface-name {
      collector interface-name;
      file-specification variant-number;
    }
  }
  retry number;
  retry-delay seconds;
  transfer-log-archive {
    archive-sites {
      ftp:url {
        password "password";
        username username;
      }
    }
    filename-prefix prefix;
    maximum-age minutes;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the flow collection.

Options The statements are explained separately.

Usage Guidelines See the topics in Flow Collection.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

ftp

See the following sections:

- ftp (Flow Collector Files) on page 1073
- ftp (Transfer Log Files) on page 1074

Copyright © 2013, Juniper Networks, Inc.

ftp (Flow Collector Files)

Syntax ftp:*url*;

Hierarchy Level [edit services flow-collector destination]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary destination FTP server addresses.

Options *url*—FTP server address. The URL can include the following macros, typed in braces:

- **{%D}**—Date
- **{%T)**—Time when the file is created
- **{%I}—**Description string for the logical interface configured using the **collector** *interface-name* statement at the **[edit services flow-collector interface-map]** hierarchy
- **{%N}**—Unique, sequential number for each new file created
- **{am_pm}**—**AM** or **PM**
- **{date}**—Current date using the **{year} {month} {day}** macros
- **{day}**—From **01** through **31**
- **{day_abbr}**—**Sun** through **Sat**
- **{day_full}**—**Sunday** through **Saturday**
- **{generation number}**—Unique, sequential number for each new file created
- **{hour_12}**—From **01** through **12**
- **{hour_24}**—From **00** through **23**
- **{ifalias}**—Description string for the logical interface configured using the **collector** statement at the **[edit services flow-collector interface-map]** hierarchy
- **{minute}**—From **00** through **59**
- **{month}**—From **01** through **12**
- **{month_abbr}**—**Jan** through **Dec**
- **{month_full}**—**January** through **December**
- **{num_zone}**—From **-2359** to **+2359**; this macro is not supported
- **{second**}—From **00** through **60**
- **{time}**—Time the file is created, using the **{hour_24} {minute} {second}** macros
- **{time_zone}**—Time zone code name of the locale; for example, **gmt** (this macro is not supported).
- **{year}**—In the format YYYY; for example, **1970**
- **{year_abbr}**—From **00** through **99**

Copyright © 2013, Juniper Networks, Inc.

Usage Guidelines See "Configuring Destination FTP Servers for Flow Records" on page 1055.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

ftp (Transfer Log Files)

Syntax ftp:*url*;

Hierarchy Level [edit services flow-collector transfer-log-archive archive-sites]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary destination FTP server addresses.

Options *url*—FTP server address.

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interface-map

Syntax

```
interface-map {
  collector interface-name;
  file-specification variant-number;
  interface-name {
    collector interface-name;
    file-specification variant-number;
  }
}
```

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Match an input interface with a flow collector interface and apply the preset file specifications to the input interface.

Options The statements are explained separately.

Usage Guidelines See "Configuring Interface Mappings" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

maximum-age

Syntax maximum-age *minutes*;

Hierarchy Level [edit services flow-collector transfer-log-archive]

Release Information Statement introduced before Junos OS Release 7.4.

Description Maximum age of transfer log file.

Options **maximum-age *minutes***—Transfer log file age.
Range: 1 through 360

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

name-format

Syntax name-format "*format*";

Hierarchy Level [edit services flow-collector file-specification variant *variant-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the name format for a specific file format. The files may include supported macros. Use macros to organize files on the external machine to which they are exported from the collector PIC.

Options *format*—Specify the filename format, within quotation marks. The name format can include the following macros, typed in braces:

- **{%D}**—Date
- **{%T)**—Time when the file is created
- **{%I}—**Description string for the logical interface configured using the **collector** statement at the **[edit services flow-collector interface-map]** hierarchy level
- **{%N}**—Unique, sequential number for each new file created
- **{am_pm}**—**AM** or **PM**
- **{date}**—Current date using the **{year} {month} {day}** macros
- **{day}**—From **01** through **31**
- **{day_abbr}**—**Sun** through **Sat**
- **{day_full}**—**Sunday** through **Saturday**
- **{generation number}**—Unique, sequential number for each new file created
- **{hour_12}**—From **01** through **12**
- **{hour_24}**—From **00** through **23**
- **{ifalias}**—Description string for the logical interface configured using the **collector** statement at the **[edit services flow-collector interface-map]** hierarchy level
- **{minute}**—From **00** through **59**
- **{month}**—From **01** through **12**
- **{month_abbr}**—**Jan** through **Dec**
- **{month_full}**—**January** through **December**
- **{num_zone}**—From **-2359** through **+2359**; this macro is not supported
- **{second**}—From **00** through **60**
- **{time}**—Time the file is created, using the **{hour_24} {minute} {second}** macros
- **{time_zone}**—Time zone code name of the locale; for example, **gmt** (this macro is not supported).

Copyright © 2013, Juniper Networks, Inc.

- **{year}**—In the format YYYY; for example, **1970**
- **{year_abbr}**—From **00** through **99**

Usage Guidelines See "Configuring File Formats" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

password

See the following sections:

- password (Flow Collector File Servers) on page 1078
- password (Transfer Log File Servers) on page 1078

password (Flow Collector File Servers)

Syntax password "*password*";

Hierarchy Level [edit services flow-collector destination ftp:*url*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary destination FTP server password.

Options ***password***—FTP server password.

Usage Guidelines See "Configuring Destination FTP Servers for Flow Records" on page 1055.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

password (Transfer Log File Servers)

Syntax password "*password*";

Hierarchy Level [edit services flow-collector transfer-log-archive archive-sites]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary and secondary destination FTP server password.

Options ***password***—FTP server password.

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

retry (Services Flow Collector)

Syntax retry *number*;

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the maximum number of attempts the flow collector interface will make to transfer log files to the FTP server.

Options ***number***—Maximum number of transfer retry attempts.
Range: 0 through 10

Usage Guidelines See "Configuring Retry Attempts" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

retry-delay

Syntax retry-delay *seconds*;

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the amount of time the flow collector interface waits between retry attempts.

Options ***seconds***—Amount of time between transfer retry attempts.
Range: 0 through 60

Usage Guidelines See "Configuring Retry Attempts" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

transfer

Syntax

```
transfer {
  record-level number;
  timeout seconds;
}
```

Hierarchy Level [edit services flow-collector file-specification variant *variant-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify when to send the flow collection file. The file is sent when either of the two conditions is met.

Options **record-level *number***—Number of flow collection files collected.

timeout *seconds*—Timeout duration.

Usage Guidelines See "Configuring File Formats" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

transfer-log-archive

Syntax

```
transfer-log-archive {
  archive-sites {
    ftp:url {
      password "password";
      username username;
    }
  }
  filename-prefix prefix;
  maximum-age minutes;
}
```

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the filename prefix, maximum age, and destination FTP server for log files containing the transfer activity history for a flow collector interface.

Options The statements are explained separately.

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

username (Services)

Syntax username *user-name*;

Hierarchy Level [edit services flow-collector transfer-log-archive archive-sites]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the username for the transfer log server.

Options ***username***—FTP server username.

Usage Guidelines See "Configuring Transfer Logs" on page 1057.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

variant

Syntax

```
variant variant-number {
  data-format format;
  name-format format;
  transfer {
    record-level number;
    timeout seconds;
  }
}
```

Hierarchy Level [edit services flow-collector file-specification]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a variant of the file format.

Options The statements are explained separately.

Usage Guidelines See "Configuring File Formats" on page 1056.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 51

Dynamic Flow Capture Configuration Guidelines

Dynamic flow capture enables you to capture packet flows on the basis of dynamic filtering criteria. Specifically, you can use this feature to forward passively monitored packet flows that match a particular filter list to one or more destinations using an on-demand control protocol.

This chapter contains the following sections:

- Dynamic Flow Capture Architecture on page 1083
- Configuring Dynamic Flow Capture on page 1085
- Example: Configuring Dynamic Flow Capture on page 1092

Dynamic Flow Capture Architecture

The architecture consists of one or more *control sources* that send requests to a Juniper Networks router to monitor incoming data, and then forward any packets that match specific filter criteria to a set of one or more *content destinations*. The architectural components are defined as follows:

- Control source—A client that monitors electronic data or voice transfer over the network. The control source sends filter requests to the Juniper Networks router using the Dynamic Task Control Protocol (DTCP), specified in draft-cavuto-dtcp-03.txt at **http://www.ietf.org/internet-drafts**. The control source is identified by a unique identifier and an optional list of IP addresses.
- Monitoring platform—A T Series or M320 router containing one or more Dynamic Flow Capture (DFC) PICs, which support dynamic flow capture processing. The monitoring platform processes the requests from the control sources, creates the filters, monitors incoming data flows, and sends the matched packets to the appropriate content destinations.
- Content destination—Recipient of the matched packets from the monitoring platform. Typically the matched packets are sent using an IP Security (IPsec) tunnel from the monitoring platform to another router connected to the content destination. The content destination and the control source can be physically located on the same host. For more information on IPsec tunnels, see IPsec Properties.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The DFC PIC (either a Monitoring Services III PIC or Multiservices 400 PIC) forwards the entire packet content to the content destination, rather than to a content record as is done with cflowd or flow aggregation version 9 templates.

Figure 13 on page 1084 shows a sample topology. The number of control sources and content destinations is arbitrary.

Figure 13: Dynamic Flow Capture Topology

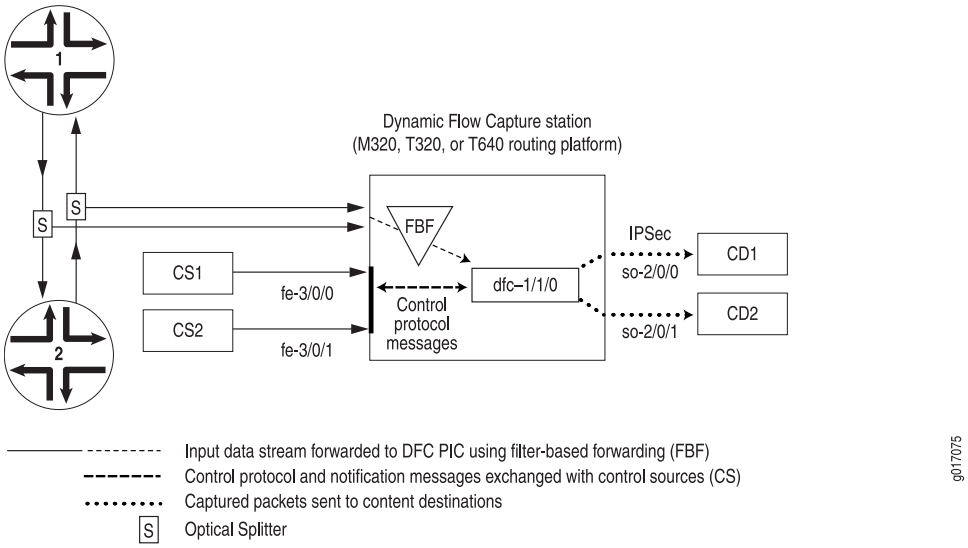

Liberal Sequence Windowing

Each DTCP packet (add, delete, list, and refresh packets) contains a 64-bit sequence number to identify the order of the packets. Because the network is connectionless, the DTCP packets can arrive out of order to the router running the DFC application.

The *liberal sequence window* feature implements a negative window for the sequence numbers received in the DTCP packets. It enables the DFC application to accept not only DTCP packets with sequence numbers greater than those previously received, but also DTCP packets with lesser sequence numbers, up to a certain limit. This limit is the negative window size; the positive and negative window sizes are +256 and –256 respectively, relative to the current maximum sequence number received. No configuration is required to activate this feature; the window sizes are hard-coded and nonconfigurable.

Intercepting IPv6 Flows

Starting with Junos OS Release 11.4, the Dynamic Flow Capture (DFC) application also supports intercepting IPv6 flows in M320, T320, T640, and T1600 routers with a Multiservices 400 or Multiservices 500 PIC. The DFC application can intercept passively monitored IPv6 traffic only. All support for IPv4 interception remains the same. The interception of IPv6 traffic happens in the same way the filters capture IPv4 flows. With

Copyright © 2013, Juniper Networks, Inc.

the introduction of IPv6 interception, both IPv4 and IPv6 filters can coexist. The mediation device , however, cannot be located in an IPv6 network.

The DFC application does not support interception of VPLS and MPLS traffic. The application cannot intercept Address Resolution Protocol (ARP) or other Layer 2 exception packets. The interception filter can be configured to timeout based on factors like total time (seconds), idle time (seconds), total packets or total data transmitted (bytes).

Configuring Dynamic Flow Capture

To configure dynamic flow capture, include the **dynamic-flow-capture** statement at the **[edit services]** hierarchy level:

```
[edit services]
dynamic-flow-capture {
  capture-group client-name {
    content-destination identifier {
      address address;
      hard-limit bandwidth;
      hard-limit-target bandwidth;
      soft-limit bandwidth;
      soft-limit-clear bandwidth;
      ttl hops;
    }
    control-source identifier {
      allowed-destinations [ destinations ];
      minimum-priority value;
      no-syslog;
      notification-targets address port port-number;
      service-port port-number;
      shared-key value;
      source-addresses [ addresses ];
    }
    duplicates-dropped-periodicity seconds;
    input-packet-rate-threshold rate;
    interfaces interface-name;
    max-duplicates number;
    pic-memory-threshold percentage percentage;
  }
  g-duplicates-dropped-periodicity seconds;
  g-max-duplicates number;
  traceoptions{
    file filename <files number> <size size> <world-readable | non-world-readable>;
  }
}
```

This section describes the following tasks for configuring dynamic flow capture:

- Configuring the Capture Group on page 1086
- Configuring the Content Destination on page 1086
- Configuring the Control Source on page 1087
- Configuring the DFC PIC Interface on page 1088
- Configuring the Firewall Filter on page 1089

Copyright © 2013, Juniper Networks, Inc.

- Configuring System Logging on page 1090
- Configuring Tracing Options for Dynamic Flow Capture Events on page 1090
- Configuring Thresholds on page 1090
- Limiting the Number of Duplicates of a Packet on page 1091

Configuring the Capture Group

A capture group defines a profile of dynamic flow capture configuration information. The static configuration includes information about control sources, content destinations, and notification destinations. Dynamic configuration is added through interaction with control sources using a control protocol.

To configure a capture group, include the **capture-group** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
capture-group client-name {
  content-destination identifier {
    address address;
    hard-limit bandwidth;
    hard-limit-target bandwidth;
    soft-limit bandwidth;
    soft-limit-clear bandwidth;
    ttl hops;
  }
  control-source identifier {
    allowed-destinations [ destinations ];
    minimum-priority value;
    no-syslog;
    notification-targets address port port-number;
    service-port port-number;
    shared-key value;
    source-addresses [ addresses ];
  }
  duplicates-dropped-periodicity seconds;
  input-packet-rate-threshold rate;
  interfaces interface-name;
  max-duplicates number;
  pic-memory-threshold percentage percentage;
}
```

To specify the **capture-group**, assign it a unique ***client-name*** that associates the information with the requesting control sources.

Configuring the Content Destination

You must specify a destination for the packets that match DFC PIC filter criteria. To configure the content destination, include the **content-destination** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
content-destination identifier {
  address address;
  hard-limit bandwidth;
  hard-limit-target bandwidth;
  soft-limit bandwidth;
```


Copyright © 2013, Juniper Networks, Inc.

```
    soft-limit-clear bandwidth;
    ttl hops;
}
```

Assign the **content-destination** a unique ***identifier***. You must also specify its IP address and you can optionally include additional settings:

- **address**—The DFC PIC interface appends an IP header with this destination address on the matched packet (with its own IP header and contents intact) and sends it out to the content destination.
- **ttl**—The time-to-live (TTL) value for the IP-IP header. By default, the TTL value is 255. Its range is 0 through 255.
- Congestion thresholds—You can specify per-content destination bandwidth limits that control the amount of traffic produced by the DFC PIC during periods of congestion. The thresholds are arranged in two pairs: **hard-limit** and **hard-limit-target**, and **soft-limit** and **soft-limit-clear**. You can optionally include one or both of these paired settings. All four settings are 10–second average bandwidth values in bits per second. Typically **soft-limit-clear** < **soft-limit** < **hard-limit-target** < **hard-limit**. When the content bandwidth exceeds the **soft-limit** setting:

  1. A congestion notification message is sent to each control source of the criteria that point to this content destination
  2. If the control source is configured for **syslog**, a system log message is generated.
  3. A latch is set, indicating that the control sources have been notified. No additional notification messages are sent until the latch is cleared, when the bandwidth falls below the **soft-limit-clear** value.

  When the bandwidth exceeds the **hard-limit** value:

  1. The dynamic flow capture application begins deleting criteria until the bandwidth falls below the **hard-limit-target** value.
  2. For each criterion deleted, a CongestionDelete notification is sent to the control source for that criterion.
  3. If the control source is configured for **syslog**, a log message is generated.

  The application evaluates criteria for deletion using the following data:

  - Priority—Lower priority criteria are purged first, after adjusting for control source minimum priority.
  - Bandwidth—Higher bandwidth criteria are purged first.
  - Timestamp—The more recent criteria are purged first.

Configuring the Control Source

You configure information about the control source, including allowed source addresses and destinations and authentication key values. To configure the control source information, include the **control-source** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
control-source identifier {
  allowed-destinations [ destination-identifiers ];
  minimum-priority value;
  no-syslog;
  notification-targets address port port-number;
  service-port port-number;
  shared-key value;
  source-addresses [ addresses ];
}
```

Assign the **control-source** statement a unique ***identifier***. You can also include values for the following statements:

- **allowed-destinations**—One or more content destination identifiers to which this control source can request that matched data be sent in its control protocol requests. If you do not specify any content destinations, all available destinations are allowed.
- **minimum-priority**—Value assigned to the control source that is added to the priority of the criteria in the DTCP ADD request to determine the total priority for the criteria. The lower the value, the higher the priority. By default, **minimum-priority** has a value of 0 and the allowed range is 0 through 254.
- **notification-targets**—One or more destinations to which the DFC PIC interface can log information about control protocol-related events and other events such as PIC bootup messages. You configure each **notification-target** entry with an IP **address** value and a User Datagram Protocol (UDP) **port** number.
- **service-port**—UDP port number to which the control protocol requests are directed. Control protocol requests that are not directed to this port are discarded by DFC PIC interfaces.
- **shared-key**—20-byte authentication key value shared between the control source and the DFC PIC monitoring platform.
- **source-addresses**—One or more allowed IP addresses from which the control source can send control protocol requests to the DFC PIC monitoring platform. These are /32 addresses.

Configuring the DFC PIC Interface

You specify the interface that interacts with the control sources configured in the same capture group. A Monitoring Services III PIC can belong to only one capture group, and you can configure only one PIC for each group.

To configure a DFC PIC interface, include the **interfaces** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
interfaces interface-name;
```

You specify DFC interfaces using the **dfc-** identifier at the **[edit interfaces]** hierarchy level. You must specify three logical units on each DFC PIC interface, numbered **0**, **1**, and **2**. You cannot configure any other logical interfaces.

- **unit 0** processes control protocol requests and responses.

Copyright © 2013, Juniper Networks, Inc.

- **unit 1** receives monitored data.
- **unit 2** transmits the matched packets to the destination address.

The following example shows the configuration necessary to set up a DFC PIC interface and intercept both IPv4 and IPv6 traffic:

```
[edit interfaces dfc-0/0/0]
unit 0 {
  family inet {
    filter {
      output high; #Firewall filter to route control packets
      # through 'network-control' forwarding class. Control packets
      # are loss sensitive.
    }
    address 10.1.0.0/32 { # DFC PIC address
      destination 10.36.100.1; # DFC PIC address used by
      # the control source to correspond with the
      # monitoring platform
    }
  }
unit 1 { # receive data packets on this logical interface
  family inet; # receive IPv4 traffic for interception
  family inet6; # receive IPv6 traffic for interception
}
unit 2 { # send out copies of matched packets on this logical interface
  family inet;
}
```

In addition, you must configure the dynamic flow capture application to run on the DFC PIC in the correct chassis location. The following example shows this configuration at the **[edit chassis]** hierarchy level:

```
fpc 0 {
  pic 0 {
    monitoring-services application dynamic-flow-capture;
  }
}
```

For more information on configuring chassis properties, see the Junos OS System Basics Configuration Guide.

Configuring the Firewall Filter

You can specify the firewall filter to route control packets through the network control forwarding class. The control packets are loss sensitive. To configure the firewall filter, include the following statements at the **[edit]** hierarchy level:

```
firewall {
  family inet {
    filter high {
      term all {
        then forwarding-class network-control;
      }
    }
  }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
```

Configuring System Logging

By default, control protocol activity is logged as a separate system log facility, **dfc**. To modify the filename or level at which control protocol activity is recorded, include the following statements at the **[edit syslog]** hierarchy level:

```
file dfc.log {
    dfc any;
}
```

To cancel logging, include the **no-syslog** statement at the **[edit services dynamic-flow-capture capture-group** *client-name* **control-source** *identifier***]** hierarchy level:

```
no-syslog;
```



NOTE: The dynamic flow capture (dfc-) interface supports up to 10,000 filter criteria. When more than 10,000 filters are added to the interface, the filters are accepted, but system log messages are generated indicating that the filter is full.

Configuring Tracing Options for Dynamic Flow Capture Events

You can enable tracing options for dynamic flow capture events by including the **traceoptions** statement at the **[edit services dynamic-flow-capture]** hierarchy level.

When you include the **traceoptions** configuration, you can also specify the trace file name, maximum number of trace files, the maximum size of trace files, and whether the trace file can be read by all users or not.

To enable tracing options for dynamic flow capture events, include the following configuration at the **[edit services dynamic-flow-capture]** hierarchy level:

```
traceoptions{
    file filename <files number> <size size> <world-readable | non-world-readable>;
}
```

To disable tracing for dynamic flow control events, delete the **traceoptions** configuration from the **[edit services dynamic-flow-capture]** hierarchy level.



NOTE: In Junos OS releases earlier than 9.2R1, tracing of dynamic flow control was enabled by default, and the logs were saved to the /var/log/dfcd directory.

Configuring Thresholds

You can optionally specify threshold values for the following situations in which warning messages will be recorded in the system log:

- Input packet rate to the DFC PIC interfaces

Copyright © 2013, Juniper Networks, Inc.

- Memory usage on the DFC PIC interfaces

To configure threshold values, include the **input-packet-rate-threshold** or **pic-memory-threshold** statements at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
input-packet-rate-threshold rate;
pic-memory-threshold percentage percentage;
```

If these statements are not configured, no threshold messages are logged. The threshold settings are configured for the capture group as a whole.

The range of configurable values for the **input-packet-rate-threshold** statement is 0 through 1 Mpps. The PIC calibrates the value accordingly; the Monitoring Services III PIC caps the threshold value at 300 Kpps and the Multiservices 400 PIC uses the full configured value. The range of values for the **pic-memory-threshold** statement is 0 to 100 percent.

Limiting the Number of Duplicates of a Packet

You can optionally specify the maximum number of duplicate packets the DFC PIC is allowed to generate from a single input packet. This limitation is intended to reduce the load on the PIC when packets are sent to multiple destinations. When the maximum number is reached, the duplicates are sent to the destinations with the highest criteria class priority. Within classes of equal priority, criteria having earlier timestamps are selected first.

To configure this limitation, include the **max-duplicates** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level:

```
max-duplicates number;
```

You can also apply the limitation on a global basis for the DFC PIC by including the **g-max-duplicates** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
g-max-duplicates number;
```

By default, the maximum number of duplicates is set to 3. The range of allowed values is 1 through 64. A setting for **max-duplicates** for an individual capture-group overrides the global setting.

In addition, you can specify the frequency with which the application sends notifications to the affected control sources that duplicates are being dropped because the threshold has been reached. You configure this setting at the same levels as the maximum duplicates settings, by including the **duplicates-dropped-periodicity** statement at the **[edit services dynamic-flow-capture capture-group *client-name*]** hierarchy level or the **g-duplicates-dropped-periodicity** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
duplicates-dropped-periodicity seconds;
g-duplicates-dropped-periodicity seconds;
```

Copyright © 2013, Juniper Networks, Inc.

As with the **g-max-duplicates** statement, the **g-duplicates-dropped-periodicity** statement applies the setting globally for the application and is overridden by a setting applied at the capture-group level. By default, the frequency for sending notifications is 30 seconds.

Example: Configuring Dynamic Flow Capture

The following example includes all parts of a complete dynamic flow capture configuration.

Configure the DFC PIC interface:

```
[edit interfaces dfc-0/0/0]
unit 0 {
   family inet {
      filter {
         output high; #Firewall filter to route control packets
         # through 'network-control' forwarding class. Control packets
         # are loss sensitive.
      }
      address 10.1.0.0/32 { # DFC PIC address
         destination 10.36.100.1; # DFC PIC address used by
         # the control source to correspond with the
         # monitoring platform
      }
   }
unit 1 { # receive data packets on this logical interface
   family inet;
   family inet6;
}
unit 2 { # send out copies of matched packets on this logical interface
   family inet;
}
```

Configure the capture group:

```
services dynamic-flow-capture {
   capture-group g1 {
      interfaces dfc-0/0/0;
      input-packet-rate-threshold 90k;
      pic-memory-threshold percentage 80;
      control-source cs1 {
         source-addresses 10.36.41.1;
         service-port 2400;
         notification-targets {
            10.36.41.1 port 2100;
         }
         shared-key "$9$ASxdsYoX7wg4aHk";
         allowed-destinations cd1;
      }
      content-destination cd1 {
         address 10.36.70.2;
         ttl 244;
      }
   }
}
```


Copyright © 2013, Juniper Networks, Inc.

Configur3 filter-based forwarding (FBF) to the DFC PIC interface, logical unit 1.

For more information about configuring passive monitoring interfaces, see "Enabling Passive Flow Monitoring" on page 968.

```
interfaces so-1/2/0 {
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode;
        family inet {
            filter {
                input catch;
            }
        }
    }
}
```

Configure the firewall filter:

```
firewall {
    filter catch {
        interface-specific;
        term def {
            then {
                count counter;
                routing-instance fbf_inst;
            }
        }
    }
    family inet {
        filter high {
            term all {
                then forwarding-class network-control;
            }
        }
    }
}
```

Configure a forwarding routing instance. The next hop points specifically to the logical interface corresponding to **unit 1**, because only this particular logical unit is expected to relay monitored data to the DFC PIC.

```
routing-instances fbf_inst {
    instance-type forwarding;
    routing-options {
        static {
            route 0.0.0.0/0 next-hop dfc-0/0/0.1;
        }
    }
}
```

Configure routing table groups:

```
[edit]
routing-options {
    interface-routes {
        rib-group inet common;
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
    rib-groups {
      common {
        import-rib [ inet.0 fbf_inst.inet.0 ];
      }
    }
    forwarding-table {
      export pplb;
    }
}
```

Configure interfaces to the control source and content destination:

```
interfaces fe-4/1/2 {
  description "to cs1 from dfc";
  unit 0 {
    family inet {
      address 10.36.41.2/30;
    }
  }
}
interfaces ge-7/0/0 {
  description "to cd1 from dfc";
  unit 0 {
    family inet {
      address 10.36.70.1/30;
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 52

Flow-Tap Configuration Guidelines

Dynamic flow capture enables you to capture packet flows on the basis of dynamic filtering criteria, using Dynamic Tasking Control Protocol (DTCP) requests. The flow-tap application extends the use of this protocol to intercept IPv4 and IPv6 packets in an active monitoring router and send a copy of packets that match filter criteria to one or more content destinations. Flow-tap data can be used in the following applications:

- Flexible trend analysis for detection of new security threats
- Lawful intercept

Flow-tap service is supported on M Series and T Series routers, except M160 and TX Matrix routers. Flow-tap filters are applied on all IPv4 traffic and do not add any perceptible delay in the forwarding path. Flow-tap filters can also be applied on IPv6 traffic. For security, filters installed by one client are not visible to others and the CLI configuration does not reveal the identity of the monitored target. A lighter version of the application is supported on MX Series routers only; for more information, see "Configuring FlowTapLite" on page 1100.



NOTE: For information about dynamic flow capture, see "Dynamic Flow Capture Configuration Guidelines" on page 1083. For information about DTCP, see **draft-cavuto-dtcp-01.txt at http://www.ietf.org/internet-drafts.**

To configure flow-tap services, include the **flow-tap** statement at the **[edit services]** hierarchy level. You can also specify whether you want to apply the flow-tap service to IPv4 traffic or IPv6 traffic by including the **family inet | inet6** statement. If the **family** statement is not included in the configuration, the flow-tap service is applied only to the IPv4 traffic.

```
flow-tap {
  interface interface-name;
  family inet | inet6;
}
```

Other statements are configured at the **[edit interfaces]** and **[edit system]** hierarchy levels.

Copyright © 2013, Juniper Networks, Inc.

This chapter contains the following sections:

- Flow-Tap Architecture on page 1096
- Configuring the Flow-Tap Service on page 1097
- Configuring FlowTapLite on page 1100
- Examples: Configuring Flow-Tap Services on page 1101

Flow-Tap Architecture

The flow-tap architecture consists of one or more *mediation devices* that send requests to a Juniper Networks router to monitor incoming data and forward any packets that match specific filter criteria to a set of one or more *content destinations*:

- Mediation device—A client that monitors electronic data or voice transfer over the network. The mediation device sends filter requests to the Juniper Networks router using the DTCP. The clients are not identified for security reasons, but have permissions defined by a set of special login classes. Each system can support up to 16 different mediation devices for each user, up to a maximum of 64 mediation devices for the whole system.
- Monitoring platform—An M Series or T Series router containing one or more Adaptive Services (AS) or Multiservices PICs, which are configured to support the flow-tap application. The monitoring platform processes the requests from the mediation devices, applies the dynamic filters, monitors incoming data flows, and sends the matched packets to the appropriate content destinations.
- Content destination—Recipient of the matched packets from the monitoring platform. Typically the matched packets are sent using an IP Security (IPsec) tunnel from the monitoring platform to another router connected to the content destination. The content destination and the mediation device can be physically located on the same host. For more information about IPsec tunnels, see IPsec Properties.
- Dynamic filters—Firewall filters automatically generated by the Packet Forwarding Engine and applied to all routing instances. Each term in the filter includes a **flow-tap** action that is similar to the existing **sample** or **port-mirroring** actions. As long as one of the filter terms matches an incoming packet, the router copies the packet and forwards it to the Adaptive Services or Multiservices PIC that is configured for flow-tap service. The Adaptive Services or Multiservices PIC runs the packet through the client filters and sends a copy to each matching content destination.

Following is a sample filter configuration; note that it is dynamically generated by the router (no user configuration required):

```
filter combined_LEA_filter {
  term LEA1_filter {
    from {
      source-address 1.2.3.4;
      destination-address 3.4.5.6;
    }
    then {
      flow-tap;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    term LEA2_filter {
      from {
        source-address 10.1.1.1;
        source-port 23;
      }
      then {
        flow-tap;
      }
    }
}
```

Figure 14 on page 1097 shows a sample topology that uses two mediation devices and two content destinations.

Figure 14: Flow-Tap Topology

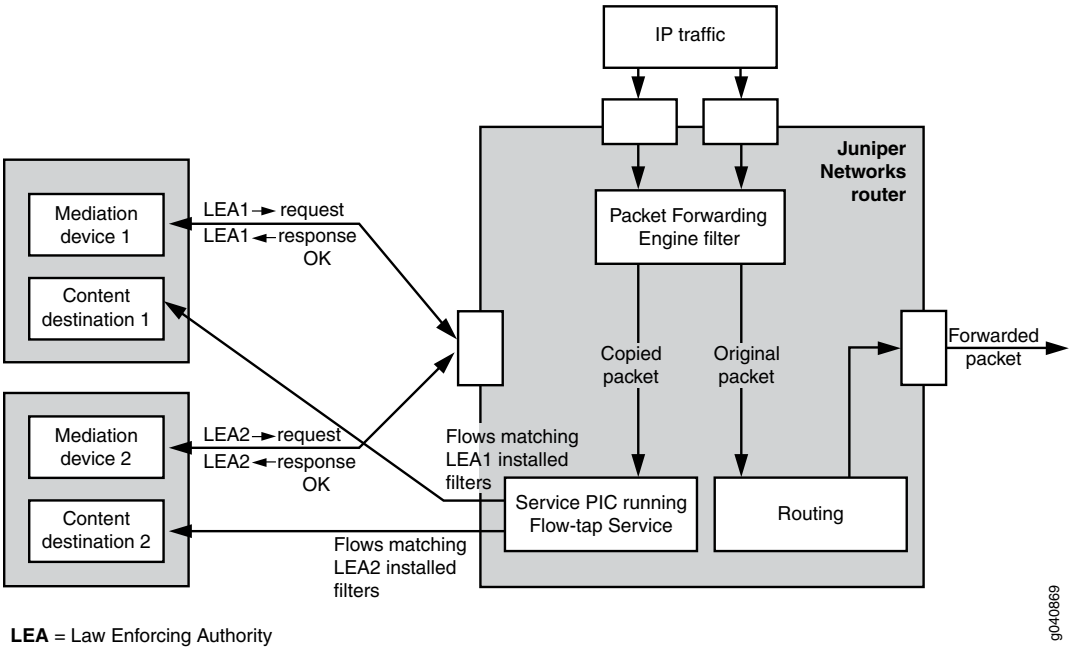

LEA = Law Enforcing Authority

Configuring the Flow-Tap Service

This section describes the following tasks for configuring flow-tap service:

- Configuring the Flow-Tap Interface on page 1097
- Strengthening Flow-Tap Security on page 1098
- Restrictions on Flow-Tap Services on page 1099

Configuring the Flow-Tap Interface

To configure an adaptive services interface for flow-tap service, include the **interface** statement at the **[edit services flow-tap]** hierarchy level:

```
interface sp-fpc/pic/port.unit-number;
family inet | inet6;
```

Copyright © 2013, Juniper Networks, Inc.

You can assign any Adaptive Services or Multiservices PIC in the active monitoring router for flow-tap service, and use any logical unit on the PIC.

You can specify the type of traffic for which you want to apply the flow-tap service by including the **family inet | inet6** statement. If the **family** statement is not included, the flow-tap service is, by default, applied to the IPv4 traffic. To apply flow-tap service to IPv6 traffic, you must include the **family inet6** statement in the configuration. To enable the flow-tap service for IPv4 and IPv6 traffic, you must explicitly configure the **family** statement for both **inet** and **inet6** families.



NOTE: You cannot configure dynamic flow capture and flow-tap features on the same router simultaneously.

You must also configure the logical interface at the **[edit interfaces]** hierarchy level:

```
interface sp-fpc/pic/port {
    unit logical-unit-number {
        family inet;
        family inet6;
    }
}
```



NOTE: If you do not include the family inet6 statement in the configuration, IPv6 flows will not be intercepted.

Strengthening Flow-Tap Security

You can add an extra level of security to Dynamic Tasking Control Protocol (DTCP) transactions between the mediation device and the router by enabling DTCP sessions on top of the SSH layer. To configure SSH settings, include the **flow-tap-dtcp** statement at the **[edit system services]** hierarchy level:

```
flow-tap-dtcp {
    ssh {
        connection-limit value;
        rate-limit value;
    }
}
```

To configure client permissions for viewing and modifying flow-tap configurations and for receiving tapped traffic, include the **permissions** statement at the **[edit system login class *class-name*]** hierarchy level:

```
permissions [permissions];
```

The permissions needed to use flow-tap features are as follows:

- **flow-tap**—Can view flow-tap configuration
- **flow-tap-control**—Can modify flow-tap configuration

Copyright © 2013, Juniper Networks, Inc.

- **flow-tap-operation**—Can tap flows

You can also specify user permissions on a RADIUS server, for example:

```
Bob Auth-Type := Local, User-Password = = “abc123”
Juniper-User-Permissions = “flow-tap-operation”
```

For details on **[edit system]** and RADIUS configuration, see the Junos OS System Basics Configuration Guide.

Restrictions on Flow-Tap Services

The following restrictions apply to flow-tap services:

- You cannot configure dynamic flow capture and flow-tap features on the same router simultaneously.
- On routers that support LMNR-based FPCs, you cannot configure the flow-tap service for IPv6 along with port mirroring or sampling of IPv6 traffic. This restriction holds good even if the router does not have any LMNR-based FPC installed on it. However, there is no restriction on configuring flow-tap service on routers that are configured for port mirroring or sampling of IPv4 traffic.
- Flow-tap service does not support interception of MPLS and virtual private LAN service (VPLS).
- Flow-tap service cannot intercept Address Resolution Protocol (ARP) and other Layer 2 exceptions.
- IPv4 and IPv6 intercept filters can coexist on a system, subject to a combined maximum of 100 filters.
- When the dynamic flow capture process or the Adaptive Services or Multiservices PIC configured for flow-tap processing restarts, all filters are deleted and the mediation devices are disconnected.
- Only the first fragment of an IPv4 fragmented packet stream is sent to the content destination.
- Port mirroring might not work in conjunction with flow-tap processing.
- Running the flow-tap application over an IPsec tunnel on the same router can cause packet loops and is not supported.
- M10i routers do not support the standard flow-tap application, but do support FlowTapLite (see “Configuring FlowTapLite” on page 1100). Flow-tap and FlowTapLite cannot be configured simultaneously on the same chassis.
- PIC-based flow-tap is not supported on M7i and M10i routers equipped with an Enhanced Compact Forwarding Engine Board (CFEB-E).
- You cannot configure flow-tap services on channelized interfaces.

Copyright © 2013, Juniper Networks, Inc.

Configuring FlowTapLite

A lighter version of the flow-tap application is available on MX Series routers and also on M320 routers with Enhanced III Flexible PIC Concentrators (FPCs). All of the functionality resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC).



NOTE: On M320 routers only, if the replacement of FPCs results in a mode change, you must restart the dynamic flow capture process manually by disabling and then re-enabling the CLI configuration.

FlowTapLite uses the same DTCP-SSH architecture to install the Dynamic Tasking Control Protocol (DTCP) filters and authenticate the users as the original flow-tap application and supports up to 3000 filters per chassis.



NOTE: The original flow-tap application and FlowTapLite cannot be used at the same time.

To configure FlowTapLite, include the **flow-tap** statement at the **[edit services]** hierarchy level:

```
flow-tap {
    tunnel-interface interface-name;
}
```

For the Packet Forwarding Engine to encapsulate the intercepted packet, it must send the packet to a tunnel logical (**vt-**) interface. You need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use. To create the tunnel interface, include the following configuration:

```
chassis {
    fpc number {
        pic number {
            tunnel-services {
                bandwidth (1g | 10g);
            }
        }
    }
}
```



NOTE: Currently FlowTapLite supports only one tunnel interface per instance.

For more information about this configuration, see the Junos OS System Basics Configuration Guide.

To configure the logical interfaces and assign them to the dynamic flow capture process, include the following configuration:

Copyright © 2013, Juniper Networks, Inc.

```
interfaces {
  vt-fpc/pic/port {
    unit 0 {
      family inet;
      family inet6;
    }
  }
}
```

NOTE: If a service PIC or DPC is available, you can use its tunnel interface for the same purpose.

NOTE: If you do not include the **family inet6** statement in the configuration, IPv6 flows will not be intercepted.

Examples: Configuring Flow-Tap Services

The following example shows all parts of a complete flow-tap configuration. The example configuration intercepts IPv4 and IPv6 flows.

```
services {
  flow-tap {
    interface sp-1/2/0.100;
  }
}
interfaces {
  sp-1/2/0 {
    unit 100 {
      family inet;
      family inet6;
    }
  }
}
system {
  services {
    flow-tap-dtcp {
      ssh {
        connection-limit 5;
        rate-limit 5;
      }
    }
  }
  login {
    class ft-class {
      permissions flow-tap-operation;
    }
    user ft-user1 {
      class ft-class;
      authentication {
        encrypted-password “xxxx”;
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
```

The following example shows a FlowTapLite configuration that intercepts IPv4 and IPv6 flows:

```
system {
  login {
    class flowtap {
      permissions flow-tap-operation;
    }
    user ftap {
      uid 2000;
      class flowtap;
      authentication {
        encrypted-password "$1$nZfwNn4L$TWi/oxFwFZyOyyxN/87Jv0"; ##
          SECRET-DATA
      }
    }
  }
  services {
    flow-tap-dtcp {
      ssh;
    }
  }
}
  chassis {
    fpc 0 {
      pic 0 {
        tunnel-services {
          bandwidth 10g;
        }
      }
    }
  }
  interfaces {
    vt-0/0/0 {
      unit 0 {
        family inet;
        family inet6;
      }
    }
  }
  services {
    flow-tap {
      tunnel-interface vt-0/0/0.0;
    }
  }
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 53

Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements

The following sections explain each of the dynamic flow capture and flow-tap configuration statements. The statements are organized alphabetically.

address (Services Dynamic Flow Capture)

Syntax address *address*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Configure an IP address for the flow capture destination.

Options ***address***—IP address for the content destination.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

allowed-destinations

Syntax allowed-destinations [*identifiers*];

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Identify flow capture destinations that are allowed in messages sent from this control source.

Options ***identifier***—Allowed content destination name.

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

capture-group

Syntax

```
capture-group client-name {
  content-destination identifier {
    address address;
    hard-limit bandwidth;
    hard-limit-target bandwidth;
    soft-limit bandwidth;
    soft-limit-clear bandwidth;
    ttl hops;
  }
  control-source identifier {
    allowed-destinations [ destinations ];
    minimum-priority value;
    no-syslog;
    notification-targets address port port-number;
    service-port port-number;
    shared-key value;
    source-addresses [ addresses ];
  }
  duplicates-dropped-periodicity seconds;
  input-packet-rate-threshold rate;
  interfaces interface-name;
  max-duplicates number;
  pic-memory-threshold percentage percentage;
}
```

Hierarchy Level [edit services dynamic-flow-capture]

Release Information Statement introduced in Junos OS Release 7.4.

Description Define the capture group values.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring the Capture Group" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

content-destination

Syntax

```
content-destination identifier {
    address (Services Dynamic Flow Capture) address;
    hard-limit bandwidth;
    hard-limit-target bandwidth;
    soft-limit bandwidth;
    soft-limit-clear bandwidth;
    ttl hops;
}
```

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Identify the destination for captured packets.

Options ***identifier***—Name of the destination.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

control-source

Syntax

```
control-source identifier {
  allowed-destinations [ destinations ];
  minimum-priority value;
  no-syslog;
  notification-targets address port port-number;
  service-port port-number;
  shared-key value;
  source-addresses [ addresses ];
}
```

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Identify the source of the dynamic flow capture request.

Options *identifier*—Name of control source.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

duplicates-dropped-periodicity

Syntax duplicates-dropped-periodicity *seconds*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the **max-duplicates** threshold has been reached.

Options *seconds*—Period for sending DuplicatesDropped notifications.

Default: 30 seconds

Usage Guidelines See "Limiting the Number of Duplicates of a Packet" on page 1091.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- g-duplicates-dropped-periodicity on page 1110, max-duplicates on page 1114

Copyright © 2013, Juniper Networks, Inc.

dynamic-flow-capture

Syntax

```
dynamic-flow-capture {
  capture-group client-name {
    content-destination identifier {
      address address;
      hard-limit bandwidth;
      hard-limit-target bandwidth;
      soft-limit bandwidth;
      soft-limit-clear bandwidth;
      ttl hops;
    }
    control-source identifier {
      allowed-destinations [ destinations ];
      minimum-priority value;
      no-syslog;
      notification-targets address port port-number;
      service-port port-number;
      shared-key value;
      source-addresses [ addresses ];
    }
    duplicates-dropped-periodicity seconds;
    input-packet-rate-threshold rate;
    interfaces interface-name;
    max-duplicates number;
    pic-memory-threshold percentage percentage;
  }
  g-duplicates-dropped-periodicity seconds;
  g-max-duplicates number;
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 7.4.

Description Define the dynamic flow capture properties to be applied to traffic.

Options The remaining statements are explained separately.

Usage Guidelines See Dynamic Flow Capture.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

flow-tap

Syntax

```
flow-tap {
  (interface interface-name | tunnel-interface interface-name);
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 8.1.

Description Enable the flow-tap or FlowTapLite application on an interface. FlowTapLite is a lighter version of the flow-tap application that is available on MX Series platforms, M120 routers, and M320 routers with Enhanced III FPCs only.

Options

interface *interface-name*—Specify the interface name for the flow-tap application.

tunnel-interface *interface-name*—Specify the tunnel interface name for the FlowTapLite application.

The remaining statements are explained separately.

Usage Guidelines See Flow-Tap.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

g-duplicates-dropped-periodicity

Syntax g-duplicates-dropped-periodicity *seconds*;

Hierarchy Level [edit services dynamic-flow-capture]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the **g-max-duplicates** threshold has been reached. This setting is applied globally; the **duplicates-dropped-periodicity** setting applied at the **capture-group** level overrides the global setting.

Default The default period for sending notifications is 30 seconds.

Options ***seconds***—Period for sending DuplicatesDropped notifications.

Usage Guidelines See "Limiting the Number of Duplicates of a Packet" on page 1091.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- duplicates-dropped-periodicity on page 1107

 Copyright © 2013, Juniper Networks, Inc.

g-max-duplicates

Syntax g-max-duplicates *number*;

Hierarchy Level [edit services dynamic-flow-capture]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the maximum number of content destinations to which DFC PICs can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting is applied globally; the **max-duplicates** setting applied at the **capture-group** level overrides the global setting.

Default If no value is configured, a default setting of 3 is used.

Options *number*—Maximum number of content destinations.
Range: 1 through 64

Usage Guidelines See "Limiting the Number of Duplicates of a Packet" on page 1091.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- max-duplicates on page 1114

hard-limit

Syntax hard-limit *bandwidth*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a bandwidth threshold at which the dynamic flow capture application begins deleting criteria, until the bandwidth falls below the **hard-limit-target** value.

Options *bandwidth*—Hard limit threshold, in bits per second.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- hard-limit-target on page 1112

Copyright © 2013, Juniper Networks, Inc.

hard-limit-target

Syntax hard-limit-target *bandwidth*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a bandwidth threshold at which the dynamic flow capture application stops deleting criteria.

Options ***bandwidth***—Target value, in bits per second.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- hard-limit on page 1111

input-packet-rate-threshold

Syntax input-packet-rate-threshold *rate*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Specify a packet rate threshold value that triggers a system log warning message.

Options ***rate***—Threshold value.

Usage Guidelines See "Configuring Thresholds" on page 1090.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

interface (Services Flow Tap)

Syntax interface sp-*fpc/pic/port.logical-unit-number*;

Hierarchy Level [edit services flow-tap]

Release Information Statement introduced in Junos OS Release 8.1.

Description Specify the AS PIC interface used with the flow-tap application. Any AS PIC available in the router can be assigned, and any logical interface on the AS PIC can be used.

Options ***interface-name***—Name of the DFC interface.

You cannot configure flow-tap services on channelized interfaces.

Usage Guidelines See "Configuring the Flow-Tap Interface" on page 1097.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interfaces (Services Dynamic Flow Capture)

Syntax interfaces *interface-name*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Specify the DFC interface used with the control source configured in the same capture group.

Options ***interface-name***—Name of the DFC interface.

Usage Guidelines See "Configuring the DFC PIC Interface" on page 1088.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

max-duplicates

Syntax max-duplicates *number*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the maximum number of content destinations to which the DFC PIC can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting overrides the globally applied **g-max-duplicates** setting.

Default If no value is configured, a default setting of 3 is used.

Options ***number***—Maximum number of content destinations.
Range: 1 through 64

Usage Guidelines See "Limiting the Number of Duplicates of a Packet" on page 1091.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- g-max-duplicates on page 1111

minimum-priority

Syntax minimum-priority *value*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the minimum priority for the control source.

Options ***value***—Minimum priority value; if not specified, defaults to 0.
Range: 0 through 254

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

no-syslog

Syntax no-syslog;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Disable system logging of control protocol requests and responses. By default, these messages are logged.

Usage Guidelines See "Configuring System Logging" on page 1090.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

notification-targets

Syntax notification-targets *address* port *port-number*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description List of destination IP addresses and User Datagram Protocol (UDP) ports to which DFC PICs log exception information and control protocol state transitions, such as timeout values.

Options **address *address***—Allowed destination IP address.

port *port-number*—Allowed destination UDP port number.

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

pic-memory-threshold

Syntax pic-memory-threshold percentage *percentage*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Specify a PIC memory usage percentage that triggers a system log warning message.

Options **percentage *percentage***—PIC memory threshold value.

Usage Guidelines See "Configuring Thresholds" on page 1090.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

service-port

Syntax service-port *port-number*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Identify the User Datagram Protocol (UDP) port number for control protocol requests.

Options ***port-number***—Port number for control protocol request messages.

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

services (Dynamic Flow Capture)

Syntax services dynamic-flow-capture { ... },
services flow-tap {...}

Hierarchy Level [edit]

Release Information **dynamic-flow-capture** statement introduced in Junos OS Release 7.4.
flow-tap statement introduced in Junos OS Release 8.1.

Description Define the services to be applied to traffic.

Options **dynamic-flow-capture**—The values configured for dynamic flow capture.

flow-tap—The values configured for the flow-tap application.

The statements are explained separately.

Usage Guidelines See "Configuring Dynamic Flow Capture" on page 1085 or "Configuring the Flow-Tap Service" on page 1097.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

shared-key

Syntax shared-key *value*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Configure the authentication key value.

Options *value*—Secret authentication value shared between a control source and destination.

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

soft-limit

Syntax soft-limit *bandwidth*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a bandwidth threshold at which congestion notifications are sent to each control source of the criteria that point to this content destination. If the control source is configured with the **syslog** statement, a log message will also be generated.

Options ***bandwidth***—Soft limit threshold, in bits per second.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

soft-limit-clear

Syntax soft-limit-clear *bandwidth*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a bandwidth threshold at which the latch set by the soft-limit threshold is cleared.

Options ***bandwidth***—Soft-limit clear threshold, in bits per second.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- soft-limit on page 1118

 Copyright © 2013, Juniper Networks, Inc.

source-addresses

Syntax source-addresses [*addresses*];

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* control-source *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description List of IP addresses from which the control source can send control protocol requests to the Juniper Networks router.

Options ***address***—Allowed IP source address.

Usage Guidelines See "Configuring the Control Source" on page 1087.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Dynamic Flow Capture)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | non-world-readable>;
}
```

Hierarchy Level [edit services dynamic-flow-capture]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable and define tracing options for dynamic flow capture events.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number for files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files.

Default: 10 files.

no-world-readable—(Optional) Restrict access to the file.

world-readable—(Optional) Enable free access to the file.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Flow Capture on page 1085

ttl

Syntax ttl *hops*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Time-to-live (TTL) value for the IP-IP header.

Options ***hops***—TTL value.

Usage Guidelines See "Configuring the Content Destination" on page 1086.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

PART 6

Link and Multilink Services

- Link and Multilink Services Overview on page 1123
- Link and Multilink Services Configuration Guidelines on page 1127
- Summary of Multilink and Link Services Configuration Statements on page 1189

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 54

Link and Multilink Services Overview

This chapter discusses the following topics:

- Link and Multilink Services Overview on page 1123

Link and Multilink Services Overview

Multilink-based protocols enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of a multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members.

The Juniper Networks Junos operating system (Junos OS) supports several multilink-based protocols (such as MLPPP, FRF.15, and FRF.16) on the services PICs such as the Multilink Services PIC, the Link Services PIC, and the link services intelligent queuing (IQ) and voice services configured on the Adaptive Services (AS) and MultiServices PICs. For more information about link services IQ, see "Layer 2 Service Package Capabilities and Interfaces" on page 348. For more information about voice services, see "Configuring Services Interfaces for Voice Services" on page 422.

Starting with Junos OS Release 12.1, the following channelized MICs on MX240, MX480, and MX960 routers support Multilink Point-to-Point Protocol (MLPPP)-based services:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

. For more information about Multilink Point-to-Point Protocol (MLPPP)-based services MICs, see "Multilink Interfaces on Channelized MICs Overview" on page 1131.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The ml- interface type is used to configure interfaces on the Multilink Services PIC and does not support class-of-service (CoS) features. The ls- interface type is used for limited CoS configurations on the Link Services PIC (except on J Series Services Routers), and the lsq- interface type is used for full CoS configurations on the Adaptive Services and MultiServices PICs. The bundle interfaces are configured on the Multiservices DPC as link services IQ (lsq) interfaces and virtual LSQ redundancy (rlsq) interfaces.

For link services IQ (lsq) interfaces, Junos OS CoS components are fully supported and are handled normally on M Series and T Series routers, as described in the Junos OS Class of Service Configuration Guide. There are some restrictions on J Series Services Routers; for more information on link services IQ configuration, see "Layer 2 Service Package Capabilities and Interfaces" on page 348.

The Link Services and Multilink Services PICs support the following encapsulation types:

- Multilink Point-to-Point Protocol (MLPPP)
- Multilink Frame Relay (MLFR)

Starting with Junos OS Release 12.1, support for the following encapsulation types and protocols has been extended to the MX240, MX480, and MX960 routers with Multiservices DPCs:

- Multilink Point-to-Point Protocol (MLPPP)
- Multiclass MLPPP
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16) (also referred to as MFR)
- Compressed Real-Time Transport Protocol (CRTP)

MLPPP enables you to bundle multiple PPP links into a single logical link. MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single logical link. MLPPP and MLFR provide service option granularity between low-speed T1 and E1 services and higher-speed T3 and E3 services. You use MLPPP and MLFR to increase bandwidth in smaller, more cost-effective increments. In addition to providing incremental bandwidth, bundling multiple links can add a level of fault tolerance to your dedicated access service, because you can implement bundling across multiple PICs, protecting against the failure of any single PIC.

NOTE: Even if the PIC can support up to 4xDS3 total throughput, each aggregate can only run a volume of traffic equal to one DS3 in bandwidth. Aggregating DS3 links is not supported.

Copyright © 2013, Juniper Networks, Inc.

At the logical unit level, the Multilink Services and Link Services PICs support the MLPPP and MLFR Frame Relay Forum (FRF) 15 encapsulation types. At the physical interface level, the Link Services PIC also supports the MLFR FRF.16 encapsulation type.

MLPPP and MLFR FRF.15 are supported on interface types **ml-*fpc/pic/port***, **ls-*fpc/pic/port***, and **lsq-*fpc/pic/port***. For MLFR FRF.15, multiple permanent virtual circuits (PVCs) are combined into one aggregated virtual circuit (AVC). This provides fragmentation over multiple PVCs on one end and reassembly of the AVC on the other end.

MLFR FRF.16 is supported on a channelized interface, **ls-*fpc/pic/port*:*channel***, which denotes a single MLFR FRF.16 bundle. For MLFR FRF.16, multiple links are combined to form one logical link. Packet fragmentation and reassembly occur on a per-VC basis. Each bundle can support multiple VCs. Link Services PICs can support up to 256 DLCIs per MLFR FRF.16 bundle. The physical connections must be E1, T1, channelized DS3-to-DS1, channelized DS3-to-DS0, channelized E1, channelized STM1, or channelized IQ interfaces. When you bundle channelized interfaces using the link services interface, the channelized interfaces require M Series Enhanced Flexible PIC Concentrators (FPCs).



NOTE: When running MLPPP or MLFR on a non-QPP interface, you cannot mix logical units that are members of an aggregate with logical units configured using other families, such as inet. For example, the following configuration is not valid:

```
interface e3-0/0/0 {
    encapsulation frame-relay;
    unit 99 {
        dlci 99;
        family mlfr-end-to-end {
            bundle ls-0/0/0.1;
        }
    }
    unit 100 { ## mixes mlfr with family inet
        dlci 100;
        family inet {
            address 192.168.164.53/30;
        }
    }
}
```

The standards for MLPPP, MLFR FRF.15, and MLFR FRF.16 are defined in the following specifications:

- RFC 1990, *The PPP Multilink Protocol (MP)*
- FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*
- FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

Copyright © 2013, Juniper Networks, Inc.



NOTE: Endpoint Discriminator Class compatibility checking is enabled on MLPPP interfaces. Prior to Junos OS Release 8.0, when a Juniper Networks router received an unsupported Endpoint Discriminator Class message from an MLPPP session peer, it returned an ACK response.

Related Documentation

- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links on page 1133
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 55

Link and Multilink Services Configuration Guidelines

To configure multilink and link services logical interfaces, include the following statements.

```
(ml-fpc/pic/port | ls-fpc/pic/port) {
  unit logical-unit-number {
    dlci dlci-identifier;
    drop-timeout milliseconds;
    encapsulation type;
    fragment-threshold bytes;
    interleave-fragments;
    minimum-links number;
    mrru bytes;
    multicast-dlci dlci-identifier;
    short-sequence;
    family family {
      address address {
        destination address;
      }
      bundle (ml-fpc/pic/port | ls-fpc/pic/port);
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

To configure link services physical interfaces, include the **mlfr-uni-nni-bundle-options** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel*]** hierarchy level:

```
[edit interfaces ls-fpc/pic/port:channel]
encapsulation type;
mlfr-uni-nni-bundle-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  hello-timer milliseconds;
```

Copyright © 2013, Juniper Networks, Inc.

```
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 number;
    t392 number;
    yellow-differential-delay milliseconds;
}
```

This chapter contains the following sections:

- Multilink and Link Services PICs Overview on page 1129
- Multilink Interfaces on Channelized MICs Overview on page 1131
- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links on page 1133
- Configuring the Number of Bundles on Link Services PICs on page 1134
- Configuring the Links in a Multilink or Link Services Bundle on page 1135
- Multilink and Link Services Logical Interface Configuration Overview on page 1136
- Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 1138
- Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 1139
- Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 1140
- Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 1141
- Configuring MRRU on Multilink and Link Services Logical Interfaces on page 1141
- Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 1142
- Configuring DLCIs on Link Services Logical Interfaces on page 1143
- Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 1144
- Configuring Link Services Physical Interfaces on page 1147
- Configuring CoS on Link Services Interfaces on page 1151
- Examples: Configuring Multilink Interfaces on page 1156
- Examples: Configuring Link Interfaces on page 1159
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175
- Example: Configuring Link Interfaces on Channelized MICs on page 1179

Copyright © 2013, Juniper Networks, Inc.

Multilink and Link Services PICs Overview

Each Multilink Services or Link Services PIC can support a number of *bundles*. A bundle can contain up to eight individual *links*.

For Multilink Services PICs, the links can be T1, E1, or DS0 physical interfaces, and each link is associated with a logical unit number that you configure. For Link Services PICs, the links can be E1, T1, channelized DS3-to-DS1, channelized DS3-to-DS0, channelized E1, channelized STM1 interfaces, or channelized IQ interfaces. For MLFR FRF.16 bundles, each link is associated with a channel number that you configure.

You must configure a link before it can join a bundle. Each bundle should consist solely of one type of link; the mixing of physical interfaces of differing speeds within a bundle is not supported.



NOTE: **On both Juniper Networks J Series Services Routers and M Series Multiservice Edge Routers, only one DS3 link is allowed in an MLFR bundle. MLPPP bundles can include two DS3 links.**

Three versions of Multilink Services and three versions of Link Services PICs are available, as shown in Table 20 on page 1129. The PIC hardware is identical, except for different faceplates that enable you to identify which version you are installing. The software limits the unit numbers and maximum number of physical interfaces you assign to the PIC.

Table 20: Multilink and Link Services PIC Capacities

PIC Capacity	Unit Numbers	Maximum Number of T1/DS0 Interfaces	Maximum Number of E1 Interfaces
4-bundle PIC	0 through 3	32 links	32 links
32-bundle PIC	0 through 31	256 links	219 links
128-bundle PIC	0 through 127	292 links	219 links

A single PIC can support an aggregate bandwidth of 450 megabits per second (Mbps).

You can configure a larger number of links, but the Multilink Services and Link Services PICs can reliably process only 450 Mbps of traffic. A higher rate of traffic might degrade performance.



NOTE: **In Junos OS releases 9.0 and above you are not allowed to configure a unit number greater than the maximum unit number available on your link services PIC. Attempting to do so will cause an error message.**

Copyright © 2013, Juniper Networks, Inc.

For configuration information, see the following sections:

- Configuring the Number of Bundles on Link Services PICs on page 1134
- Configuring the Links in a Multilink or Link Services Bundle on page 1135

Copyright © 2013, Juniper Networks, Inc.

Multilink Interfaces on Channelized MICs Overview

Copyright © 2013, Juniper Networks, Inc.

Multiservices Modular Interface Cards (MICs) enable you to perform multiple services on the same MIC by configuring a set of services and applications such as voice services and Layer 2 Tunneling Protocol (L2TP) services. On Juniper Networks MX Series 3D Universal Edge Routers, the Multiservices DPC provides essentially the same capabilities as the Multiservices PIC. The interfaces on both platforms are configured in the same way. The Multilink interfaces are hosted on a channelized MIC. The bundle interfaces are configured on Multiservices DPC as virtual LSQ redundancy (rlsq) interfaces.

Starting with Junos OS Release 12.1, the following channelized MICs on MX240, MX480, and MX960 routers support Multilink Point-to-Point Protocol (MLPPP)-based services:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

The following encapsulations, interfaces, protocol, and packet types are supported on the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)—Supports Priority-based Flow Control (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and Multiclass MLPPP are supported for both data and control packets.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Multilink Frame Relay (MFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Link fragmentation and interleaving (LFI) non multilink MLPPP and MLFR packets.

Layer 2 services and voice services functionality are implemented on the Multiservices Dense Port Concentrators which supports the following two kinds of traffic that are routed by the Packet Forwarding Engine:

- Customer-end to provider-end (also, known as customer traffic)—Here, the Multilink fragments from the customer end arrive at the Multiservices interfaces configured on the channelized MIC. These fragments are then transmitted to the Multiservices DPC for Layer 2 processing such as CoS and are reassembled by the Multiservices software running on the Multiservices DPC. These reassembled packets are sent to the Packet Forwarding Engine where they go through the regular router lookup process and are finally sent over the Internet to the provider end. The voice packets also go through the same process.
- Provider-end to customer-end (also, known as Internet traffic)—Here, the data packets that are sent from the Internet provider end are received at any generic ingress interface in the Packet Forwarding Engine. These packets are then sent to the Multiservices DPC for Layer 2 processing. The Multiservices software running on Multiservices DPC

 Copyright © 2013, Juniper Networks, Inc.

fragment these data packets and send it to the Packet Forwarding Engine. These Multilink fragments are sent over the channelized MIC interfaces to the customer end. The voice packets also go through the same process.



NOTE: All the features that are supported on Multilink and Link Services PICs are also supported on the Multilink Services or Link Services MICs. For more information about Multilink and Link Services PICs, see "Multilink and Link Services PICs Overview" on page 1129.

Support for the following encapsulations, interfaces, protocol, and packet types are now extended to the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)—Supports priority-based flow control (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and multiclass MLPPP are supported for both data and control packets.
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI) and Consortium LMI (C-LMI) for data and control packets.
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
- Link fragmentation and interleaving (LFI) on multilink MLPPP and MLFR packets—Reduces delay and jitter on links by breaking up large data packets and interleaving delay-sensitive voice packets with the resulting smaller packets.

Related Documentation

- Link and Multilink Services Overview on page 1123
- Example: Configuring Link Interfaces on Channelized MICs on page 1179
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175

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

MX240, MX480, and MX960 3D Universal Edge Routers support MLPPP and MLFR multilink encapsulations. MLPPP enables you to bundle multiple PPP links into a single multilink bundle, and MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single multitlink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

You configure multilink bundles as logical units or channels on the link services interface **lsq-0/0/0**:

Copyright © 2013, Juniper Networks, Inc.

- With MLPPP and MLFR FRF.15, multilink bundles are configured as logical units on **lsq-0/0/0**—for example, **lsq-0/0/0.0** and **lsq-0/0/0.1**.
- With MLFR FRF.16, multilink bundles are configured as channels on **lsq-0/0/0**—for example, **lsq-0/0/0:0** and **lsq-0/0/0:1**.

After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated. You can create 64 multilink bundles, and on each multilink bundle you can add up to 8 constituent links. The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.
- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle, and only interfaces with a Frame Relay encapsulation can be added to an MLFR bundle.
- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

Configuring a multilink bundle on the two serial links increases the bandwidth by 70 percent from approximately 1 Mbps to 1.7 Mbps and prepends each packet with a multilink header as specified in the FRF.12 standard. To increase the bandwidth further, you can add up to eight serial links to the bundle. In addition to a higher bandwidth, configuring the multilink bundle provides load balancing and redundancy. If one of the serial links fails, traffic continues to be transmitted on the other links without any interruption. In contrast, independent links require routing policies for load balancing and redundancy. Independent links also require IP addresses for each link as opposed to one IP address for the bundle. In the routing table, the multilink bundle is represented as a single interface.

Related Documentation

- Link and Multilink Services Overview on page 1123
- Multilink Interfaces on Channelized MICs Overview on page 1131
- Example: Configuring Link Interfaces on Channelized MICs on page 1179
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175

Configuring the Number of Bundles on Link Services PICs

You can combine MLFR FRF.16, MLPPP, and MLFR FRF.15 bundles on a single Link Services PIC. For a sample configuration, see "Example: Configuring a Link Services Interface with Two Links" on page 1160.

To configure the number of bundles on a Link Services PIC, include the **mlfr-uni-nni-bundles** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

```
mlfr-uni-nni-bundles number;
```


Copyright © 2013, Juniper Networks, Inc.

Each Link Services PIC can accommodate a maximum of 256 MLFR UNI NNI bundles. For more information, see the Junos OS System Basics Configuration Guide.

A link can associate with one link services bundle only. All Link Services PICs support up to 256 single-link bundles and up to 256 DLCIs. For an example configuration, see the configuration examples.

NOTE: When one or more links in a bundle are put in loopback, reassembly buffering and hence processing are reduced so as to not affect other bundles. This prevents packet loss on other bundles, while reducing the reassembly buffers available for the bundle with looped links.

Related Documentation

- Example: Configuring a Link Services Interface with Two Links on page 1160
- Example: Configuring a Link Services Interface with MLPPP on page 1161
- Example: Configuring a Link Services Interface with MLFR FRF.15 on page 1162
- Example: Configuring a Link Services PIC with MLFR FRF.16 on page 1162
- Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types on page 1163

Configuring the Links in a Multilink or Link Services Bundle

To complete a multilink or link services interface configuration, you need to configure both the physical interface and the multilink or link services bundle. For multilink interfaces, you configure the link bundle on the logical unit. For link services interfaces, you configure the link bundle as a channel (see Figure 15 on page 1135). The physical interface is usually connected to networks capable of supporting MLPPP or MLFR (FRF.15 or FRF.16).

Figure 15: Multilink Interface Configuration

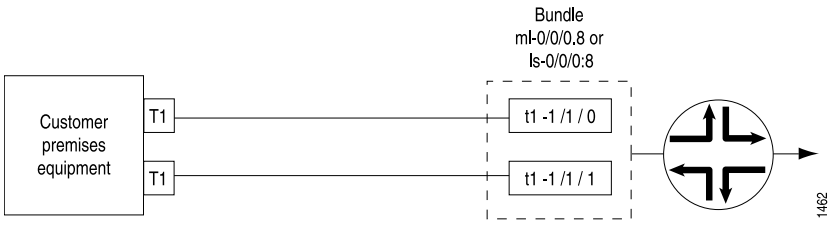

The following sample configuration refers to the topology in Figure 15 on page 1135 and configures a multilink or link services bundle over a T1 connection (for which the T1 physical interface is already configured).

1. To configure a physical T1 link for MLPPP, include the following statements at the **[edit interfaces t1-*fpc/pic/port*]** hierarchy level:

```
unit 0 {
  family mlppp {
    bundle (ml-fpc/pic/port | ls-fpc/pic/port);
  }
```

```
}
```

You do not need to configure an IP address on this link.

To configure a physical T1 link for MLFR FRF.16, include the following statements at the **[edit interfaces t1-*fpc/pic/port*]** hierarchy level:

```
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
  family mlfr-uni-nni {
    bundle ls-fpc/pic/port:channel;
  }
}
```

You do not need to configure an IP address or a DLCI on this link.

2. To configure the logical address for the MLPPP, MLFR FRF.15, or MLFR FRF.16 bundle, include the **address** and **destination** statements:

```
address address {
  destination address;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family inet]**

When you add statements such as **mrru** to the configuration and commit, the T1 interface becomes part of the multilink bundle.



NOTE: For MLPPP and MLFR (FRF.15 and FRF.16) links, you must specify the subnet address as /32 or /30. Any other subnet designation is treated as a mismatch.

Multilink and Link Services Logical Interface Configuration Overview

You configure multilink and link services interface properties at the logical unit level. Default settings for multilink and link services logical interface properties are described in "Default Settings for Multilink and Link Services Logical Interfaces" on page 1137.

For general information about logical unit properties or **family inet** properties, see the Junos® OS Network Interfaces. For information about multilink and link services properties you configure at the **family inet** hierarchy level, see "Configuring the Links in a Multilink or Link Services Bundle" on page 1135.

Copyright © 2013, Juniper Networks, Inc.

NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the router does not use the bandwidth value if the interfaces are included in an MLPPP or MLFR bundle. The bandwidth is calculated internally according to the time slots, framing, and byte-encoding of the interface. For more information about logical interface properties, see the Junos® OS Network Interfaces.

Default Settings for Multilink and Link Services Logical Interfaces

Table 21 on page 1137 lists the default settings for multilink and link services statements, together with the other permitted values or value ranges.

Table 21: Multilink and Link Services Logical Interface Statements

Option	Default Value	Possible Values
DLCI	None	16 through 1022
Drop timeout period	500 ms for bundles greater than or equal to the T1 bandwidth value and 1500 ms for other bundles.	0 through 2000 milliseconds
Encapsulation	For multilink interfaces, **multilink-ppp**. For link services interfaces, **multilink-frame-relay-end-to-end**.	**multilink-frame-relay-end-to-end**, **multilink-ppp**
Fragmentation threshold	0 bytes	128 through 16,320 bytes (*N*x64)
Interleave fragments	disabled	enabled, disabled
Minimum links	1 link	1 through 8 links
Maximum received reconstructed unit (MRRU)	1504 bytes	1500 through 4500 bytes
Sequence ID format for MLPPP	24 bits	12 or 24 bits
Sequence ID format for MLFR FRF.15 and FRF.16	12 bits	12 bits

See "Default Settings for Link Services Interfaces" on page 1147 for statements that apply to link services physical interfaces only.

Related Documentation

- Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 1138
- Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 1139

Copyright © 2013, Juniper Networks, Inc.

- Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 1140
- Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 1141
- Configuring MRRU on Multilink and Link Services Logical Interfaces on page 1141
- Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 1142
- Configuring DLCIs on Link Services Logical Interfaces on page 1143
- Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 1144

Configuring Encapsulation for Multilink and Link Services Logical Interfaces

Multilink and link services interfaces support the following logical interface encapsulation types:

- MLPPP
- MLFR end-to-end

By default, the logical interface encapsulation type on multilink interfaces is MLPPP. The default logical interface encapsulation type on link services interfaces is MLFR end-to-end. For general information on encapsulation, see the Junos® OS Network Interfaces.

You can also configure physical interface encapsulation on link services interfaces. For more information, see "Configuring Encapsulation for Link Services Physical Interfaces" on page 1148.

To configure multilink or link services encapsulation, include the **encapsulation** statement:

```
encapsulation type;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

You must also configure the T1, E1, or DS0 physical interface with the same encapsulation type.



CAUTION: When you configure the first MLFR encapsulated unit or delete the last MLFR encapsulated unit on a port, it triggers an interface encapsulation change on the port, which causes an interface flap on the other units within the port that are configured with generic Frame Relay.

Copyright © 2013, Juniper Networks, Inc.

Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces

By default, the drop timeout parameter is disabled. You can configure a drop timeout value to provide a recovery mechanism if individual links in the multilink or link services bundle drop one or more packets. Drop timeout is not a differential delay tolerance setting, and does not limit the overall latency. However, you need to make sure the value you set is larger than the expected differential delay across the links, so that the timeout period does not elapse under normal jitter conditions, but only when there is actual packet loss. You can configure differential delay tolerance for link services interfaces only. For more information, see "Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16" on page 1149.

To configure the drop timeout value, include the **drop-timeout** statement:

```
drop-timeout milliseconds;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

For link services interfaces, you also can configure the drop timeout value at the physical interface level by including the **drop-timeout** statement at the **[edit interfaces ls-***fpc/pic/port***:***channel* **mlfr-uni-nni-bundle-options]** hierarchy level:

```
drop-timeout milliseconds;
```

By default, the drop timer has a value of 500 ms for bundles greater than or equal to the T1 bandwidth value, and 1500 ms for other bundles. Any CLI-configured value overrides these defaults. Values can range from 1 through 2000 milliseconds. Values less than 5 milliseconds are not recommended, and a configured value of 0 reverts to the default value of 2000 milliseconds.



NOTE: For multilink or link services interfaces, if a packet or fragment encounters an error condition and is destined for a disabled bundle or link, it does not contribute to the dropped packet and frame counts in the per-bundle statistics. The packet is counted under the global error statistics and is not included in the global output bytes and output packet counts. This unusual accounting happens only if the error conditions are generated inside the multilink interface, not if the packet encounters errors on the wire or elsewhere in the network.

If you configure the **drop-timeout** statement with a value of **0**, it disables any resequencing by the PIC for the specified class of MLPPP traffic. Packets are forwarded with the assumption that they arrived in sequence, and forwarding of fragmented packets is disabled for all classes. Fragments dropped as a result of this setting will increment the counter at the class level.

Copyright © 2013, Juniper Networks, Inc.

Alternatively, you can configure the **drop-timeout** statement at the **[edit class-of-service fragmentation-maps** *map-name* **forwarding-class** *class***]** hierarchy level. The behavior and the default and range values are identical, but the setting applies only to the specified forwarding class. Configuration at the bundle level overrides configuration at the class-of-service level.

By default, compression of the inner PPP header in the MLPPP payload is enabled. To disable compression, include the **disable-mlppp-inner-ppp-pfc** statement at the **[edit interfaces** *interface-name* **unit** *logical-unit-number***]** hierarchy level. For example:

```
interfaces lsq-1/2/0 {
  unit 0 {
    encapsulation multilink-ppp;
    disable-mlppp-inner-ppp-pfc;
    multilink-max-classes 4;
    family inet {
      address 10.50.1.2/30;
    }
  }
}
```

For more information about CoS configuration, see the Junos OS Class of Service Configuration Guide. You can view the configured drop-timeout value and the status of inner PPP header compression by issuing the **show interfaces** *interface-name* **extensive** command.

Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces

For multilink and link services logical interfaces with MLPPP encapsulation only, you can configure a *fragmentation threshold* to limit the size of packet payloads transmitted across the individual links within the multilink circuit. The software splits any incoming packet that exceeds the fragmentation threshold into smaller units suitable for the circuit size; it reassembles the fragments at the other end, but does not affect the output traffic stream. The threshold value affects the payload only; it does not affect the MLPPP header. By default, the fragmentation threshold parameter is disabled.



NOTE: To ensure proper load balancing:

- For Link Services MLFR (FRF.15 and FRF.16) interfaces, do not include the **fragment-threshold** statement in the configuration.
- For MLPPP interfaces, do not include both the **fragment-threshold** statement and the **short-sequence** statement in the configuration.
- For MLFR (FRF.15 and FRF.16) and MLPPP interfaces, if the MTU of links in a bundle is less than the bundle MTU plus encapsulation overhead, then fragmentation is automatically enabled. You should avoid this situation for MLFR (FRF.15 and FRF.16) interfaces and for MLPPP interfaces on which short-sequencing is enabled.

To configure a fragmentation threshold value, include the **fragment-threshold** statement:

Copyright © 2013, Juniper Networks, Inc.

fragment-threshold *bytes*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

For link services interfaces, you also can configure a fragmentation threshold value at the physical interface level by including the **fragment-threshold** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

fragment-threshold *bytes*;

The maximum fragment size can be from 128 through 16,320 bytes. The Junos OS automatically subdivides packet payloads that exceed this value. Any value you set must be a multiple of 64 bytes (*N*x64). The default value, 0, results in no fragmentation.

Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces

You can set the minimum number of links that must be up for the multilink bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up. A member link is considered up when the PPP Link Control Protocol (LCP) phase transitions to open state.

The **minimum-links** value should be identical on both ends of the bundle.

To set the minimum number, include the **minimum-links** statement:

minimum-links *number*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

For link services interfaces, you also can configure the minimum number of links at the physical interface level by including the **minimum-links** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

minimum-links *number*;

The number can be from 1 through 8. The maximum number of links supported in a bundle is 8. When 8 is specified, all configured links of a bundle must be up.

Configuring MRRU on Multilink and Link Services Logical Interfaces

The *maximum received reconstructed unit* (*MRRU*) is similar to a maximum transmission unit (MTU), but applies only to multilink bundles; it is the maximum packet size that the multilink interface can process. By default, the MRRU is set to 1500 bytes; you can configure a different MRRU value if the peer equipment allows this. The MRRU accounts

Copyright © 2013, Juniper Networks, Inc.

for the original payload, for example the Layer 3 protocol payload, but does not include the 2-byte PPP header or the additional MLPPP or MLFR header applied while the individual multilink packets are traversing separate links in the bundle.

To configure a different MRRU value, include the **mrru** statement:

mrru *bytes*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

For link services interfaces, you also can configure a different MRRU at the physical interface level by including the **mrru** statement at the **[edit interfaces ls-***fpc/pic/port***:***channel* **mlfr-uni-nni-bundle-options]** hierarchy level:

mrru *bytes*;

The MRRU size can range from 1500 through 4500 bytes.



NOTE: If you set the MRRU on a bundle to a value larger than the MTU of the individual links within it, you must enable a fragmentation threshold for that bundle. Set the threshold to a value no larger than the smallest MTU of any link included in the bundle.

Determine the appropriate MTU size for the bundle by ensuring that the MTU size does not exceed the sum of the encapsulation overhead and the MTU sizes for the links in the bundle.

You can configure separate **family mtu** values on the following protocol families under bundle interfaces: **inet**, **inet6**, **iso**, and **mpls**. If not configured, the default value of 1500 is used on all except for **mpls** configurations, in which the value 1488 is used.



NOTE: The effective family MTU might be different from the MTU value specified for MLPPP configurations, because it is adjusted downward by the remote MRRU's constraints. The remote MRRU configuration is not supported on M120 routers.

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces

For MLPPP, the sequence header format is set to 24 bits by default. You can configure an alternative value of 12 bits, but 24 bits is considered the more robust value for most networks.

To configure a different sequence header value, include the **short-sequence** statement:

Copyright © 2013, Juniper Networks, Inc.

short-sequence;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

For MLFR FRF.15, the sequence header format is set to 24 bits by default. This is the only valid option.

Configuring DLCIs on Link Services Logical Interfaces

For link services interfaces only, you can configure multiple DLCIs for MLFR FRF.16 or MLPPP bundles.

DLCIs are not supported on multilink interfaces.

Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles

For link services interfaces only, you can configure multiple point-to-point DLCIs for each MLFR FRF.16 or MLPPP bundle. A channelized interface, such as **ls-1/1/1:0**, denotes a single MLFR FRF.16 bundle. To configure a DLCI, include the **dlci** statement:

dlci *dlci-identifier*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

The DLCI identifier is a value from 16 through 1022. Numbers 1 through 15 are reserved for future use.

When you configure point-to-point connections, the maximum transmission unit (MTU) sizes on both sides of the connection must be the same.

Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles

For link services interfaces only, you can configure multiple multicast-capable DLCIs for each MLFR FRF.16 bundle. A channelized interface, such as **ls-1/1/1:0**, denotes a single MLFR FRF.16 bundle. By default, Frame Relay connections assume unicast traffic. If your Frame Relay switch performs multicast replication, you can configure the link services connection to support multicast traffic by including the **multicast-dlci** statement:

multicast-dlci *dlci-identifier*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

Copyright © 2013, Juniper Networks, Inc.

The DLCI identifier is a value from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

You can configure multicast support only on point-to-multipoint link services connections. Multicast-capable DLCIs are not supported on multilink interfaces.

If keepalives are enabled, causing the interface to send Local Management Interface (LMI) messages during idle times, the number of possible DLCI configurations is limited by the MTU selected for the interface. For more information, see "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces

For link services FRF.15 and MLPPP interfaces only, you can configure link fragment interleaving (LFI). LFI reduces excessive delays of Frame Relay packets by fragmenting long packets into smaller packets and interleaving them with real-time frames. This allows real-time and non-real-time data frames to be carried together on lower-speed links without causing excessive delays to the real-time traffic. When the peer interface receives the smaller fragments, it reassembles the fragments into their original packet. For example, short delay-sensitive packets, such as packetized voice, can race ahead of larger delay-insensitive packets, such as common data packets.



NOTE: All Link Services PICs (4-multilink bundle, 32-multilink bundle, and 128-multilink bundle) support up to 256 link services interfaces with LFI enabled, if those link services interfaces contain only one constituent link each. For the Link Services PIC, multiple-link LFI bundles are simply multilink bundles, and are limited based on the type of PIC (4-multilink bundle, 32-multilink bundle, and 128-multilink bundle).

In addition, the multilink bundles you configure subtract from the total of 256 possible LFI-enabled link services interfaces. For example, if a 32-multilink bundle Link Services PIC has 24 multilink bundles configured and active, then you can configure 256 – 24 = 232 LFI-enabled link services interfaces, each with a single constituent link.

For link services IQ interfaces (lsq), the interleave-fragments statement is not valid. Instead, you can enable LFI by configuring fragmentation maps. For more information, see "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 365.

You can configure multiple links in a bundle and configure packet interleaving. However, if you use packet interleaving, high-priority, nonmultilink-encapsulated packets use a hash-based algorithm to choose a single link.

For detailed information about link services CoS, see "Configuring CoS on Link Services Interfaces" on page 1151.

Copyright © 2013, Juniper Networks, Inc.

Per-bundle CoS queuing is supported on link services IQ interfaces (**lsq**). For more information about link services IQ interfaces, see "Layer 2 Service Package Capabilities and Interfaces" on page 348.

The Junos OS supports end-to-end fragmentation in compliance with the FRF.12 *Frame Relay Fragmentation Implementation Agreement* standard. Unlike user-to-network interface (UNI) and network-to-network (NNI) fragmentation, end-to-end supports fragmentation only at the endpoints.

By default, packet interleaving is disabled. To enable packet interleaving, include the **interleave-fragments** statement:

```
interleave-fragments;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Configuring LFI with DLCI Scheduling

For Link Services and Channelized DS3 IQ PICs, you can configure LFI and DLCI scheduling. For channelized DS3 interfaces, LFI is supported with FRF.15 only, and on M10i and M20 platforms only.

Configuring LFI with DLCI scheduling enables packets entering the Link Services PIC to be fragmented before being transmitted to the Channelized DS3 IQ PIC. Once the fragmented packets enter the Channelized DS3 IQ PIC, they are scheduled at the DLCI level, to allow priority transmission for real-time applications.

For more information about associating a scheduler with a DLCI, see the Junos OS Class of Service Configuration Guide.

Example: Configuring LFI with DLCI Scheduling

Configure packets entering the Link Services PIC to be fragmented before being transmitted to the Channelized DS3 IQ PIC. Once the fragmented packets enter the Channelized DS3 IQ PIC, they are scheduled at the DLCI level, to allow priority transmission for real-time applications.

```
[edit interfaces]
ls-1/0/0 {
    unit 1 {
        encapsulation multilink-frame-relay-end-to-end;
        interleave-fragments;
        family inet {
            address 192.168.5.2/32 {
                destination 192.168.5.3;
            }
        }
    }
    t3-1/0/0:1 {
```

Copyright © 2013, Juniper Networks, Inc.

```
    per-unit-scheduler;
    unit 0 {
      dlci 16;
      encapsulation multilink-frame-relay-end-to-end;
      family mlfr-end-to-end {
        bundle ls-1/0/0.1;
      }
    }
}
[edit class-of-service]
interfaces {
    t3-1/0/0:1 {
      unit 0 {
        scheduler-map sched-map-logical-0;
        shaping-rate 10m;
      }
      unit 1 {
        scheduler-map sched-map-logical-1;
        shaping-rate 20m;
      }
    }
}
scheduler-maps {
    sched-map-logical-0 {
      forwarding-class best-effort scheduler sched-best-effort-0;
      forwarding-class assured-forwarding scheduler sched-bronze-0;
      forwarding-class expedited-forwarding scheduler sched-silver-0;
      forwarding-class network-control scheduler sched-gold-0;
    }
    sched-map-logical-1 {
      forwarding-class best-effort scheduler sched-best-effort-1;
      forwarding-class assured-forwarding scheduler sched-bronze-1;
      forwarding-class expedited-forwarding scheduler sched-silver-1;
      forwarding-class network-control scheduler sched-gold-1;
    }
    schedulers {
      sched-best-effort-0 {
        transmit-rate 4m;
      }
      sched-bronze-0 {
        transmit-rate 3m;
      }
      sched-silver-0 {
        transmit-rate 2m;
      }
      sched-gold-0 {
        transmit-rate 1m;
      }
      sched-best-effort-1 {
        transmit-rate 8m;
      }
      sched-bronze-1 {
        transmit-rate 6m;
      }
      sched-silver-1 {
        transmit-rate 4m;
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
            sched-gold-1 {
               transmit-rate 2m;
            }
         }
      }
   }
```

Configuring Link Services Physical Interfaces

You configure link services interface properties at the logical unit and physical interface level. Default settings for link services physical interface properties are described in "Default Settings for Link Services Interfaces" on page 1147.

The following sections explain how to configure link services physical interfaces:

- Default Settings for Link Services Interfaces on page 1147
- Configuring Encapsulation for Link Services Physical Interfaces on page 1148
- Configuring Acknowledgment Timers on Link Services Physical Interfaces on page 1148
- Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 on page 1149
- Configuring Keepalives on Link Services Physical Interfaces on page 1150

For information about link services physical interface properties that can also be configured at the logical unit level, see "Multilink and Link Services Logical Interface Configuration Overview" on page 1136.

Default Settings for Link Services Interfaces

Table 22 on page 1147 lists the default settings for link services statements, together with the other permitted values or value ranges.

Table 22: Link Services Physical Interface Statements for MLFR FRF.16

Option	Default Value	Possible Values
Action red differential delay	**remove-link**	**disable-tx**, **remove-link**
Red differential delay	120 ms	1 through 2000 ms
Yellow differential delay	72 ms	1 through 2000 ms
Drop timeout period	0 ms	0 through 2000 ms
Encapsulation	**multilink-frame-relay-uni-nni**	**multilink-frame-relay-uni-nni**
Fragmentation threshold	0 bytes	128 through 16,320 bytes (*N*x64)
LMI type	itu	ansi, itu
Minimum links	1 link	1 through 8 links

Copyright © 2013, Juniper Networks, Inc.

Table 22: Link Services Physical Interface Statements for MLFR FRF.16 *(continued)*

Option	Default Value	Possible Values
MRRU	1504 bytes	1500 through 4500 bytes
n391 (full status polling counter)	6	1 through 255
n392 (LMI error threshold)	3	1 through 10
n393 (LMI monitored event count)	4	1 through 10
t391 (link integrity verify polling timer)	10	5 through 30
t392 (polling verification timer)	15	5 through 30
Sequence ID format for MLFR	12 bits	12 bits

Configuring Encapsulation for Link Services Physical Interfaces

Link services interfaces support the physical interface encapsulation MLFR UNI NNI. By default, the physical interface encapsulation on link services interfaces is MLFR UNI NNI. Multilink interfaces do not support physical interface encapsulation.

For more information, see the Junos® OS Network Interfaces.

You can also configure logical interface encapsulation on multilink and link services interfaces. For more information, see "Configuring Encapsulation for Multilink and Link Services Logical Interfaces" on page 1138.

To explicitly configure link services physical interface encapsulation, include the **encapsulation** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel*]** hierarchy level:

```
encapsulation type;
```

You must also configure the T1, E1, or DS0 physical and physical interface with the same encapsulation type.

Configuring Acknowledgment Timers on Link Services Physical Interfaces

For link services interfaces configured with MLFR FRF.16, each link end point in a bundle initiates a request for bundle operation with its peer by transmitting an add link message. A hello message notifies the peer end point that the local end point is up. Both ends of a link generate a hello message periodically, or as configured with the hello timer. A remove link message notifies the peer that the local end management is removing the link from bundle operation. End points respond to add link, remove link, and hello messages by sending acknowledgment messages.

You can configure the maximum period to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment by including the **acknowledge-timer**

Copyright © 2013, Juniper Networks, Inc.

statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
acknowledge-timer milliseconds;
```

The acknowledgment timer can be from 1 through 10 milliseconds. The default is 4 milliseconds.

For link services interfaces, you can configure the number of retransmission attempts to be made for consecutive hello or remove link messages after the expiration of the acknowledgment timer by including the **acknowledge-retries** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
acknowledge-retries number;
```

acknowledgment-retries can be a value from 1 through 5. The default is 2.

You can configure the rate at which hello messages are sent by including the **hello-timer** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
hello-timer milliseconds;
```

A hello message is transmitted after the specified period (in milliseconds) has elapsed. The hello timer can be from 1 through 180 milliseconds; the default is 10 milliseconds. When the hello timer expires, a link end point generates an add-link message.

Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16

For link services interfaces configured with MLFR FRF.16, the differential delay between links in a bundle is measured and warning is given when a link has a substantially greater differential delay than other links in the same bundle. The implementing endpoint can determine if the differential delay is in an acceptable range and decide to remove the link from the bundle, or to stop transmission on the link.

You can configure the yellow differential delay for links in a bundle by including the **yellow-differential-delay** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
yellow-differential-delay milliseconds;
```

The yellow differential delay can be from 1 through 2000 milliseconds. The default is 72 milliseconds.

You can configure the red differential delay for links in a bundle to give warning by including the **red-differential-delay** statements at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
red-differential-delay milliseconds;
```

The red differential delay can be from 1 through 2000 milliseconds. The default is 120 milliseconds.

Copyright © 2013, Juniper Networks, Inc.

You can configure the action to be taken when differential delay exceeds the red limit by including the **action-red-differential-delay red** statements at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
action-red-differential-delay (disable-tx | remove-link);
```

The **disable-tx** option disables transmission on the link. The **remove-link** option removes the link from the bundle. The default action is **remove-link**.

You can view these settings in the output of the **show interfaces extensive lsq-*fpc/pic/port*:*channel*** command.

Configuring Keepalives on Link Services Physical Interfaces

You can tune the keepalive settings on the physical link-services interface. By default, the Junos OS uses ITU Q.933 Annex A LMIs for FRF.16. To instead use ITU Annex A LMIs (ANSI), include the **lmi-type ansi** statement at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level. LMI type ANSI is used in the following example:

```
lmi-type ansi;
```

To configure Frame Relay keepalive parameters on a link services interface, include the **n391**, **n392**, **n393**, **t391** and **t392** statements at the **[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]** hierarchy level:

```
[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]
n391 number;
n392 number;
n393 number;
t391 number;
t392 number;
```

The statements determine the indicated keepalive settings:

- **n391**—Full status polling interval. The data terminal equipment (DTE) sends a status inquiry to the data communication equipment (DCE) at the interval specified by the **t391** statement. This statements sets the frequency at which the DTE requests full status report; for example, the value **10** means that the DTE requests full status report in every tenth inquiry. The intermediate inquiries request a keepalive response only. The range is **1** through **255**, with a default of **6**.
- **n392**—Error threshold, which is the maximum number of errors that can occur during the number of events set by the **n393** statement before the link is marked inoperative. The range is **1** through **10**, with a default of **3**.
- **n393**—Monitored event count. The range is **1** through **10**, with a default of **4**.
- **t391**—The interval at which the DTE requests a keepalive response from the DCE and updates status, depending on the error threshold value. The range is 5 through 30 seconds, with a default of 10 seconds.
- **t392**—The period during which the DCE checks for keepalive responses from the DTE and updates status, depending on the DCE error threshold value. The range is from 5 through 30 seconds, with a default of 15 seconds.

Copyright © 2013, Juniper Networks, Inc.

NOTE: For the LMI to work properly, you must configure one side of a link services bundle to be a DCE.

Configuring CoS on Link Services Interfaces

For link services IQ (**lsq-**) interfaces, Junos class of service (CoS) is fully supported and functions as described in the Junos OS Class of Service Configuration Guide. For more information and detailed configuration examples, see "Layer 2 Service Package Capabilities and Interfaces" on page 348.

On SRX Series and J Series devices, the **lsq-** interface is an internal interface, which is not associated with a physical interface. For information about link services on SRX Series and J Series devices, see the *Junos OS Interfaces Configuration Guide for Security Devices*.

For information about CoS functions and link services on M Series or T Series routers, see the following sections:

- CoS for Link Services Interfaces on M Series and T Series Routers on page 1151
- Example: Configuring CoS on Link Services Interfaces on page 1152

CoS for Link Services Interfaces on M Series and T Series Routers

For Link Services PIC interfaces (**ls**) on M Series and T Series routers, queue 0 is the only queue that you should configure to receive fragmented packets. Configure all other queues to be higher-priority queues.

Table 23 on page 1151 summarizes how CoS queues work on link services (**ls**) interfaces.

Table 23: Link Services CoS Queues

Supported Bundling Type	Queue 0	Higher-Priority Queues
Hash-based load balancing	No	Yes
MLFR FRF.15	Yes	No
MLFR FRF.16	Yes	No
MLPPP	Yes	No

For M Series and T Series routers, CoS on link services (**ls**) interfaces works as follows:

- On all platforms, the Link Services PIC currently supports up to four queues: 0, 1, 2, and 3.
- Queue 0 uses MLFR FRF.15, MLFR FRF.16, or MLPPP to bundle packets.
- Higher-priority queues (1, 2, and 3) use hash-based load balancing to bundle packets. IP and MPLS header information is included in the hash.

Copyright © 2013, Juniper Networks, Inc.

- MLPPP packets traversing link services interfaces using queue 0 are fragmented and distributed across the constituent links. Queue 0 packets are sent on the least utilized link, proportional to its bandwidth. The queue 0 load balancer attempts to maintain even distribution of all traffic across all constituent links. In situations with a small number of high-priority traffic flows (queues 1, 2, and 3), queue 0 traffic might be unevenly distributed.
- For the MLFR FRF.16 protocol, only queue 0 works. If you configure a bundled interface to use MLFR FRF.16 with queue 0, then you must ensure the classifier does not send any traffic to queues 1, 2, and 3 on that interface.
- To carry high-priority traffic correctly on MLFR FRF.16 interfaces, you must configure an output firewall filter that forces all traffic into queue 0 on the ls-***fpc/pic/port.channel*** interface.
- MLFR FRF.15 and MLPPP interfaces support CoS through packet interleaving. The MLFR FRF.16 standard does not support packet interleaving, so all packets destined for an FRF.16 PVC interface must egress from the same queue.
- For constituent link interfaces of Link Services PICs, you can configure standard scheduler maps.
- For input packets and fragments received from constituent links, you can use regular input firewall filters and standard CoS classifiers on the link services interface.
- For packets that pass through a link services interface and are destined for a constituent link interface, all traffic using queue 0 is fragmented. Traffic using higher-priority queues (1, 2, and 3) is not fragmented.
- For MLFR FRF.15 and MLPPP, routing protocol packets smaller than 128 bytes are sent to queue 3; routing protocol packets that exceed 128 bytes are sent to queue 0 and fragmented accordingly. For MLFR FRF.16, queue 0 is used for all packet sizes.
- You must configure output firewall classification for egress traffic on the link services interface, not directly on the constituent link interface directly.
- Inverse multiplexing for ATM (IMA) is not supported on link services interfaces.

For more information, see "Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces" on page 1144 and the Routing Policy Configuration Guide.

Example: Configuring CoS on Link Services Interfaces

Configure CoS on a link services interface and its constituent link interfaces.



NOTE: This example applies to M Series and T Series routers. For examples that apply to SRX Series and J Series devices, see the *Junos OS Interfaces Configuration Guide for Security Devices*.

Packets that do not match the firewall filters are sent to a queue that performs load balancing by sending fragments to all constituent links.

Copyright © 2013, Juniper Networks, Inc.

Packets that match the firewall filters are sent to a queue that does not support packet fragmentation and reassembly; instead, this traffic is load-balanced by sending each packet flow to a different constituent link. Each packet that matches a firewall filter is subjected to a hash on the IP source address and the IP destination address to determine the packet flow to which each packet belongs.

When you configure the MLPPP encapsulation type or the multilink FRF.15 Frame Relay end-to-end encapsulation type, routing protocol packets smaller than 128 bytes are sent to the network-control queue on the constituent link interface. This keeps routing protocols operating normally, even when low-speed links are congested by regular packets.

```
[edit interfaces]
ls-7/0/0 {
   unit 0 {
      encapsulation multilink-ppp;
      interleave-fragments;
      family inet {
         filter {
            output lfi_ls_filter;
         }
         address 10.54.0.2/32 {
            destination 10.54.0.1;
         }
      }
   }
}
ge-7/2/0 {
   unit 0 {
      family inet {
         address 192.168.1.1/24;
      }
   }
}
ce1-7/3/6 {
   no-partition interface-type e1;
}
e1-7/3/6 {
   encapsulation ppp;
   unit 0 {
      family mlppp {
         bundle ls-7/0/0.0;
      }
   }
}
ce1-7/3/7 {
   no-partition interface-type e1;
}
e1-7/3/7 {
   encapsulation ppp;
   unit 0 {
      family mlppp {
         bundle ls-7/0/0.0;
      }
   }
}
```

Copyright © 2013, Juniper Networks, Inc.

```
[edit class-of-service]
classifiers {
   dscp dscp_default {
      import default;
   }
   inet-precedence inet-precedence_default {
      import default;
   }
}
code-point-aliases {
   dscp {
      af11 001010;
      af12 001100;
      af13 001110;
      af21 010010;
      af22 010100;
      af23 010110;
      af31 011010;
      af32 011100;
      af33 011110;
      af41 100010;
      af42 100100;
      af43 100110;
      be 000000;
      cs1 001000;
      cs2 010000;
      cs3 011000;
      cs4 100000;
      cs5 101000;
      cs6 110000;
      cs7 111000;
      ef 101110;
   }
   inet-precedence {
      af11 001;
      af21 010;
      af31 011;
      af41 100;
      be 000;
      cs6 110;
      cs7 111;
      ef 101;
      nc1 110;
      nc2 111;
   }
}
forwarding-classes {
   queue 0 be;
   queue 1 ef;
   queue 2 af;
   queue 3 nc;
}
interfaces {
   ge-7/2/0 {
      scheduler-map sched-map;
      unit 0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
        classifiers {
          dscp dscp_default;
        }
      }
    }
    e1-7/3/6 {
      scheduler-map sched-map;
    }
    e1-7/3/7 {
      scheduler-map sched-map;
    }
    ls-7/0/0 {
      scheduler-map sched-map;
      unit 0 {
        classifiers {
          inet-precedence inet-precedence_default;
        }
      }
    }
}
scheduler-maps {
    sched-map {
      forwarding-class af scheduler af-scheduler;
      forwarding-class be scheduler be-scheduler;
      forwarding-class ef scheduler ef-scheduler;
      forwarding-class nc scheduler nc-scheduler;
    }
}
schedulers {
    af-scheduler {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
    be-scheduler {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
    ef-scheduler {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
    nc-scheduler {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
}
[edit firewall]
filter lfi_ls_filter {
    term term0 {
      from {
        destination-address {
          192.168.1.3/32;
        }
        precedence 5;
      }
      then {
```

Copyright © 2013, Juniper Networks, Inc.

```
        count count-192-168-1-3;
        forwarding-class af;
        accept;
      }
    }
    term default {
      then {
        log;
        forwarding-class best effort;
        accept;
      }
    }
}
```

Examples: Configuring Multilink Interfaces

The examples in this section include only the configuration of multilink interfaces. For information about configuring the constituent interfaces, see the Junos® OS Network Interfaces.

The examples in this section show the following configurations:

- Example: Configuring a Multilink Interface with MLPPP on page 1156
- Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces on page 1157
- Example: Configuring a Multilink Interface with MLFR FRF.15 on page 1158

Example: Configuring a Multilink Interface with MLPPP

```
[edit interfaces]
ml-1/0/0 {
  unit 1 {
    fragment-threshold 128;
    family inet {
      address 192.168.5.1/32 {
        destination 192.168.200.200;
      }
    }
  }
  unit 10 {
    family inet {
      address 10.1.1.3/32 {
        destination 10.1.1.2;
      }
    }
  }
}
t1-5/1/0 {
  unit 0 {
    family mlppp {
      bundle ml-1/0/0.1;
    }
  }
}
t1-5/1/1 {
```


Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family mlppp {
        bundle ml-1/0/0.1;
      }
    }
}
t1-5/1/2 {
    unit 0 {
      family mlppp {
        bundle ml-1/0/0.1;
      }
    }
}
```

Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces

```
[edit interfaces]
at-0/0/0 {
    atm-options {
      pic-type atm2;
      vpi 10;
    }
    unit 0 {
      encapsulation atm-mlppp-llc;
      ppp-options {
        chap {
          access-profile pe-B-ppp-clients;
          local-name “pe-A-at-0/0/0”;
        }
      }
      keepalive interval 5 up-count 6 down-count 4;
      vci 10.120;
      family mlppp {
        bundle ls-0/3/0.0;
      }
    }
}
at-0/0/1 {
    atm-options {
      pic-type atm2;
      vpi 11;
    }
    unit 1 {
      encapsulation atm-mlppp-llc;
      ppp-options {
        chap {
          access-profile pe-B-ppp-clients;
          local-name “ pe-A-at-0/0/0”;
        }
      }
      keepalive interval 5 up-count 6 down-count 4;
      vci 11.120;
      family mlppp {
        bundle ls-0/3/0.0;
      }
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
at-1/2/3 {
   atm-options {
      pic-type atm2;
      vpi 12;
   }
   unit 2 {
      encapsulation atm-mlppp-llc;
      ppp-options {
         chap {
            access-profile pe-B-ppp-clients;
            local-name " pe-A-at-0/0/0";
         }
      }
      keepalive interval 5 up-count 6 down-count 4;
      vci 12.120;
      family mlppp {
         bundle ls-0/3/0.0;
      }
   }
}
...
ls-0/3/0 {
   encapsulation multilink-ppp;
   interleave-fragments;
   keepalive;
   unit 0 {
      mrru 4500;
      short-sequence;
      fragment-threshold 16320;
      drop-timeout 2000;
      encapsulation multilink-ppp;
      interleave-fragments;
      minimum-links 8;
      family inet {
         address 10.10.0.1/32 {
            destination 10.10.0.2;
         }
      }
      family iso;
      family inet6 {
         address 2001:DB8:0:1/32 {
            destination 2001:DB8:0:2;
         }
      }
   }
   ...
}
```

Example: Configuring a Multilink Interface with MLFR FRF.15

```
[edit interfaces]
ml-1/0/0 {
   unit 1 {
      encapsulation multilink-frame-relay-end-to-end;
      family inet {
```


Copyright © 2013, Juniper Networks, Inc.

```
        address 192.168.5.2/32 {
          destination 192.168.5.3;
        }
      }
    }
    unit 10 {
      encapsulation multilink-frame-relay-end-to-end;
      family inet {
        address 10.1.1.3/32 {
          destination 10.1.1.2;
        }
      }
    }
  }
  t1-5/1/0 {
    unit 0 {
      dlci 16;
      encapsulation multilink-frame-relay-end-to-end;
      family mlfr-end-to-end {
        bundle ml-1/0/0.1;
      }
    }
  }
  t1-5/1/1 {
    unit 0 {
      dlci 17;
      encapsulation multilink-frame-relay-end-to-end;
      family mlfr-end-to-end {
        bundle ml-1/0/0.10;
      }
    }
  }
  t1-5/1/2 {
    unit 0 {
      dlci 26;
      encapsulation multilink-frame-relay-end-to-end;
      family mlfr-end-to-end {
        bundle ml-1/0/0.10;
      }
    }
  }
```

Examples: Configuring Link Interfaces

The examples in this section include only the configuration of link interfaces. For information about configuring the constituent interfaces, see the Junos® OS Network Interfaces

- Example: Configuring a Link Services Interface with Two Links on page 1160
- Example: Configuring a Link Services Interface with MLPPP on page 1161
- Example: Configuring a Link Services Interface with MLFR FRF.15 on page 1162

Copyright © 2013, Juniper Networks, Inc.

- Example: Configuring a Link Services PIC with MLFR FRF.16 on page 1162
- Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types on page 1163

Example: Configuring a Link Services Interface with Two Links

This example uses the MLFR UNI NNI protocol between Router A and Router B and logically connects link services bundles **ls-1/1/0.3** and **ls-0/0/0.10**, as specified in Table 24 on page 1160.

Table 24: Link Services Bundle

Router A	Router B
t1-0/1/0 (ls-1/1/0:3)	t1-0/3/0 (ls-0/0/0:10)
t1-0/1/1 (ls-1/1/0:3)	t1-0/3/1 (ls-0/0/0:10)

For LMI to work properly, you must configure one router to be a DCE.

Configuration on Router A

```
[edit interfaces]
ls-1/1/0:3 {
    dce;
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        dlci 16;
        family inet {
            address 10.3.3.1/32 {
                destination 10.3.3.2;
            }
        }
    }
}
t1-0/1/0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/1/0:3;
        }
    }
}
t1-0/1/1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
            bundle ls-1/1/0:3;
        }
    }
}
```

Configuration on Router B

```
[edit interfaces]
ls-0/0/0:10 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
      dlci 16;
      family inet {
        address 10.3.3.2/32 {
          destination 10.3.3.1;
        }
      }
    }
  }
  t1-0/3/0 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      family mlfr-uni-nni {
        bundle ls-0/0/0:10;
      }
    }
  }
  t1-0/3/1 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      family mlfr-uni-nni {
        bundle ls-0/0/0:10;
      }
    }
  }
```

Example: Configuring a Link Services Interface with MLPPP

```
[edit interfaces]
t1-0/0/0 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
t1-0/0/1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
ls-0/3/0 {
  unit 0 {
    encapsulation multilink-ppp;
    family inet {
      address 10.16.1.2/32 {
        destination 10.16.1.1;
      }
    }
    family iso;
    family inet6 {
      address 2001:DB8:1:2/126;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
```

Example: Configuring a Link Services Interface with MLFR FRF.15

```
[edit interfaces]
t1-0/0/0 {
  encapsulation frame-relay;
  unit 0 {
    dlci 16;
    family mlfr-end-to-end {
      bundle ls-0/3/0.0;
    }
  }
}
t1-0/0/1 {
  encapsulation frame-relay;
  unit 0 {
    dlci 16;
    family mlfr-end-to-end {
      bundle ls-0/3/0.0;
    }
  }
}
ls-0/3/0 {
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 10.16.1.2/32 {
        destination 10.16.1.1;
      }
    }
    family iso;
    family inet6 {
      address 2001:DB8:1:2/12;
    }
  }
}
```

Example: Configuring a Link Services PIC with MLFR FRF.16

```
[edit chassis]
fpc 1 {
  pic 2 {
    mlfr-uni-nni-bundles 5;
  }
}
[edit interfaces]
t1-0/0/0 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle ls-1/2/0:0;
    }
  }
```


Copyright © 2013, Juniper Networks, Inc.

```
}
t1-0/0/1 {
   encapsulation multilink-frame-relay-uni-nni;
   unit 0 {
      family mlfr-uni-nni {
         bundle ls-1/2/0:0;
      }
   }
}
ls-1/2/0:0 {
   dce;
   encapsulation multilink-frame-relay-uni-nni;
   unit 0 {
      dlci 26;
      family inet {
         address 10.26.1.1/32 {
            destination 10.26.1.2;
         }
      }
   }
}
```

Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types

```
[edit chassis]
fpc 1 {
   pic 3 {
      mlfr-uni-nni-bundles 4;
   }
}
[edit interfaces]
t1-0/2/0:0 {
   encapsulation multilink-frame-relay-uni-nni;
   unit 0 {
      family mlfr-uni-nni {
         bundle ls-1/3/0:0;
      }
   }
}
t1-0/2/0:1 {
   encapsulation multilink-frame-relay-uni-nni;
   unit 0 {
      family mlfr-uni-nni {
         bundle ls-1/3/0:0;
      }
   }
}
t1-0/2/0:5 {
   unit 0 {
      family mlppp {
         bundle ls-1/3/0.2;
      }
   }
}
t1-0/2/0:6 {
   unit 0 {
```

Copyright © 2013, Juniper Networks, Inc.

```
      family mlppp {
        bundle ls-1/3/0.2;
      }
    }
}
t1-0/2/0:7 {
    encapsulation frame-relay;
    unit 0 {
      dlci 20;
      family mlfr-end-to-end {
        bundle ls-1/3/0.1;
      }
    }
}
t1-0/2/0:8 {
    encapsulation frame-relay;
    unit 0 {
      dlci 20;
      family mlfr-end-to-end {
        bundle ls-1/3/0.1;
      }
    }
}
t1-0/2/0:10 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
      family mlppp {
        bundle lsq-1/1/0.0;
      }
    }
}
t3-1/0/0 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
      family mlppp {
        bundle lsq-1/1/0.2;
      }
    }
}
lsq-1/1/0 {
    unit 0 {
      encapsulation multilink-ppp;
      compression {
        rtp {
          f-max-period 100;
          queues [ q1 q2 ];
          port minimum 2000 maximum 6000;
        }
      }
      family inet {
        address 10.5.5.5/24;
      }
    }
    unit 1 {
```


Copyright © 2013, Juniper Networks, Inc.

```
      encapsulation multilink-ppp;
      compression {
        rtp {
          port minimum 2000 maximum 6000;
        }
      }
      family inet {
        address 10.6.6.1/24;
      }
    }
    unit 2 {
      encapsulation multilink-ppp;
      compression {
        rtp {
          port minimum 2000 maximum 6000;
        }
      }
      family inet {
        address 10.9.9.1/24;
      }
    }
  }
  t1-1/2/0 {
    no-keepalives;
    unit 0 {
      family mlppp {
        bundle lsq-1/1/0.1;
      }
    }
  }
  ls-1/3/0 {
    unit 1 {
      encapsulation multilink-frame-relay-end-to-end;
      family inet {
        address 10.1.4.1/24;
      }
    }
    unit 2 {
      encapsulation multilink-ppp;
      family inet {
        address 10.7.4.1/24;
      }
    }
  }
  ls-1/3/0:0 {
    encapsulation multilink-frame-relay-uni-nni;
    mlfr-uni-nni-bundle-options {
      debug-flags 15;
    }
    unit 0 {
      dlci 20;
      family inet {
        address 10.5.4.1/24;
      }
    }
  }
```

Copyright © 2013, Juniper Networks, Inc.

```
[edit routing-options]
static {
  route 10.12.12.0/24 next-hop 10.1.1.9;
}
```

On Router B:

```
[edit chassis]
fpc 1 {
  pic 3 {
    mlfr-uni-nni-bundles 4;
  }
}
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.1.1.1/24;
    }
  }
}
so-0/1/1 {
  encapsulation ppp;
  unit 0 {
    family inet {
      address 10.7.7.7/24;
    }
  }
}
t1-0/2/0:0 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle ls-1/3/0:0;
    }
  }
}
t1-0/2/0:1 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle ls-1/3/0:0;
    }
  }
}
t1-0/2/0:5 {
  no-keepalives;
  unit 0 {
    family mlppp {
      bundle ls-1/3/0.2;
    }
  }
}
t1-0/2/0:6 {
  no-keepalives;
  unit 0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
      family mlppp {
        bundle ls-1/3/0.2;
      }
    }
  }
  t1-0/2/0:7 {
    dce;
    encapsulation frame-relay;
    unit 0 {
      dlci 20;
      family mlfr-end-to-end {
        bundle ls-1/3/0.1;
      }
    }
  }
  t1-0/2/0:8 {
    dce;
    encapsulation frame-relay;
    unit 0 {
      dlci 20;
      family mlfr-end-to-end {
        bundle ls-1/3/0.1;
      }
    }
  }
  t1-0/2/0:10 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
      family mlppp {
        bundle lsq-1/1/0.0;
      }
    }
  }
  t3-0/3/0 {
    no-keepalives;
    encapsulation ppp;
    unit 0 {
      family mlppp {
        bundle lsq-1/1/0.2;
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
      family inet {
        address 10.2.2.1/24;
      }
    }
  }
  lsq-1/1/0 {
    unit 0 {
      compression {
        rtp {
          port minimum 2000 maximum 6000;
        }
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
        family inet {
          address 10.5.5.1/24;
        }
      }
      unit 1 {
        encapsulation multilink-ppp;
        compression {
          rtp {
            port minimum 16384 maximum 20102;
          }
        }
        family inet {
          address 10.3.4.1/24;
        }
      }
      unit 2 {
        encapsulation multilink-ppp;
        compression {
          rtp {
            port minimum 2000 maximum 6000;
          }
        }
        family inet {
          address 10.9.9.9/24;
        }
      }
    }
    t1-1/2/2 {
      no-keepalives;
      unit 0 {
        family mlppp {
          bundle ls-1/3/0.1;
        }
      }
    }
    t1-1/2/3 {
      no-keepalives;
      unit 0 {
        family mlppp {
          bundle lsq-1/1/0.1;
        }
      }
    }
    ls-1/3/0 {
      unit 1 {
        encapsulation multilink-frame-relay-end-to-end;
        family inet {
          address 10.1.4.4/24;
        }
        family iso;
      }
      unit 2 {
        encapsulation multilink-ppp;
        family inet {
          address 10.7.4.4/24;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
    ls-1/3/0:0 {
      dce;
      encapsulation multilink-frame-relay-uni-nni;
      unit 0 {
        dlci 20;
        family inet {
          address 10.5.4.4/24;
        }
      }
    }
    [edit routing-options]
    static {
      route 10.12.12.0/24 next-hop 10.3.4.4;
    }
```

Example: Configuring an MLPPP Bundle

This example shows how to configure an MLPPP bundle to increase traffic bandwidth.

- Requirements on page 1169
- Overview on page 1169
- Configuration on page 1170
- Verification on page 1172

Requirements

Before you begin, you should have two MX Series routers (MX240, MX480, or MX960 routers) configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you create the MLPPP bundle lsq-0/0/0.0 at the logical unit level of the link services interface lsq-0/0/0 on the MX Series routers R0 and R1. You then add the two serial interfaces se-1/0/0 and se-1/0/1 as constituent links to the multilink bundle. In Figure 16 on page 1169, your company's branch office is connected to its main branch using routers R0 and R1. You transmit data and voice traffic on two low-speed 1-Mbps serial links. To increase bandwidth, you configure MLPPP and join the two serial links se-1/0/0 and se-1/0/1 into the multilink bundle lsq-0/0/0.0. Then you configure LFI and CoS on R0 and R1 to enable them to transmit voice packets ahead of data packets.

Figure 16: Configuring MLPPP and LFI on Serial Links

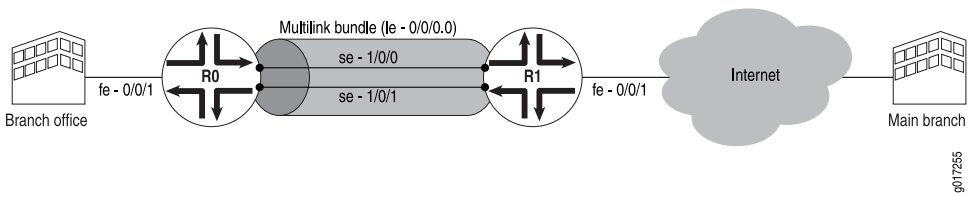

Configuration

CLI Quick Configuration

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

For device R0
```
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.10/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
set interfaces se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz
```

For device R1
```
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.9/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
```

Step-by-Step Procedure

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

To configure MLPPP bundle:

1. Create an interface on both the routers.

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

2. Configure a family inet and define the IP address on device R0.

   ```
   [edit interfaces lsq-0/0/0 unit 0]
   user@host# set family inet address 10.0.0.10/24
   ```

3. Configure a family inet and define the IP address on device R1.

   ```
   [edit interfaces lsq-0/0/0 unit 0]
   user@host# set family inet address 10.0.0.9/24
   ```

4. Specify the names of the constituent links to be added to the multilink bundle on both the routers.

   ```
   [edit interfaces]
   user@host# edit se-1/0/0 unit 0
   user@host# set family mlppp bundle lsq-0/0/0.0
   [edit interfaces]
   user@host# edit se-1/0/1 unit 0
   user@host# set family mlppp bundle lsq-0/0/0.0
   ```

5. Set the serial options to the same values for both interfaces on R0.



NOTE: R0 is set as a DCE device. The serial options are not set for interfaces on R1. You can set the serial options according to your network setup.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
user@host# set se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz
```

Results From configuration mode, confirm your configuration by entering the **show interfaces lsq-0/0/0**, **show interfaces se-1/0/0**, and **show interfaces se-1/0/1** commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
  user@host# show interfaces lsq-0/0/0
family inet {
  address 10.0.0.10/24;
    }
  }
[edit]
user@host# show interfaces se-1/0/0
  clocking-mode dce;
    clock-rate 2.0mhz;
    }
    unit 0 {
      family mlppp {
      bundle lsq-0/0/0.0;
    }
  }
[edit]
  user@host# show interfaces se-1/0/1
serial-options {
  clocking-mode dce;
    clock-rate 2.0mhz;
    }
    unit 0 {
      family mlppp {
      bundle lsq-0/0/0.0;
    }
  }

For device R1
[edit]
user@host# show interfaces lsq-0/0/0
  family inet {
    address 10.0.0.9/24;
    }
  }
[edit]
  user@host# show interfaces se-1/0/0
  unit 0 {
  family mlppp {
    bundle lsq-0/0/0.0;
  }
}
[edit]
  user@host# show interfaces se-1/0/1
  unit 0 {
```

Copyright © 2013, Juniper Networks, Inc.

```
    family mlppp {
      bundle lsq-0/0/0.0;
    }
  }
```

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

Verification

Confirm that the configuration is working properly:

- Verifying the MLPPP Bundle on page 1172

Verifying the MLPPP Bundle

Purpose Verify that the constituent links are added to the bundle correctly.

Action From operational mode, enter the **show interfaces lsq-0/0/0 statistics** command.

Related Documentation

- Link and Multilink Services Overview on page 1123
- Multilink Interfaces on Channelized MICs Overview on page 1131
- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links on page 1133
- Example: Configuring Link Interfaces on Channelized MICs on page 1179
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175

Example: Configuring Multilink Frame Relay FRF.15

This example shows how to configure MLFR FRF.15 for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1, E1, and serial links.

- Requirements on page 1172
- Overview on page 1172
- Configuration on page 1173
- Verification on page 1175

Requirements

Before you begin, you should have two MX Series 3D Universal Edge Routers (MX240, MX480, or MX960 routers) configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you aggregate two T1 links to create the MLFR FRF.15 bundle on two MX Series routers, R0 and R1, and set the interface to lsq-0/0/0. You configure a logical unit

Copyright © 2013, Juniper Networks, Inc.

on the lsq-0/0/0 interface and set the family type to **inet** with address 10.0.0.4/24. Then you configure an IP address for the multilink bundle on the unit level of the interface.

You define the multilink bundle as an MLFR FRF.15 bundle by specifying the MLFR end-to-end encapsulation type. You specify the names of the constituent links to be added to the multilink bundle as t1-2/0/0 and t1-2/0/1 and set the encapsulation type to **frame-relay**. You then define R0 as a DCE device and R1 as a DTE device. You set the DLCI value to 100 (range is from 16 through 1022). Finally, you set the multilink bundle to lsq-0/0/0.0.

Configuration

CLI Quick Configuration

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

For device R0

```
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.4/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

For device R1

```
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.5/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

Step-by-Step Procedure

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

To configure the MLFR FRF.15 bundle:

1. Create an interface on both the routers.

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

2. Set a logical unit on the interface and define the family type for the routers R0 and R1.

   ```
   [edit interfaces lsq-0/0/0 unit 0]
   user@host# set family inet address 10.0.0.4/24
   user@host# set family inet address 10.0.0.5/24
   ```

3. Define the multilink bundle as an MLFR FRF.15 bundle.

   ```
   [edit interfaces lsq-0/0/0 unit 0]
   user@host# set encapsulation multilink-frame-relay-end-to-end
   ```

4. Specify the names of the constituent links to be added to the multilink bundle.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation frame-relay
user@host# set t1-2/0/1 encapsulation frame-relay
```

5. Define the router R0 as a DCE device.

```
[edit interfaces]
user@host# edit lsq-0/0/0
user@host# set dce
```

6. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces lsq-0/0/0**, **show interfaces t1-2/0/0**, and **show interfaces t1-2/0/1** commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
  user@host# show interfaces lsq-0/0/0
  dce;
unit 0 {
  encapsulation multilink-frame-relay-end-to-end;
    dlci 100;
  family inet {
    address 10.0.0.4/24;
      }
    family mlfr-end-to-end {
    bundle lsq-0/0/0.0;
      }
    }
[edit]
  user@host#show interfaces t1-2/0/0
encapsulation frame-relay;
  [edit]
    user@host# show interfaces t1-2/0/1
  encapsulation frame-relay;

For device R1
[edit]
  user@host# show interfaces lsq-0/0/0
  unit 0 {
  encapsulation multilink-frame-relay-end-to-end;
    dlci 100;
    family inet {
  address 10.0.0.5/24;
  }
  family mlfr-end-to-end {
    bundle lsq-0/0/0.0;
  }
}
[edit]
  user@host# show interfaces t1-2/0/0
encapsulation frame-relay;
```


Copyright © 2013, Juniper Networks, Inc.

```
[edit]
  user@host# show interfaces t1-2/0/1
encapsulation frame-relay;
```

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

Verification

Confirm that the configuration is working properly:

- Verifying the MLFR FRF.15 Configuration on page 1175

Verifying the MLFR FRF.15 Configuration

Purpose Verify the MLFR FRF.15 configuration.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- Link and Multilink Services Overview on page 1123
- Multilink Interfaces on Channelized MICs Overview on page 1131
- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links on page 1133
- Example: Configuring Link Interfaces on Channelized MICs on page 1179
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175

Example: Configuring Multilink Frame Relay FRF.16

This example shows how to configure MLFR FRF.16 for additional bandwidth, load balancing, and redundancy.

- Requirements on page 1175
- Overview on page 1175
- Configuration on page 1176
- Verification on page 1178

Requirements

Before you begin, you should have two MX Series 3D Universal Edge Routers configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you aggregate two T1 interfaces to create an MLFR FRF.16 bundle on two MX Series, R0 and R1. You configure the chassis interface and specify the number of MLFR FRF.16 bundles to be created on the interface. You then specify the channel to

Copyright © 2013, Juniper Networks, Inc.

be configured as a multilink bundle and create interface lsq-0/0/0:0. You set the multilink bundle as an MLFR FRF.16 bundle by specifying the MLFR UNI NNI encapsulation type.

Then you define R0 as a DCE device and R1 as a DTE device. You configure a logical unit on the multilink bundle lsq-0/0/0:0, and set the family type to inet. You then assign a DLCI of 400 and an IP address of 10.0.0.10/24 to the multilink bundle. You create the T1 interfaces, t1-2/0/0 and t1-2/0/1, that are to be added as constituent links to the multilink bundle and define the Frame Relay encapsulation type. Finally, you set the multilink bundle to lsq-0/0/0:0.

Configuration

CLI Quick Configuration

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

```
For device R0
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.10/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
For device R1
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.9/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Step-by-Step Procedure

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

To configure an MLFR FRF.16 bundle:

1. Configure a chassis interface.

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

2. Specify the number of MLFR bundles.

   ```
   [edit chassis]
   user@host# set fpc 0 pic 0 mlfr-uni-nni-bundles 1
   ```

3. Create an interface.

   ```
   [edit]
   user@host# edit interfaces lsq-0/0/0:0
   ```

4. Specify the MLFR encapsulation type.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces lsq-0/0/0:0]
user@host# set encapsulation multilink-frame-relay-uni-nni
```

5. Set the router R0 as a DCE device.

```
[edit]
user@host# edit interfaces lsq-0/0/0
user@host# set dce
```

6. Specify a logical unit on the multilink bundle and set the family type.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 400 family inet address 10.0.0.10/24
```

7. Create the T1 interfaces and set the Frame Relay encapsulation.

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
user@host# set t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
```

8. Specify the multilink bundle to which the interface is to be added as a constituent link on device R0.

```
[edit interfaces t1-2/0/0]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

9. Specify the multilink bundle to which the interface is to be added as a constituent link on device R1.

```
[edit interfaces t1-2/0/1]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Results From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces lsq-0/0/0**, **show interfaces lsq-0/0/0:0**, **show interfaces t1-2/0/0**, and **show interfaces t1-2/0/1** commands for the routers R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
  user@host# show chassis
  fpc 0 {
  pic 0 {
  mlfr-uni-nni-bundles 1;
    }
  }
  [edit]
    user@host# show interfaces lsq-0/0/0
    dce;
  unit 0 {
  dlci 400;
    family inet {
    address 10.0.0.10/24;
  }
  }
  [edit]
    user@host# show interfaces lsq-0/0/0:0
  encapsulation multilink-frame-relay-uni-nni;
```

Copyright © 2013, Juniper Networks, Inc.

```
  [edit]
    user@host# show interfaces t1-2/0/0
    encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
  family mlfr-uni-nni {
    bundle lsq-0/0/0:0;
  }
}
[edit]
  user@host# show interfaces t1-2/0/1
  encapsulation multilink-frame-relay-uni-nni;
unit 0 {
family mlfr-uni-nni {
  bundle lsq-0/0/0:0;
}
}
```

For device R1

```
[edit]
  user@host# show chassis
  unit 0 {
  dlci 400;
  family inet {
    address 10.0.0.9/24;
  }
  }
  [edit]
    user@host# show interfaces lsq-0/0/0:0
  encapsulation multilink-frame-relay-uni-nni;
  [edit]
    user@host# show interfaces t1-2/0/0
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      family mlfr-uni-nni {
    bundle lsq-0/0/0:0;
  }
  }
  [edit]
    user@host# show interfaces t1-2/0/1
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      family mlfr-uni-nni {
    bundle lsq-0/0/0:0;
  }
}
```

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

Verification

Confirm that the configuration is working properly:

- Verifying the MLFR FRF.16 Configuration on page 1179

 Copyright © 2013, Juniper Networks, Inc.

Verifying the MLFR FRF.16 Configuration

Purpose Verify the MLFR FRF.16 configuration.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- Link and Multilink Services Overview on page 1123
- Multilink Interfaces on Channelized MICs Overview on page 1131
- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links on page 1133
- Example: Configuring Link Interfaces on Channelized MICs on page 1179
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172

Example: Configuring Link Interfaces on Channelized MICs

- Requirements on page 1179
- Overview on page 1179
- Configuration on 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP on page 1180
- Verification on page 1187

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 or later for MX240, MX480, and MX960 routers
- One MX240, MX480, or MX960 router

Overview

This example provides information about configuring the link interfaces on the following channelized MICs:

- 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
- 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
- 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

You need to first partition each port on the MICs to configure the link interfaces T1, T3, and DS, and then you configure the link interfaces for bundles. An MLPPP bundle involves "bundling" multiple T1/T3/DS interfaces into a single, logical interface that uses only one IP address. For more information about MLPPP bundles, see "Understanding MLPPP

Copyright © 2013, Juniper Networks, Inc.

Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links" on page 1133. Similarly, you can partition the ports to configure the MICs to the E1/E3 interfaces by setting the framing mode to SDH.

For more information about multilink-based protocols on MX240, MX480, and MX960 routers with Multiservices DPC, see "Multilink Interfaces on Channelized MICs Overview" on page 1131.

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



NOTE: You can set the values for each parameter according to your requirement. The values given in this example are for illustration purposes only.

Configuration on 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP

To partition each port on the MIC and configure the link interfaces T1, T3, and DS on it and to configure the link interfaces for bundles, perform the following tasks:

- Partitioning Ports on the Channelized MICs and Configuring the Link Interfaces T1, T3, and DS on page 1181
- Configuring MLPPP, MLFR FRF.15, and MLFR FRF.16 on Link Interfaces for Bundles on page 1183
- Results on page 1184

CLI Quick Configuration

To quickly configure synchronization on the aforementioned routers, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

[edit]

```
set interfaces coc12-5/2/0 partition 1 interface-type coc1
set interfaces coc12-5/2/0 partition 1 oc-slice 1
set interfaces coc12-5/2/0 partition 2 oc-slice 2 interface-type coc1
set interfaces coc12-5/2/0 partition 3 oc-slice 3 interface-type coc1
set interfaces coc1-5/2/0:1 no-partition interface-type t3
set interfaces coc1-5/2/0:3 no-partition interface-type t3
set interfaces coc1-5/2/0:2 partition 1 interface-type ct1
set interfaces coc1-5/2/0:2 partition 2 interface-type t1
set interfaces coc1-5/2/0:2 partition 3 interface-type ct1
set interfaces coc1-5/2/0:2 partition 4 interface-type t1
set interfaces ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
set interfaces ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
set interfaces ct3-2/0/0 no-partition interface-type t3
set interfaces ct3-2/0/0 partition 1 interface-type t1
set interfaces ct3-2/0/0 partition 2 interface-type t1
set interfaces ct3-2/0/0 partition 3 interface-type t1
```


Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ct3-2/0/0 partition 4 interface-type ct1
set interfaces ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
set interfaces t1-5/2/0:2:2 unit 0 family mlppp bundle rlsq0.1
set interfaces ds-5/2/0:2:1:1 unit 0 family mlppp bundle rlsq0.0
set interfaces ds-5/2/0:2:1:2 unit 0 family mlppp bundle rlsq0.0
set interfaces t3-5/2/0:3 unit 0 family mlppp bundle rlsq0.2
set interfaces t3-5/2/0:1 unit 0 family mlppp bundle rlsq0.2
set interfaces t1-5/2/0:2:2 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces t1-5/2/0:2:4 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces ds-5/2/0:2:1:1 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces ds-5/2/0:2:1:2 encapsulation multilink-frame-relay-uni-nni unit 0 family
mlfr-uni-nni bundle rlsq0:0
set interfaces t1-5/2/0:2:2 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end
bundle rlsq0.0
set interfaces t1-5/2/0:2:4 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end
bundle rlsq0.0
set interfaces ds-5/2/0:2:1:1 encapsulation frame-relay unit 0 dlci 10 family
mlfr-end-to-end bundle rlsq0.0
set interfaces ds-5/2/0:2:1:2 encapsulation frame-relay unit 0 dlci 11 family
mlfr-end-to-end bundle rlsq0.0
set interfaces t3-5/2/0:1 encapsulation frame-relay unit 0 dlci 11 family mlfr-end-to-end
bundle rlsq0.1
set interfaces t3-5/2/0:3 encapsulation frame-relay unit 0 dlci 10 family mlfr-end-to-end
bundle rlsq0.1
```

Partitioning Ports on the Channelized MICs and Configuring the Link Interfaces T1, T3, and DS

Step-by-Step Procedure

To partition each port on the channelized MICs:

1. Configure the **coc12-5/2/0** interface by setting the **partition** option to **1** with the sublevel interface type set to **coc1**.

   ```
   [edit interfaces]
   user@host# set coc12-5/2/0 partition 1 interface-type coc1
   ```

2. Configure the **coc12-5/2/0** interface with the OC-slice range (OC-slice range specifies the bandwidth size required for the interface type you are configuring) set to **1**.

   ```
   [edit interfaces]
   user@host# set coc12-5/2/0 partition 1 oc-slice 1
   ```

3. Configure the **coc12-5/2/0** interface by setting the **partition** option to **2** with the sublevel interface type set to **coc1** and the OC-slice range set to **2**.

   ```
   [edit interfaces]
   user@host# set coc12-5/2/0 partition 2 oc-slice 2 interface-type coc1
   ```

4. Configure the **coc12-5/2/0** interface by setting the **partition** option to **3** with the sublevel interface type set to **coc1** and the OC-slice range set to **3**.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set coc12-5/2/0 partition 3 oc-slice 3 interface-type coc1
```

5. Configure the **coc1-5/2/0:1** interface as a clear channel by setting the **no-partition** option for the sublevel interface type **t3** . (A clear channel consolidates the entire bandwidth of a channelized interface into a single unpartioned stream that looks like a standard interface.)

```
[edit interfaces]
user@host# set coc1-5/2/0:1 no-partition interface-type t3
```

6. Configure the **coc1-5/2/0:3** interface as a clear channel by setting the **no-partition** option for the sublevel interface type **t3**.

```
[edit interfaces]
user@host# set coc1-5/2/0:3 no-partition interface-type t3
```

7. Configure the **coc1-5/2/0:2** interface by setting the **partition** option to **1** and **3** with the sublevel interface type set to **ct1**. Configure the **coc1-5/2/0:2** interface by setting the **partition** option to **2** and **4** with the sublevel interface type set to **t1**.

```
[edit interfaces]
user@host# set coc1-5/2/0:2 partition 1 interface-type ct1
user@host# set coc1-5/2/0:2 partition 2 interface-type t1
user@host# set coc1-5/2/0:2 partition 3 interface-type ct1
user@host# set coc1-5/2/0:2 partition 4 interface-type t1
```

8. Configure the **ct1-5/2/0:2:1** interface by setting the **partition** option to **1** and **2** with the sublevel interface type set to **ds**. Configure the time slots for the partitions.

```
[edit interfaces]
user@host# set ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
user@host# set ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
```

9. Configure a clear channel on the channelized interface **ct3-2/0/0** by setting the **no-partition** option to the sublevel interface type **t3** (a clear channel consolidates the entire bandwidth of a channelized interface into a single unpartitioned stream that looks like a standard interface).

```
[edit interfaces]
user@host# set ct3-2/0/0 no-partition interface-type t3
```

10. Configure a clear channel on the channelized interface **ct3-2/0/0** by setting the **partition** option to **1**, **2**, and **3** with the sublevel interface type **ds**. Configure the **ct3-2/0/0** interface by setting the **partition** option to **4** with the sublevel interface type **ct1**.

```
[edit interfaces]
user@host# set ct3-2/0/0 partition 1 interface-type t1
user@host# set ct3-2/0/0 partition 2 interface-type t1
user@host# set ct3-2/0/0 partition 3 interface-type t1
user@host# set ct3-2/0/0 partition 4 interface-type ct1
```

11. Configure the **ct1-2/0/0:4** interface by setting the **partition** option to **1** and **2** with the sublevel interface type set to **ds**. Configure the time slots for the partitions.

```
[edit interfaces]
user@host# set ct1-5/2/0:2:1 partition 1 timeslots 1 interface-type ds
user@host# set ct1-5/2/0:2:1 partition 2 timeslots 2 interface-type ds
```


Copyright © 2013, Juniper Networks, Inc.

Results Display the results of partitioning each port on the MIC and configuring the link interfaces T1, T3, and DS:

Results for CHOC12/3 interfaces

```
user@host# show interfaces
coc12-5/2/0 {
    partition 1 oc-slice 1 interface-type coc1;
    partition 2 oc-slice 2 interface-type coc1;
    partition 3 oc-slice 3 interface-type coc1;
}
coc1-5/2/0:1 {
    no-partition interface-type t3;
}
coc1-5/2/0:3 {
    no-partition interface-type t3;
}
coc1-5/2/0:2 {
    partition 1 interface-type ct1;
    partition 2 interface-type t1;
    partition 3 interface-type ct1;
    partition 4 interface-type t1;
}
```

Results for CHDS3 MIC interfaces

```
user@host# show interfaces
ct1-5/2/0:2:1 {
    partition 1 timeslots 1 interface-type ds;
    partition 2 timeslots 2 interface-type ds;
}
ct3-2/0/0 {
    no-partition interface-type t3;
    partition 1 interface-type t1;
    partition 2 interface-type t1;
    partition 3 interface-type t1;
    partition 4 interface-type ct1;
}
ct1-2/0/0:4 {
    partition 1 timeslots 1 interface-type ds;
    partition 2 timeslots 2 interface-type ds;
}
```

Configuring MLPPP, MLFR FRF.15, and MLFR FRF.16 on Link Interfaces for Bundles

Step-by-Step Procedure To configure MLPPP, MLFR FRF.15, and MLFR FRF.16 on the link interfaces T1, T3, and DS for bundles:

1. Configure the MLPPP encapsulation on the T1 link interfaces **t1-5/2/0:2:2** and **t1-5/2/0:2:4**.

   ```
   [edit interfaces]
   user@host# set t1-5/2/0:2:2 unit 0 family mlppp bundle rlsq0.1
   user@host# set t1-5/2/0:2:4 unit 0 family mlppp bundle rlsq0.1
   ```

2. Configure the MLPPP encapsulation on the DS link interfaces **ds-5/2/0:2:1:1** and **ds-5/2/0:2:1:2**.

   ```
   [edit interfaces]
   user@host# set ds-5/2/0:2:1:1 unit 0 family mlppp bundle rlsq0.0
   ```

Copyright © 2013, Juniper Networks, Inc.

```
user@host# set ds-5/2/0:2:1:2 unit 0 family mlppp bundle rlsq0.0
```

3. Configure the MLPPP encapsulation on the T3 link interfaces **t3-5/2/0:3** and **t3-5/2/0:1**.

```
[edit interfaces]
user@host# set t3-5/2/0:3 unit 0 family mlppp bundle rlsq0.2
user@host# set t3-5/2/0:1 unit 0 family mlppp bundle rlsq0.2
```

4. Configure the MLFR FRF.16 encapsulation on the T1 link interfaces **t1-5/2/0:2:2** and **t1-5/2/0:2:4**.

```
[edit interfaces]
user@host# set t1-5/2/0:2:2 encapsulation multilink-frame-relay-uni-nni unit 0
  family mlfr-uni-nni bundle rlsq0:0
user@host# set t1-5/2/0:2:4 encapsulation multilink-frame-relay-uni-nni unit 0
  family mlfr-uni-nni bundle rlsq0:0
```

5. Configure the MLFR FRF.16 encapsulation on the DS link interfaces **ds-5/2/0:2:1:1** and **ds-5/2/0:2:1:2**.

```
[edit interfaces]
user@host# set ds-5/2/0:2:1:1 encapsulation multilink-frame-relay-uni-nni unit 0
  family mlfr-uni-nni bundle rlsq0:0
user@host# set ds-5/2/0:2:1:2 encapsulation multilink-frame-relay-uni-nni unit 0
  family mlfr-uni-nni bundle rlsq0:0
```

6. Configure the MLFR FRF.15 encapsulation on the T1 link interfaces **t1-5/2/0:2:2** and **t1-5/2/0:2:4**.

```
[edit interfaces]
user@host# set t1-5/2/0:2:2 encapsulation frame-relay unit 0 dlci 10 family
  mlfr-end-to-end bundle rlsq0.0
user@host# set t1-5/2/0:2:4 encapsulation frame-relay unit 0 dlci 11 family
  mlfr-end-to-end bundle rlsq0.0
```

7. Configure the MLFR FRF.15 encapsulation on the DS link interfaces **ds-5/2/0:2:1:1** and **ds-5/2/0:2:1:2**.

```
[edit interfaces]
user@host# set ds-5/2/0:2:1:1 encapsulation frame-relay unit 0 dlci 10 family
  mlfr-end-to-end bundle rlsq0.0
user@host# set ds-5/2/0:2:1:2 encapsulation frame-relay unit 0 dlci 11 family
  mlfr-end-to-end bundle rlsq0.0
```

8. Configure the MLFR FRF.15 encapsulation on the T3 link interfaces **t3-5/2/0:1** and **t3-5/2/0:3**.

```
[edit interfaces]
user@host# set t3-5/2/0:1 encapsulation frame-relay unit 0 dlci 11 family
  mlfr-end-to-end bundle rlsq0.1
user@host# set t3-5/2/0:3 encapsulation frame-relay unit 0 dlci 10 family
  mlfr-end-to-end bundle rlsq0.1
```

Results

Display the results of the configuration of link interfaces for bundles:

MLPPP on T1 links

```
user@host# show interfaces
```


Copyright © 2013, Juniper Networks, Inc.

```
t1-5/2/0:2:2 {
  unit 0 {
    family mlppp {
      bundle rlsq0.1;
    }
  }
}
t1-5/2/0:2:4 {
  unit 0 {
    family mlppp {
      bundle rlsq0.1;
    }
  }
}
```

MLPPP on DS links

```
user@host# show interfaces
ds-5/2/0:2:1:1 {
  unit 0 {
    family mlppp {
      bundle rlsq0.0;
    }
  }
}
ds-5/2/0:2:1:2 {
  unit 0 {
    family mlppp {
      bundle rlsq0.0;
    }
  }
}
```

MLPPP on T3 links

```
user@host# show interfaces
t3-5/2/0:3 {
  unit 0 {
    family mlppp {
      bundle rlsq0.2;
    }
  }
}
t3-5/2/0:1 {
  unit 0 {
    family mlppp {
      bundle rlsq0.2;
    }
  }
}
```

MLFR FRF.15 on T1 links

```
user@host# show interfaces
t1-5/2/0:2:2 {
  encapsulation frame-relay;
  unit 0 {
    dlci 10;
    family mlfr-end-to-end {
      bundle rlsq0.0;
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  t1-5/2/0:2:4 {
    encapsulation frame-relay;
      unit 0 {
      dlci 11;
      family mlfr-end-to-end {
        bundle rlsq0.0;
      }
    }
  }
```

MLFR FRF.15 on DS links

```
user@host# show interfaces
ds-5/2/0:2:1:1 {
  encapsulation frame-relay;
  unit 0 {
    dlci 10;
    family mlfr-end-to-end {
      bundle rlsq0.0;
    }
  }
}
ds-5/2/0:2:1:2 {
  encapsulation frame-relay;
    unit 0 {
    dlci 11;
    family mlfr-end-to-end {
      bundle rlsq0.0;
    }
  }
}
```

MLFR FRF.15 on T3 links

```
user@host# show interfaces
t3-5/2/0:1 {
  encapsulation frame-relay;
  unit 0 {
    dlci 11;
    family mlfr-end-to-end {
      bundle rlsq0.1;
    }
  }
}
t3-5/2/0:3 {
  encapsulation frame-relay;
    unit 0 {
    dlci 10;
    family mlfr-end-to-end {
      bundle rlsq0.1;
    }
  }
}
```

MLFR FRF.16 on T1 links

```
user@host# show interfaces
t1-5/2/0:2:2 {
  encapsulation multilink-frame-relay-uni-nni;
```


Copyright © 2013, Juniper Networks, Inc.

```
    unit 0 {
      family mlfr-uni-nni {
        bundle rlsq0:0;
      }
    }
  }
  t1-5/2/0:2:4 {
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      family mlfr-uni-nni {
        bundle rlsq0:0;
      }
    }
  }
```

MLFR FRF.16 on DS links

```
user@host# show interfaces
ds-5/2/0:2:1:1 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle rlsq0:0;
    }
  }
}
ds-5/2/0:2:1:2 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    family mlfr-uni-nni {
      bundle rlsq0:0;
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

- Verifying the MLPPP Bundle on page 1187
- Verifying the MLFR FRF.15 Configuration on page 1188
- Verifying the MLFR FRF.16 Configuration on page 1188

Verifying the MLPPP Bundle

Purpose Verify that the constituent links are added to the bundle correctly.

Action From operational mode, enter the **show interfaces lsq-*fpc/pic/port*** command.

Meaning The output displays the constituent links that are added to the bundle. For more information about the **show interfaces lsq-*fpc/pic/port*** operational command, see the Junos OS Operational Mode Commands.

Copyright © 2013, Juniper Networks, Inc.

Verifying the MLFR FRF.15 Configuration

Purpose Verify the MLFR FRF.15 configuration.

Action From operational mode, enter the **show interfaces lsq-*fpc/pic/port*** command.

Meaning The output displays the standard status information about the specified link services IQ interface. For more information about the **show interfaces lsq-*fpc/pic/port*** operational command, see the Junos OS Operational Mode Commands.

Verifying the MLFR FRF.16 Configuration

Purpose Verify the MLFR FRF.16 configuration.

Action From operational mode, enter the **show interfaces lsq-*fpc/pic/port*** command.

Meaning The output displays the standard status information about the specified link services IQ interface. For more information about the **show interfaces lsq-*fpc/pic/port*** operational command, see the Junos OS Operational Mode Commands.

Related Documentation

- Link and Multilink Services Overview on page 1123
- Multilink Interfaces on Channelized MICs Overview on page 1131
- Example: Configuring an MLPPP Bundle on page 1169
- Example: Configuring Multilink Frame Relay FRF.15 on page 1172
- Example: Configuring Multilink Frame Relay FRF.16 on page 1175

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 56

Summary of Multilink and Link Services Configuration Statements

The following sections explain each of the multilink and link services statements. The statements are organized alphabetically.

acknowledge-retries

Syntax acknowledge-retries *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, configure the number of retransmission attempts to be made for consecutive hello or remove link messages following the expiration of the acknowledgment timer.

Options *number*—Number of retransmission attempts to be made following the expiration of the acknowledgment timer.

Range: 1 through 5

Default: 2

Usage Guidelines See "Configuring Acknowledgment Timers on Link Services Physical Interfaces" on page 1148.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- action-red-differential-delay on page 1191, hello-timer on page 1200

Copyright © 2013, Juniper Networks, Inc.

acknowledge-timer

Syntax acknowledge-timer *milliseconds*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, configure the maximum time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.

Options ***milliseconds***—Time to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.

Range: 1 through 10 milliseconds

Default: 4 milliseconds

Usage Guidelines See "Configuring Acknowledgment Timers on Link Services Physical Interfaces" on page 1148.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- address (Interfaces) on page 1192, hello-timer on page 1200

 Copyright © 2013, Juniper Networks, Inc.

action-red-differential-delay

Syntax action-red-differential-delay (disable-tx | remove-link);

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, configure the action to be taken when the differential delay exceeds the red limit.

Options **disable-tx**—Disable transmission on the bundle link.

remove-link—Remove the bundle link from service.

Default: remove-link

Usage Guidelines See "Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16" on page 1149.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- yellow-differential-delay on page 1211

Copyright © 2013, Juniper Networks, Inc.

address (Interfaces)

Syntax

```
address address {
  destination address;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the interface address.

Options ***address***—Address of the interface.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the Links in a Multilink or Link Services Bundle" on page 1135; for a general discussion of **address** statement options, see the Junos® OS Network Interfaces.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

bundle

Syntax

```
bundle (ml-fpc/pic/port | ls-fpc/pic/port);
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family mlfr-end-to-end],
[edit interfaces *interface-name* unit *logical-unit-number* family mlfr-uni-nni]

Release Information Statement introduced before Junos OS Release 7.4.

Description Associate the multilink interface with the logical interface it is joining.

Options **ml-*fpc/pic/port***—Name of the multilink interface you are linking.

ls-*fpc/pic/port*—Name of the link services interface you are linking.

Usage Guidelines See "Configuring the Links in a Multilink or Link Services Bundle" on page 1135.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination (Interfaces)

Syntax destination *address*;

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number* tunnel]
[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],
[edit interfaces *interface-name* unit *logical-unit-number* tunnel]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For CoS on ATM interfaces, specify the remote address of the connection.

For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.

For tunnel and encryption interfaces, specify the remote address of the tunnel.

Options ***address***—Address of the remote side of the connection.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Linear RED Profiles on ATM Interfaces
- Multilink and Link Services Logical Interface Configuration Overview on page 1136
- Configuring Encryption Interfaces on page 881
- Configuring Traffic Sampling on page 912
- Configuring Flow Monitoring on page 920
- Configuring Unicast Tunnels on page 1273

Copyright © 2013, Juniper Networks, Inc.

disable-mlppp-inner-ppp-pfc

Syntax disable-mlppp-inner-ppp-pfc;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For MLPPP interfaces only, disable compression of the inner PPP header in the MLPPP payload. By default, compression is enabled.

Usage Guidelines See "Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces" on page 1139.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

dlci

Syntax dlci *dlci-identifier*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Frame Relay and Multilink Frame Relay user-to-network interface (UNI) network-to-network interface (NNI) encapsulation only, and for link services and point-to-point interfaces only, configure the data-link connection identifier (DLCI) for a permanent virtual circuit (PVC) or a switched virtual circuit (SVC).

To configure a DLCI for a point-to-multipoint interface, use the **multipoint-destination** statement to specify the DLCI.

Options *dlci-identifier*—Data-link connection identifier.

Range: 16 through 1022

Usage Guidelines See "Configuring DLCIs on Link Services Logical Interfaces" on page 1143; for general information about Frame Relay DLCIs, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

drop-timeout

Syntax drop-timeout *milliseconds*;

Hierarchy Level
[edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit interfaces (ls-*fpc/pic/port* | ml-*fpc/pic/port*) unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit logical-systems *logical-system-name* interfaces (ls-*fpc/pic/port*| ml-*fpc/pic/port*) unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink and link services interfaces only, configure the drop timeout period, in milliseconds.

Options ***milliseconds***—Drop timeout period.

Range: 1 through 2000 milliseconds

Default: 500 ms for bundles greater than or equal to the T1 bandwidth value, and 1500 ms for other bundles. Any CLI-configured value overrides these defaults. Setting a value of 0 reverts to the default.

Usage Guidelines See "Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces" on page 1139.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

encapsulation

See the following sections:

- encapsulation (Logical Interface) on page 1196
- encapsulation (Physical Interface) on page 1197

encapsulation (Logical Interface)

Syntax encapsulation (atm-mlppp-llc | multilink-frame-relay-end-to-end | multilink-ppp | ...);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Logical link-layer encapsulation type.

Options

atm-mlppp-llc—For ATM 2 interfaces, use Multilink Point-to-Point Protocol (MLPPP) over ATM Adaptation Layer 5 (AAL5) logical link control (LLC) encapsulation, as described in RFC 2364, *PPP over AAL5*.

multilink-frame-relay-end-to-end—Use Multilink Frame Relay (MLFR) FRF.15 encapsulation. This encapsulation is usedon multilink "link services interfaces and their constituent T1 or E1 interfaces", and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink and link services interfaces and their constituent T1 or E1 interfaces.

Usage Guidelines See "Configuring Encapsulation for Multilink and Link Services Logical Interfaces" on page 1138; for information about **encapsulation** statement options used with other interface types, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

encapsulation (Physical Interface)

Syntax encapsulation (multilink-frame-relay-uni-nni | ...);

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces rlsq*number:number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Physical link-layer encapsulation type.

Default MLFR UNI NNI encapsulation (on link services interfaces).

Options **multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used only on link services interfaces functioning as FRF.16 bundles and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

Usage Guidelines See "Configuring Encapsulation for Link Services Physical Interfaces" on page 1148; for information about **encapsulation** statement options used with other interface types, see the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

family

Syntax

```
family family {
  address address {
    destination address;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure protocol family information for the logical interface.

Options *family*—Protocol family:

- **ccc**—Circuit cross-connect protocol suite
- **inet**—IP version 4 (IPv4)
- **inet6**—IP version 6 (IPv6)
- **iso**—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—MPLS
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpls**—Virtual private LAN service

The remaining statements are explained separately.

Usage Guidelines See the topics in Link and Multilink Properties; for a general discussion of **family** statement options, see the Junos® OS Network Interfaces .

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

fragment-threshold

Syntax fragment-threshold *bytes*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit interfaces (ls-*fpc/pic/port*| ml-*fpc/pic/port*) unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit logical-systems *logical-system-name* interfaces (ls-*fpc/pic/port*| ml-*fpc/pic/port*) unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink and link services interfaces only, set the fragmentation threshold, in bytes.

Options ***bytes***—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.

Range: 128 through 16,320 bytes

Default: 0 bytes (no fragmentation)

Usage Guidelines See "Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces" on page 1140.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

hello-timer

Syntax hello-timer *milliseconds*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, configure the rate at which hello messages are sent. A hello message is transmitted after a period defined in milliseconds has elapsed.

Options ***milliseconds***—The rate at which hello messages are sent.
Range: 1 through 180 milliseconds
Default: 10 milliseconds

Usage Guidelines See "Configuring Acknowledgment Timers on Link Services Physical Interfaces" on page 1148.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- address (Interfaces) on page 1192, acknowledge-timer on page 1190

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines See the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

interleave-fragments

Syntax interleave-fragments;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces ls-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, interleave long packets with high-priority packets.

Allows small delay-sensitive packets, such as voice over IP (VoIP) packets, to interleave with long fragmented packets. This minimizes the latency of delay-sensitive packets.

Usage Guidelines See "Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces" on page 1144.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

lmi-type

Syntax lmi-type (ansi | itu);

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the Frame Relay Local Management Interface (LMI) type.

Options **ansi**—Use American National Standards Institute (ANSI) T1.167 Annex D LMIs.

itu—Use ITU Q933 Annex A LMIs.
Default: itu

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

minimum-links

Syntax minimum-links *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit interfaces (ls-*fpc/pic/port* | ml-*fpc/pic/port*) unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit logical-systems *logical-system-name* interfaces (ls-*fpc/pic/port* | ml-*fpc/pic/port*) unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink or link services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up. A member link is considered up when the PPP Link Control Protocol (LCP) phase transitions to open state.

The **minimum-links** value should be identical on both ends of the bundle.

Options *number*—Number of links.

Range: 1 through 8

Default: 1

Usage Guidelines See "Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces" on page 1141.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

mlfr-uni-nni-bundle-options

Syntax

```
mlfr-uni-nni-bundle-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  cisco-interoperability send-lip-remove-link-for-link-reject;
  drop-timeout milliseconds;
  fragment-threshold bytes;
  hello-timer milliseconds;
  lmi-type (ansi | itu | c-lmi);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 number;
  t392 number;
  yellow-differential-delay milliseconds;
}
```

Hierarchy Level [edit interfaces ls-*fpc/pic/port* :*channel*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure link services interface management properties.

The statements are explained separately.

Usage Guidelines See "Configuring Encapsulation for Link Services Physical Interfaces" on page 1148.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Encapsulation for Link Services Physical Interfaces on page 1148

Copyright © 2013, Juniper Networks, Inc.

mrru

Syntax mrru *bytes*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit interfaces (ml-*fpc/pic/port*| ls-*fpc/pic/port*) unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options],
[edit logical-systems *logical-system-name* interfaces (ml-*fpc/pic/port*| ls-*fpc/pic/port*) unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink or link services interfaces only, set the maximum received reconstructed unit (MRRU). The MRRU is similar to the maximum transmission unit (MTU), but is specific to multilink interfaces.

Options *bytes*—MRRU size.

Range: 1500 through 4500 bytes

Default: 1500 bytes

Usage Guidelines See "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

mtu

Syntax mtu *bytes*;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values.

Options ***bytes***—MTU size.

Range: 0 through 5012 bytes

Default: 1500 bytes (**inet**, **inet6**, and **iso** families), 1448 bytes (**mpls**)

Usage Guidelines See "Configuring MRRU on Multilink and Link Services Logical Interfaces" on page 1141.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces

multicast-dlci

Syntax multicast-dlci *dlci-identifier*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For point-to-multipoint link services interfaces only, enable multicast support on the interface. You can configure multicast support on the interface if the Frame Relay switch performs multicast replication.

Options ***dlci-identifier***—DLCI identifier, a number from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

Usage Guidelines See "Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles" on page 1143.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

n391

Syntax n391 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, set the Frame Relay full status polling interval.

Options *number*—Polling interval.
Range: 1 through 255
Default: 6

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- n392 on page 1206, n393 on page 1207, t391 on page 1208, and t392 on page 1209

n392

Syntax n392 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, set the Frame Relay error threshold, in number of errors.

Options *number*—Error threshold.
Range: 1 through 10
Default: 3

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- n391 on page 1206, n393 on page 1207, t391 on page 1208, and t392 on page 1209

 Copyright © 2013, Juniper Networks, Inc.

n393

Syntax n393 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, set the Frame Relay monitored event count.

Options *number*—Event count.
Range: 1 through 10
Default: 4

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- n391 on page 1206, n392 on page 1206, t391 on page 1208, and t392 on page 1209

red-differential-delay

Syntax red-differential-delay *milliseconds*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, configure the red differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.

Options *milliseconds*—Red differential delay threshold.
Range: 1 through 2000 milliseconds
Default: 120 milliseconds

Usage Guidelines See "Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16" on page 1149.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- action-red-differential-delay on page 1191, yellow-differential-delay on page 1211

Copyright © 2013, Juniper Networks, Inc.

short-sequence

Syntax short-sequence;

Hierarchy Level [edit interfaces (ls-*fpc/pic/port* | ml-*fpc/pic/port*) unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces (ls-*fpc/pic/port* | ml-*fpc/pic/port*) unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink interfaces only, set the length of the packet sequence identification number to 12 bits.

Default If not included in the configuration, the length is set to 24 bits.

Usage Guidelines See "Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces" on page 1142.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

t391

Syntax t391 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, set the Frame Relay link integrity polling interval.

Options *number*—Link integrity polling interval.
Range: 5 through 30 seconds
Default: 10 seconds

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- n391 on page 1206, n392 on page 1206, n393 on page 1207, and t392 on page 1209

Copyright © 2013, Juniper Networks, Inc.

t392

Syntax t392 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, set the Frame Relay polling verification interval.

Options *number*—Polling verification interval.
Range: 5 through 30 seconds
Default: 15 seconds

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 1150.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- n391 on page 1206, n392 on page 1206, n393 on page 1207, and t391 on page 1208

Copyright © 2013, Juniper Networks, Inc.

unit (Interfaces)

Syntax

```
unit logical-unit-number {
  disable-mlppp-inner-ppp-pfc;
  dlci dlci-identifier;
  drop-timeout milliseconds;
  encapsulation type;
  fragment-threshold bytes;
  interleave-fragments;
  minimum-links number;
  mrru bytes;
  multicast-dlci dlci-identifier;
  short-sequence;
  family family {
    address address {
      destination address;
    }
    bundle (ml-fpc/pic/port | ls-fpc/pic/port);
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Usage Guidelines See Link and Multilink Properties; for a general discussion of logical interface properties, see the Junos® OS Network Interfaces .

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

yellow-differential-delay

Syntax yellow-differential-delay *milliseconds*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, configure the yellow differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.

Options ***milliseconds***—Yellow differential delay threshold.

Range: 1 through 2000 milliseconds

Default: 72 milliseconds

Usage Guidelines See "Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16" on page 1149.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- action-red-differential-delay on page 1191, red-differential-delay on page 1207

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 7

Real-Time Performance Monitoring Services

- Real-Time Performance Monitoring Services Overview on page 1215
- Real-Time Performance Monitoring Configuration Guidelines on page 1217
- Summary of Real-Time Performance Monitoring Configuration Statements on page 1237

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 57

Real-Time Performance Monitoring Services Overview

Real-Time Performance Monitoring Services Overview

Real-Time Performance Monitoring (RPM) enables you to configure active probes to track and monitor traffic. Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets. RPM provides Management Information Base (MIB) support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

You can also configure RPM services to determine automatically whether a path exists between a host router and its configured BGP neighbors. You can view the results of the discovery using an SNMP client. Results are stored in **pingResultsTable**, **jnxPingResultsTable**, **jnxPingProbeHistoryTable**, and **pingProbeHistoryTable**.

Probe configuration and probe results are supported by the command-line interface (CLI) and SNMP.

The following probe types are supported with DSCP marking:

- ICMP echo
- ICMP timestamp
- HTTP get (not available for BGP RPM services)
- UDP echo
- TCP connection
- UDP timestamp

With probes, you can monitor the following:

- Minimum round-trip time
- Maximum round-trip time
- Average round-trip time

Copyright © 2013, Juniper Networks, Inc.

- Standard deviation of the round-trip time
- Jitter of the round-trip time—The difference between the minimum and maximum round-trip time

One-way measurements for ICMP timestamp probes include the following:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probes sent
- Number of probe responses received
- Percentage of lost probes

You can configure the following RPM thresholds:

- Round-trip time
- Ingress/egress delay
- Standard deviation
- Jitter
- Successive lost probes
- Total lost probes (per test)

Support is also implemented for user-configured CoS classifiers and for prioritization of RPM packets over regular data packets received on an input interface.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 58

Real-Time Performance Monitoring Configuration Guidelines

To configure Real-Time Performance Monitoring (RPM) services, include the **rpm** statement at the **[edit services]** hierarchy level:

```
[edit services]
rpm {
  bgp {
    data-fill data;
    data-size size;
    destination-port port;
    history-size size;
    logical-system logical-system-name [routing-instances routing-instance-name];
    moving-average-size number;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instances instance-name;
    test-interval interval;
  }
  probe owner {
    test test-name {
      data-fill data;
      data-size size;
      destination-interface interface-name;
      destination-port port;
      dscp-code-point dscp-bits;
      hardware-timestamp;
      history-size size;
      moving-average-size number;
      one-way-hardware-timestamp;
      probe-count count;
      probe-interval seconds;
      probe-type type;
      routing-instance instance-name;
      source-address address;
      target (url url | address address);
      test-interval interval;
      thresholds thresholds;
      traps traps;
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
    probe-server {
      tcp {
        destination-interface interface-name;
        port number;
      }
      udp {
        destination-interface interface-name;
        port number;
      }
    }
    probe-limit limit;
    twamp {
      server {
        authentication-mode (authenticated | encrypted | none);
        client-list list-name {
          [ address address ];
        }
        inactivity-timeout seconds;
        maximum-connections-duration hours;
        maximum-connections count;
        maximum-connections-per-client count;
        maximum-sessions count;
        maximum-sessions-per-connection count;
        port number;
        server-inactivity-timeout minutes;
      }
    }
  }
```



NOTE: RPM does not require an Adaptive Services (AS) or Multiservices PIC or Multiservices Dense Port Concentrator (DPC) unless you are configuring RPM timestamping as described in "Configuring RPM Timestamping" on page 1225.

This chapter includes the following sections:

- Configuring BGP Neighbor Discovery Through RPM on page 1218
- Configuring Real-Time Performance Monitoring on page 1220
- Enabling RPM for the Services SDK on page 1230
- Examples: Configuring BGP Neighbor Discovery Through RPM on page 1231
- Examples: Configuring Real-Time Performance Monitoring on page 1232

Configuring BGP Neighbor Discovery Through RPM

BGP neighbors can be configured at the following hierarchy levels:

Copyright © 2013, Juniper Networks, Inc.

- **[edit protocols bgp group *group-name*]**—Default logical system and default routing instance.
- **[edit routing-instances *instance-name* protocols bgp group *group-name*]**—Default logical system with a specified routing instance.
- **[edit logical-systems *logical-system-name* protocols bgp group *group-name*]**—Configured logical system and default routing instance.
- **[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name*]**—Configured logical system with a specified routing instance.

When you configure BGP neighbor discovery through RPM, if you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. You can explicitly configure RPM probes to apply only to the default logical system, the default routing instance, or to a particular logical system or routing instance.

To configure BGP neighbor discovery through RPM, configure the probe properties at the **[edit services rpm bgp]** hierarchy:

```
data-fill data;
data-size size;
destination-port port;
history-size size;
logical-system logical-system-name [routing-instances routing-instance-name];
moving-average-size number;
probe-count count;
probe-interval seconds;
probe-type type;
routing-instances instance-name;
test-interval interval;
```

- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the **data-fill** statement at the **[edit services rpm bgp]** hierarchy level. The value can be a hexadecimal value.
- To specify the size of the data portion of ICMP probes, include the **data-size** statement at the **[edit services rpm bgp]** hierarchy level. The size can be from **0** through **65507** and the default size is **0**.
- To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the **destination-port** statement at the **[edit services rpm bgp]** hierarchy level. The **destination-port** statement is used only for the UDP and TCP probe types. The value can be **7** or from **49160** through **65535**.
- To specify the number of stored history entries, include the **history-size** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from **0** to **255**. The default is **50**.
- To specify the logical system used by ICMP probes, include the **logical-system *logical-system-name*** statement at the **[edit services rpm bgp]** hierarchy level. If you do not specify a logical system, the RPM probe applies to configured BGP neighbors for

Copyright © 2013, Juniper Networks, Inc.

all logical systems. To apply the probe to only the default logical system, you must set the value of ***logical-system-name*** to **null**.

- To specify a number of samples for making statistical calculations, include the moving-average-size statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the **probe-count** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from **1** through **15**.
- To specify the time to wait between sending packets, include the **probe-interval** statement at the **[edit services rpm bgp]** hierarchy level. Specify a value from **1** through **255** seconds.
- To specify the packet and protocol contents of the probe, include the **probe-type** statement at the **[edit services rpm bgp]** hierarchy level. The following probe types are supported:
  - **icmp-ping**—Sends ICMP echo requests to a target address.
  - **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.
  - **tcp-ping**—Sends TCP packets to a target.
  - **udp-ping**—Sends UDP packets to a target.
  - **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.

  

  NOTE: Some probe types require additional parameters to be configured. For example, when you specify the **tcp-ping** or **udp-ping** option, you must configure the destination port using the **destination-port** *port* statement. The **udp-ping-timestamp** option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

- To specify the routing instance used by ICMP probes, include the **routing-instances** statement at the **[edit services rpm bgp]** hierarchy level. The default routing instance is Internet routing table **inet.0**. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. To apply the RPM probe to only the default routing instance, you must explicitly set the value of ***instance-name*** to **default**.
- To specify the time to wait between tests, include the **test-interval** statement at the **[edit services bgp probe]** hierarchy level. Specify a value from **0** through **86400** seconds.

Configuring Real-Time Performance Monitoring

This section describes the following tasks for configuring RPM:

- Configuring RPM Probes on page 1221
- Configuring RPM Receiver Servers on page 1224
- Limiting the Number of Concurrent RPM Probes on page 1225

Copyright © 2013, Juniper Networks, Inc.

- Configuring RPM Timestamping on page 1225
- Configuring TWAMP on page 1228

Configuring RPM Probes

The owner name and test name identifiers of an RPM probe together represent a single RPM configuration instance. When you specify the test name, you also can configure the test parameters.

To configure the probe owner, test name, and test parameters, include the **probe** statement at the **[edit services rpm]** hierarchy level:

```
probe owner {
  test test-name {
    data-fill data;
    data-size size;
    destination-interface interface-name;
    destination-port port;
    dscp-code-point dscp-bits;
    hardware-timestamp;
    history-size size;
    moving-average-size number;
    one-way-hardware-timestamp;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    source-address address;
    target (url url | address address);
    test-interval interval;
    thresholds thresholds;
    traps traps;
  }
}
```

- To specify a probe owner, include the **probe** statement at the **[edit services rpm]** hierarchy level. The probe owner identifier can be up to 32 characters in length.
- To specify a test name, include the **test** statement at the **[edit services rpm probe *owner*]** hierarchy level. The test name identifier can be up to 32 characters in length. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.
- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the **data-fill** statement at the **[edit services rpm probe *owner*]** hierarchy level. The value can be a hexadecimal value. The **data-fill** statement is not valid with the **http-get** or **http-metadata-get** probe types.
- To specify the size of the data portion of ICMP probes, include the **data-size** statement at the **[edit services rpm probe *owner*]** hierarchy level. The size can be from **0** through **65507** and the default size is **0**. The **data-size** statement is not valid with the **http-get** or **http-metadata-get** probe types.

Copyright © 2013, Juniper Networks, Inc.



NOTE: If you configure the hardware timestamp feature (see "Configuring RPM Timestamping" on page 1225), the **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.

- On M Series and T Series routers, you configure the **destination-interface** statement to enable hardware timestamping of RPM probe packets. You specify an **sp-** interface to have the AS or Multiservices PIC add the hardware timestamps; for more information, see "Configuring RPM Timestamping" on page 1225. You can also include the **one-way-hardware-timestamp** statement to enable one-way delay and jitter measurements.
- To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the **destination-port** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. The **destination-port** statement is used only for the UDP and TCP probe types. The value can be **7** or from **49160** through **65535**.
- To specify the value of the Differentiated Services (DiffServ) field within the IP header, include the **dscp-code-point** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. The DiffServ code point (DSCP) bits value can be set to a valid 6-bit pattern; for example, **001111**. It also can be set using an alias configured at the **[edit class-of-service code-point-aliases dscp]** hierarchy level. The default is **000000**.
- To specify the number of stored history entries, include the **history-size** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. Specify a value from **0** to **255**. The default is **50**.
- To specify a number of samples for making statistical calculations, include the **moving-average-size** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. Specify a value from **0** through **255**.
- To specify the number of probes within a test, include the **probe-count** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. Specify a value from **1** through **15**.
- To specify the time to wait between sending packets, include the **probe-interval** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. Specify a value from **1** through **255** seconds.
- To specify the packet and protocol contents of the probe, include the **probe-type** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. The following probe types are supported:
  - **http-get**—Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.
  - **http-metadata-get**—Sends an HTTP get request for metadata to a target URL.
  - **icmp-ping**—Sends ICMP echo requests to a target address.
  - **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.

Copyright © 2013, Juniper Networks, Inc.

- **tcp-ping**—Sends TCP packets to a target.
- **udp-ping**—Sends UDP packets to a target.
- **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.

The following probe types support hardware timestamping of probe packets: **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, **udp-ping-timestamp**.



NOTE: Some probe types require additional parameters to be configured. For example, when you specify the **tcp-ping** or **udp-ping** option, you must configure the destination port using the **destination-port** statement. The **udp-ping-timestamp** option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

- To specify the routing instance used by ICMP probes, include the **routing-instance** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. The default routing instance is Internet routing table **inet.0**.
- To specify the source IP address used for ICMP probes, include the **source-address** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. If the source IP address is not one of the router's assigned addresses, the packet will use the outgoing interface's address as its source.
- To specify the destination address used for the probes, include the **target** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level.
  - For HTTP probe types, specify a fully formed URL that includes **http://** in the URL address.
  - For all other probe types, specify an IP version 4 (IPv4) address for the target host.
- To specify the time to wait between tests, include the **test-interval** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. Specify a value from **0** through **86400** seconds.
- To specify thresholds used for the probes, include the **thresholds** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded. The following options are supported:
  - **egress-time**—Measures maximum source-to-destination time per probe.
  - **ingress-time**—Measures maximum destination-to-source time per probe.
  - **jitter-egress**—Measures maximum source-to-destination jitter per test.
  - **jitter-ingress**—Measures maximum destination-to-source jitter per test.
  - **jitter-rtt**—Measures maximum jitter per test, from 0 through 60000000 microseconds.
  - **rtt**—Measures maximum round-trip time per probe, in microseconds.

Copyright © 2013, Juniper Networks, Inc.

- **std-dev-egress**—Measures maximum source-to-destination standard deviation per test.
- **std-dev-ingress**—Measures maximum destination-to-source standard deviation per test.
- **std-dev-rtt**—Measures maximum standard deviation per test, in microseconds.
- **successive-loss**—Measures successive probe loss count, indicating probe failure.
- **total-loss**—Measures total probe loss count indicating test failure, from 0 through 15.

- Traps are sent if the configured threshold is met or exceeded. To set the trap bit to generate traps, include the **traps** statement at the **[edit services rpm probe *owner* test *test-name*]** hierarchy level. The following options are supported:
  - **egress-jitter-exceeded**—Generates traps when the jitter in egress time threshold is met or exceeded.
  - **egress-std-dev-exceeded**—Generates traps when the egress time standard deviation threshold is met or exceeded.
  - **egress-time-exceeded**—Generates traps when the maximum egress time threshold is met or exceeded.
  - **ingress-jitter-exceeded**—Generates traps when the jitter in ingress time threshold is met or exceeded.
  - **ingress-std-dev-exceeded**—Generates traps when the ingress time standard deviation threshold is met or exceeded.
  - **ingress-time-exceeded**—Generates traps when the maximum ingress time threshold is met or exceeded.
  - **jitter-exceeded**—Generates traps when the jitter in round-trip time threshold is met or exceeded.
  - **probe-failure**—Generates traps for successive probe loss thresholds crossed.
  - **rtt-exceeded**—Generates traps when the maximum round-trip time threshold is met or exceeded.
  - **std-dev-exceeded**—Generates traps when the round-trip time standard deviation threshold is met or exceeded.
  - **test-completion**—Generates traps when a test is completed.
  - **test-failure**—Generates traps when the total probe loss threshold is met or exceeded.

Configuring RPM Receiver Servers

The RPM TCP and UDP probes are proprietary to Juniper Networks and require a receiver to receive the probes. To configure a server to receive the probes, include the **probe-server** statement at the **[edit services rpm]** hierarchy level:

```
[edit services rpm]
probe-server {
```


Copyright © 2013, Juniper Networks, Inc.

```
    tcp {
      destination-interface interface-name;
      port number;
    }
    udp {
      port number;
    }
  }
```

The port number specified for the UDP and TCP server can be **7** or from **49160** through **65535**.

Limiting the Number of Concurrent RPM Probes

To configure the maximum number of concurrent probes allowed, include the **probe-limit** statement at the **[edit services rpm]** hierarchy level:

```
probe-limit limit;
```

Specify a limit from **1** through **500**. The default maximum number is 100.

Configuring RPM Timestamping

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You can timestamp the following RPM probe types: **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp**.

On M Series and T Series routers with an Adaptive Services (AS) or Multiservices PIC, on MX Series routers with a Multiservices DPC and on EX Series switches, you can enable hardware timestamping of RPM probe messages. The timestamp is applied on both the RPM client (the router or switch that originates the RPM probes) and the RPM probe server and applies only to IPv4 traffic. It is supported in the Layer 2 service package on all Multiservices PICs and DPCs and in the Layer 3 service package on AS and Multiservices PICs and Multiservices DPCs.

To configure two-way timestamping on M Series and T Series routers, include the **destination-interface** statement at the **[edit services rpm probe** *probe-owner* **test** *test-name***]** hierarchy level:

```
destination-interface sp-fpc/pic/port.logical-unit-number;
```

Specify the RPM client router and the RPM server router on the adaptive services logical interface by including the **rpm** statement at the **[edit interfaces** *interface-name* **unit** *logical-unit-number***]** hierarchy level:

```
rpm (client | server);
```

The logical interface must be dedicated to the RPM task. It requires configuration of the **family inet** statement and a **/32** address, as shown in the example. This configuration is also needed for other services such as NAT and stateful firewall. You cannot configure RPM service on **unit 0** because RPM requires a dedicated logical interface; the same unit cannot support both RPM and other services. Because active flow monitoring requires **unit 0**, but RPM can function on any logical interface, a constraint check prevents you from committing an RPM configuration there.

Copyright © 2013, Juniper Networks, Inc.

NOTE: If you configure RPM timestamping on an AS PIC, you cannot configure the source-address statement at the **[edit services rpm probe *probe-name* test *test-name*]** hierarchy level.

On MX Series routers and EX Series switches, you include the **hardware-timestamp** statement at the **[edit services rpm probe *probe-name* test *test-name*]** hierarchy level to specify that the probes are to be timestamped in the Packet Forwarding Engine host processor:

```
hardware-timestamp;
```

On the client side, these probes are timestamped in the Packet Forwarding Engine host processor on the egress DPC on the MX Series router or EX Series switch originating the RPM probes (RPM client). On the responder side (RPM server), the RPM probes to be timestamped are handled by Packet Forwarding Engine host processor, which generates the response instead of the RPM process. The RPM probes are timestamped only on the router that originates them (RPM client). As a result, only round-trip time is measured for these probes.

NOTE: The Packet Forwarding Engine based RPM feature does not support any stateful firewall configurations. If you need to combine RPM timestamping with stateful firewall, you should use the interface-based RPM timestamping service described earlier in this section. Multiservices DPCs support stateful firewall processing as well as RPM timestamping.

To configure one-way timestamping, you must also include the **one-way-hardware-timestamp** statement at the **[edit services rpm probe *probe-owner* test *test-name*]** hierarchy level:

```
one-way-hardware-timestamp;
```


Copyright © 2013, Juniper Networks, Inc.

NOTE: If you configure RPM probes for a services interface (sp-), you need to announce local routes in a specific way for the following routing protocols:

- For OSPF, you can announce the local route by including the services interface in the OSPF area. To configure this setting, include the **interface sp-*fpc/pic/port*** statement at the **[edit protocols ospf area *area-number*]** hierarchy level.
- For BGP and IS-IS, you must export interface routes and create a policy that accepts the services interface local route. To export interface routes, include the **point-to-point** and **lan** statements at the **[edit routing-options interface-routes family inet export]** hierarchy level. To configure an export policy that accepts the services interface local route, include the **protocol local**, **rib inet.0**, and **route-filter *sp-interface-ip-address*/32 exact** statements at the **[edit policy-options policy-statement *policy-name* term *term-name* from]** hierarchy level and the **accept** action at the **[edit policy-options policy-statement *policy-name* term *term-name* then]** hierarchy level. For the export policy to take effect, apply the policy to BGP or IS-IS with the **export *policy-name*** statement at the **[edit protocols *protocol-name*]** hierarchy level.

For more information about these configurations, see the Routing Policy Configuration Guide or the Junos OS Routing Protocols Configuration Guide.

Example: Configuring RPM Timestamping

Routing the probe packets through the AS or Multiservices PIC also enables you to filter the probe packets to particular queues. The following example shows the RPM configuration and the filter that specifies queuing:

```
services rpm {
   probe p1 {
      test t1 {
         probe-type icmp-ping;
         target address 10.8.4.1;
         probe-count 10;
         probe-interval 10;
         test-interval 10;
         dscp-code-points af11;
         data-size 100;
         destination-interface sp-1/2/0.0;
      }
   }
}
firewall {
   filter f1 {
      term t1 {
         from {
            dscp af11;
         }
         then {
            forwarding-class assured-forwarding;
         }
      }
   }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
interfaces sp-1/2/0 {
    unit 2 {
        rpm client;
        family inet {
            address 10.8.4.2/32;
            filter {
                input f1;
            }
        }
    }
}
interfaces sp-1/2/1 {
    unit 2 {
        rpm server;
        family inet {
            address 10.8.3.2/32;
            filter {
                input f1;
            }
        }
    }
}
```

For more information about firewall filters, see the Routing Policy Configuration Guide; for more information about queuing, see the Junos OS Class of Service Configuration Guide.

Configuring TWAMP

You can configure the Two-Way Active Measurement Protocol (TWAMP) on on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), and on MX Series routers with or without a Multiservices DPC. Only the responder (server) side of TWAMP is supported.

For more information on TWAMP, see RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)*.

To configure TWAMP properties, include the **twamp** statement at the **[edit services rpm]** hierarchy level:

```
[edit services rpm]
twamp {
    server {
        client-list list-name {
            [ address address ];
        }
        authentication-mode mode;
        max-connection-duration hours;
        maximum-connections count;
        maximum-connections-per-client count;
        maximum-sessions count;
        maximum-sessions-per-connection count;
        port number;
        server-inactivity-timeout minutes;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
}
```

The TWAMP configuration process includes the following tasks:

- Configuring TWAMP Interfaces on page 1229
- Configuring TWAMP Servers on page 1229

Configuring TWAMP Interfaces

To specify the service PIC logical interface that provides the TWAMP service, include the **twamp-server** statement at the **[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number***] hierarchy level:

```
twamp-server;
```



NOTE: On MX Series routers that do not include a Multiservices DPC, you can configure TWAMP properties, but you can omit specifying the **twamp-server** statement.

Configuring TWAMP Servers

You can specify a number of TWAMP server properties, some of which are optional, by including the **server** statement at the **[edit services rpm twamp]** hierarchy level:

```
[edit services rpm twamp]
server {
    client-list list-name {
        [ address address ];
    }
    authentication-mode mode;
    max-connection-duration hours;
    maximum-connections count;
    maximum-connections-per-client count;
    maximum-sessions count;
    maximum-sessions-per-connection count;
    port number;
    server-inactivity-timeout minutes;
}
```

- To specify the list of allowed control client hosts that can connect to this server, include the **client-list** statement at the **[edit services rpm twamp server]** hierarchy level. Each value you include must be a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can include multiple client lists, each of which can contain a maximum of 64 entries. You must configure at least one client address to enable TWAMP.
- You must specify the authentication mode by including the **authentication-mode** statement at the **[edit services rpm twamp server]** hierarchy level. There is no default value. You can configure **authenticated** or **encrypted** mode, based on RFC 4656; if there is no authentication or encryptions mode specified, you should set the value to **none**. This statement is required in the TWAMP configuration.

Copyright © 2013, Juniper Networks, Inc.

- To specify the inactivity timeout period in seconds, include the **inactivity-timeout** statement at the **[edit services rpm twamp server]** hierarchy level. By default, the value is **1800**; the range is 0 through 3600 seconds.
- To specify the maximum number of concurrent connections the server can have to client hosts, include the **maximum-connections** statement at the **[edit services rpm twamp server]** hierarchy level. The allowed range of values is 1 through 2048 and the default value is 64. You can also limit the number of connections the server can make to a particular client host by including the **maximum-connections-per-client** statement.
- To specify the maximum number of sessions the server can have running at one time, include the **maximum-sessions** statement at the **[edit services rpm twamp server]** hierarchy level. The allowed range of values is 1 through 2048 and the default value is 64. You can also limit the number of sessions the server can have on a single connection by including the **maximum-sessions-per-connection** statement.
- To specify the TWAMP server listening port, include the **port** statement at the **[edit services rpm twamp server]** hierarchy level. The range is 1 through 65,535. This statement is mandatory.
- To specify the server inactivity timeout period in minutes, include the **server-inactivity-timeout** statement at the **[edit services rpm twamp server]** hierarchy level. The range is 0 through 30 minutes.

For examples of TWAMP configuration, see "Examples: Configuring Real-Time Performance Monitoring" on page 1232.

Enabling RPM for the Services SDK

Real-time performance monitoring (RPM), which has been supported on the adaptive services interface, is now supported by the Services SDK. RPM is supported on all platforms and service PICs that support the Services SDK.

To enable RPM for the Services SDK on the adaptive services interface, configure the **object-cache-size**, **policy-db-size**, and **package** statements at the **[edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider]** hierarchy level. For the Services SDK, ***package-name*** in the **package *package-name*** statement is **jservices-rpm**.

For more information about the Services SDK, see the *SDK Applications Configuration Guide and Command Reference*.

The following example shows how to enable RPM for the Services SDK on the adaptive services interface:

```
chassis fpc 1 {
  pic 2 {
    adaptive-services {
      service-package {
        extension-provider {
          control-cores 1;
          data-cores 1;
          object-cache-size 512;
```


Copyright © 2013, Juniper Networks, Inc.

```
            policy-db-size 64;
            package jservices-rpm;
            syslog daemon any;
        }
      }
    }
  }
}
```

Related Documentation

- Examples: Configuring Real-Time Performance Monitoring on page 1232
- destination-interface on page 1241

Examples: Configuring BGP Neighbor Discovery Through RPM

Configure BGP neighbor discovery through RPM for all logical systems and all routing instances:

```
[edit services rpm]
bgp {
  probe-type icmp-ping;
  probe-count 5;
  probe-interval 1;
  test-interval 60;
  history-size 10;
  data-size 255;
  data-fill 0123456789;
}
```

Configure BGP neighbor discovery through RPM for only the following logical systems and routing instances: **LS1/RI1**, **LS1/RI2**, **LS2**, and **RI3**:

```
[edit services rpm]
bgp {
  probe-type icmp-ping;
  probe-count 5;
  probe-interval 1;
  test-interval 60;
  history-size 10;
  data-size 255;
  data-fill 0123456789;
  logical-system {
    LS1 {
      routing-instances {
        RI1;
        RI2;
      }
    }
    LS2;
  }
  routing-instance {
    RI3;
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

Configure BGP neighbor discovery through RPM for only the default logical system and default routing instance:

```
[edit services rpm]
bgp {
   probe-type icmp-ping;
   probe-count 5;
   probe-interval 1;
   test-interval 60;
   history-size 10;
   data-size 255;
   data-fill 0123456789;
   logical-system {
      null {
         routing-instances {
            default;
         }
      }
   }
}
```

Examples: Configuring Real-Time Performance Monitoring

Configure an RPM instance identified by the probe name **probe1** and the test name **test1**:

```
[edit services rpm]
probe probe1{
   test test1 {
      dscp-code-points 001111;
      probe-interval 1;
      probe-type icmp-ping;
      target address 172.17.20.182;
      test-interval 20;
      thresholds rtt 10;
      traps rtt-exceeded;
   }
}
probe-server {
   tcp {
      destination-interface lt-0/0/0.0
      port 50000;
   }
   udp {
      destination-interface lt-0/0/0.0
      port 50001;
   }
}
probe-limit 200;
```

Configure packet classification, using **lt-** interfaces to send the probe packets to a logical tunnel input interface. By sending the packet to the logical tunnel interface, you can configure regular and multifield classifiers, firewall filters, and header rewriting for the probe packets. To use the existing tunnel framework, the **dlci** and **encapsulation** statements must be configured.

```
[edit services rpm]
```


Copyright © 2013, Juniper Networks, Inc.

```
probe p1 {
  test t1 {
    probe-type icmp-ping;
    target address 10.8.4.1;
    probe-count 10;
    probe-interval 10;
    test-interval 10;
    source-address 10.8.4.2;
    dscp-code-points ef;
    data-size 100;
    destination-interface lt-0/0/0.0;
  }
}
[edit interfaces]
lt-0/0/0 {
  unit 0 {
    encapsulation frame-relay;
    dlci 10;
    peer-unit 1;
    family inet;
  }
  unit 1 {
    encapsulation frame-relay;
    dlci 10;
    peer-unit 0;
    family inet;
  }
}
[edit class-of-service]
interfaces {
  lt-0/0/0 {
    unit 1 {
      classifiers {
        dscp default;
      }
    }
  }
}
```

Configure an input filter on the interface on which the RPM probes are received. This filter enables prioritization of the received RPM packets, separating them from the regular data packets received on the same interface.

```
[edit firewall]
filter recos {
  term recos {
    from {
      source-address {
        10.8.4.1/32;
      }
      destination-address {
        10.8.4.2/32;
      }
    }
    then {
      loss-priority high;
```

Copyright © 2013, Juniper Networks, Inc.

```
        forwarding-class network-control;
      }
    }
  }
  [edit interfaces]
  fe-5/0/0 {
    unit 0 {
      family inet {
        filter {
          input recos;
        }
        address 10.8.4.2/24;
      }
    }
  }
```

Configure an RPM instance and enable RPM for the Services SDK on the adaptive services interface:

```
  [edit services rpm]
  probe probe1{
    test test1 {
      data-size 1024;
      data-fill 0;
      destination-interface ms-1/2/0.10;
      dscp-code-points 001111;
      probe-count 10;
      probe-interval 1;
      probe-type icmp-ping;
      target address 172.17.20.182;
      test-interval 20;
      thresholds rtt 10;
      traps rtt-exceeded;
    }
  }
  [edit interfaces]
  ms-1/2/0 {
    unit 0 {
      family inet;
    }
    unit 10 {
      rpm client;
      family inet {
        address 1.1.1.1/32;
      }
  }
  [edit chassis]
  fpc 1 {
    pic 2 {
      adaptive-services {
        service-package {
          extension-provider {
            control-cores 1;
            data-cores 1;
            object-cache-size 512;
            policy-db-size 64;
```


Copyright © 2013, Juniper Networks, Inc.

```
            package jservices-rpm;
            syslog {
              daemon any;
            }
          }
        }
      }
    }
```

Configure the minimum statements necessary to enable TWAMP:

```
[edit services]
rpm {
  twamp {
    server {
      authentication-mode none;
      port 10000; # Twamp server's listening port
      client-list LIST-1 { # LIST-1 is the name of the client-list. Multiple lists can be
        configured.
        address {
          20.0.0.2/30; # IP address of the control client.
        }
      }
    }
  }
[edit interfaces sp-5/0/0]
unit 0 {
  family inet;
}
unit 10 {
  rpm {
    twamp-server; # You must configure a separate logical interface on the service PIC
      interface for the TWAMP server.
  }
  family inet {
    address 50.50.50.50/32; # This address must be a host address with a 32-bit mask.
  }
}
[edit chassis]
fpc 5 {
  pic 0 {
    adaptive-services {
      service-package layer-2;     # Configure the service PIC to run in Layer 2 mode.
    }
  }
}
```

Configure additional TWAMP settings:

```
[edit services]
rpm {
  twamp {
    server {
      maximum-sessions 5;
      maximum-sessions-per-connection 2;
      maximum-connections 3;
```

Copyright © 2013, Juniper Networks, Inc.

```
        maximum-connections-per-client 1;
        port 10000;
        server-inactivity-timeout ;
        client-list LIST-1 {
          address {
            20.0.0.2/30;
          }
        }
      }
    }
  }
```


Copyright © 2013, Juniper Networks, Inc.

CHAPTER 59

Summary of Real-Time Performance Monitoring Configuration Statements

The following sections explain each of the Real-Time Performance Monitoring (RPM) statements. The statements are organized alphabetically.

authentication-mode

Syntax authentication-mode (authenticated | control-only-encrypted | encrypted | none);

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the authentication or encryption mode support for the TWAMP test protocol. This statement is required in the configuration; if no authentication or encryption is specified, you should set the value to **none**.

Options

authenticated—Data packets are authenticated.

control-only-encrypted—TWAMP control packets are encrypted. TWAMP data packets are in plain text format.

encrypted—Data packets are encrypted.

none—No authentication or encryption.

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

bgp

Syntax

```
bgp {
  data-fill data;
  data-size size;
  destination-port port;
  history-size size;
  logical-system logical-system-name <routing-instances routing-instance-name>;
  moving-average-size size;
  probe-count count;
  probe-interval seconds;
  probe-type type;
  routing-instances instance-name;
  test-interval interval;
}
```

Hierarchy Level

[edit services rpm bgp]
[edit protocols bgp group *group-name*]
[edit routing-instances *instance-name* protocols bgp group *group-name*]
[edit logical-system *logical-system-name* protocols bgp group *group-name*]
[edit logical-system *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure BGP neighbor discovery through Real-Time Performance Monitoring (RPM).

Options

bgp—Define properties for configuring BGP neighbor discovery.

The remaining statements are explained separately.

NOTE: On MX Series routers, you can configure all the statements. On M Series and T Series routers, you can configure only the **logical-system** and **routing-instances** statements.

Usage Guidelines

See "Configuring BGP Neighbor Discovery Through RPM" on page 1218.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

client-list

Syntax
```
client-list list-name {
    address address;
}
```

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description List of allowed control client hosts that can connect to this server. Each entry is a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can configure more than one list, but you must configure at least one client address to enable TWAMP. Each list can contain up to 64 entries.

Options ***list-name***—Name of client address list.

address—Address and mask for an allowed client.

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

data-fill

Syntax data-fill *data*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes.

Options ***data***—A hexadecimal value; for example, **0-9**, **A-F**.

Usage Guidelines The **data-fill** statement is not valid with the **http-get** or **http-metadata-get** probe types. See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

data-size

Syntax data-size *size*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the size of the data portion of ICMP probes.

Options *data*—The size can be from 0 through 65507

Default: 0



NOTE: If you configure the hardware timestamp feature (see "Configuring RPM Timestamping" on page 1225), the **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.

Usage Guidelines The **data-size** statement is not valid with the **http-get** or **http-metadata-get** probe type. See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

destination-interface

Syntax destination-interface *interface-name*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*],
[edit services rpm probe-server (tcp | udp)]

Release Information Statement introduced in Junos OS Release 7.5.

Description On M Series and T Series routers, specify a services (**sp-**) interface that adds a timestamp to RPM probe messages. This feature is supported only with **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp** probe types. You must also configure the **rpm** statement on the **sp-** interface and include the **unit 0 family inet** statement with a **/32** address.

On M Series, MX Series, and T Series routers, specify a multiservices (**ms-**) interface that adds a timestamp to RPM probe messages. This feature is supported only with **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp** probe types. You must also configure the **rpm** statement on the **ms-** interface and include the **unit 0 family inet** statement with a **/32** address.

To enable RPM for the Services SDK on the adaptive services interface, configure the **object-cache-size**, **policy-db-size**, and **package** statements at the **[edit chassis fpc** *slot-number* **pic** *pic-number* **adaptive-services service-package extension-provider]** hierarchy level. For the Services SDK, ***package-name*** in the **package** ***package-name*** statement is **jservices-rpm**.

Options *interface-name*—Name of the adaptive services interface.

Usage Guidelines See "Configuring RPM Probes" on page 1221, "Configuring RPM Receiver Servers" on page 1224, or "Configuring RPM Timestamping" on page 1225.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- hardware-timestamp on page 1244
- rpm
- Enabling RPM for the Services SDK on page 1230

Copyright © 2013, Juniper Networks, Inc.

destination-port

Syntax destination-port *port*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to which a probe is sent. This statement is used only for TCP or UDP probe types.

Options *port*—The port number can be **7** or from **49,160** to **65,535**.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

dscp-code-point

Syntax dscp-code-point *dscp-bits*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.

Options *dscp-bits*—A valid 6-bit pattern; for example, **001111**, or one of the following configured DSCP aliases:

- **af11**—Default: 001010
- **af12**—Default: 001100
- **af13**—Default: 001110
- **af21**—Default: 010010
- **af22**—Default: 010100
- **af23** —Default: 010110
- **af31** —Default: 011010
- **af32** —Default: 011100
- **af33** —Default: 011110
- **af41** —Default: 100010
- **af42** —Default:100100
- **af43** —Default:100110
- **be**—Default: 000000
- **cs1**—Default: 001000
- **cs2**—Default: 010000
- **cs3**—Default: 011000
- **cs4**—Default: 100000
- **cs5**—Default: 101000
- **cs6**—Default: 110000
- **cs7**—Default: 111000
- **ef**—Default: 101110
- **nc1**—Default: 110000
- **nc2**—Default: 111000

Copyright © 2013, Juniper Networks, Inc.

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

hardware-timestamp

Syntax hardware-timestamp;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced in Junos OS Release 8.1.
Statement applied to MX Series routers in Junos OS Release 10.0.
Statement introduced in Junos OS Release 10.3 for EX Series switches.

Description On MX Series routers and EX Series switches only, enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp** probe types.

Usage Guidelines See "Configuring RPM Timestamping" on page 1225.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

history-size

Syntax history-size *size*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the number of stored history entries.

Options *size*—A value from 0 to 255.
Default: 50

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

inactivity-timeout (Services RPM)

Syntax inactivity-timeout *seconds*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description Inactivity timeout period, in seconds.

Options ***seconds***—Length of time the session is inactive before it times out.
Default: 1800 seconds

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

logical-system

Syntax
```
logical-system logical-system-name {
    [ routing-instances instance-name ];
}
```

Hierarchy Level [edit services rpm bgp]

Release Information Statement introduced in Junos OS Release 7.6.

Description Specify the logical system used by the probes.

The remaining statements are explained separately.

Options ***logical-system-name***—Logical system name.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

max-connection-duration

Syntax max-connection-duration *hours*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 11.1.

Description Specify the maximum time a connection can exist between a client and the server.

Options ***hours***—Number of hours a connection can exist between a client and the server.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring TWAMP on page 1228

maximum-connections

Syntax maximum-connections *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description Maximum number of allowed connections between the server and all control client hosts.

Options ***count***—Maximum number of connections.
Range: 1 through 2048
Default: 64

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

maximum-connections-per-client

Syntax maximum-connections-per-client *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description Maximum number of allowed connections between the server and a single control client host.

Options ***count***—Maximum number of connections.
Default: 64

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

maximum-sessions

Syntax maximum-sessions *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description Maximum number of allowed test sessions the server can have running at one time.

Options ***count***—Maximum number of sessions.
Range: 1 through 2048
Default: 64

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

maximum-sessions-per-connection

Syntax maximum-sessions-per-connection *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description Maximum number of allowed sessions the server can open on a single client connection.

Options *count*—Maximum number of sessions.
Default: 64

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

moving-average-size

Syntax moving-average-size *number*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Enable statistical calculation operations to be performed across a configurable number of the most recent samples.

Options *number*—Number of samples to be used in calculations.
Range: 0 through 255

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

one-way-hardware-timestamp

Syntax one-way-hardware-timestamp;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Enable timestamping of RPM probe messages for one-way delay and jitter measurements. You must configure this statement along with the **destination-interface** statement to invoke timestamping. This feature is supported only with **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp** probe types.

Usage Guidelines See "Configuring RPM Timestamping" on page 1225.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- destination-interface on page 1241, hardware-timestamp on page 1244

Copyright © 2013, Juniper Networks, Inc.

port

See the following sections:

- port (RPM) on page 1250
- port (TWAMP) on page 1250

port (RPM)

Syntax port *number*;

Hierarchy Level [edit services rpm probe-server (tcp | udp)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the port number for the probe server.

Options ***number***—Port number for the probe server. The value can be **7** or **49,160** through **65,535**.

Usage Guidelines See "Configuring RPM Receiver Servers" on page 1224.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

port (TWAMP)

Syntax port *number*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 9.3.

Description TWAMP server listening port. You must configure this statement to enable TWAMP.

Options ***number***—Port number.
Range: 1 through 65,535

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

probe

Syntax

```
probe owner {
  test test-name {
    data-fill data;
    data-size size;
    destination-interface interface-name;
    destination-port port;
    dscp-code-point dscp-bits;
    hardware-timestamp;
    history-size size;
    moving-average-size number;
    one-way-hardware-timestamp;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    source-address address;
    target (url | address);
    test-interval interval;
    thresholds thresholds;
    traps traps;
  }
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.

Options *owner*—Specify an owner name up to 32 characters in length.

The remaining statements are explained separately.

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

probe-count

Syntax probe-count *count*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the number of probes within a test.

Options ***count***—A value from 1 through 15.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

probe-interval

Syntax probe-interval *interval*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the time to wait between sending packets, in seconds.

Options ***interval***—Number of seconds, from 1 through 255.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

probe-limit

Syntax probe-limit *limit;*

Hierarchy Level [edit services rpm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the maximum number of concurrent probes allowed.

Options ***limit***—A value from 1 through 500.
Default: 100.

Usage Guidelines See "Limiting the Number of Concurrent RPM Probes" on page 1225.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

probe-server

Syntax

```
probe-server {
  tcp {
    destination-interface interface-name;
    port number;
  }
  udp {
    destination-interface interface-name;
    port number;
  }
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the server to act as a receiver for the probes.

The remaining statements are explained separately.

Usage Guidelines See "Configuring RPM Receiver Servers" on page 1224.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

probe-type

Syntax probe-type *type*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the packet and protocol contents of a probe.

Options *type*—Specify one of the following probe type values:

- **http-get**—(Not available at the [edit services rpm bgp] hierarchy level.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.
- **http-metadata-get**—(Not available at the [edit services rpm bgp] hierarchy level.) Sends an HTTP get request for metadata to a target URL.
- **icmp-ping**—Sends ICMP echo requests to a target address.
- **icmp-ping-timestamp**—Sends ICMP timestamp requests to a target address.
- **tcp-ping**—Sends TCP packets to a target.
- **udp-ping**—Sends UDP packets to a target.
- **udp-ping-timestamp**—Sends UDP timestamp requests to a target address.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

routing-instance

Syntax routing-instance *instance-name*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the routing instance used by the probes.

Options ***instance-name***—A routing instance configured at the **[edit routing-instance]** hierarchy level.

Default: Internet routing table **inet.0**.

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

routing-instances

Syntax routing-instances *instance-name*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm bgp logical-system *logical-system-name*]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the routing instance used by the probes.

Options ***instance-name***—A routing instance configured at the **[edit routing-instances]** hierarchy level.

Default: Internet routing table **inet.0**.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

rpm

Syntax

```
rpm {
  bgp {
    data-fill data;
    data-size size;
    destination-port port;
    history-size size;
    logical-system logical-system-name <routing-instances routing-instance-name>;
    moving-average-size number;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instances instance-name;
    test-interval interval;
  }
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure BGP neighbor discovery through RPM.

The remaining statements are explained separately.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

server

Syntax

```
server {
  client-list list-name {
    [ address address ];
  }
  inactivity-timeout seconds;
  maximum-connections count;
  maximum-connections-per-client count;
  maximum-sessions count;
  maximum-sessions-per-connection count;
  port number;
}
```

Hierarchy Level [edit services rpm twamp]

Release Information Statement introduced in Junos OS Release 9.3.

Description TWAMP server configuration settings.

Options The remaining statements are described separately.

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

server-inactivity-timeout

Syntax server-inactivity-timeout *minutes*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 11.1.

Description The maximum time the Two-Way Active Measurement Protocol (TWAMP) server has to finish the TWAMP control protocol negotiation.

Options *minutes*—Number of minutes the TWAMP server has to finish the TWAMP control protocol negotiation.

Default: 15 minutes

Range: 1-30 minutes

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

services (RPM)

Syntax services rpm { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service rules to be applied to traffic.

Options **rpm**—Identifies the RPM set of rules statements.

Usage Guidelines See Real-Time Performance Monitoring Services.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

source-address

Syntax source-address *address*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet will use the outgoing interface's address as its source.

Options ***address***—Valid IP address.

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

target (Services RPM)

Syntax target (url *url* | address *address*);

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the destination address or URL used for the probes.

Options **url *url***—For HTTP probe types, specify a fully formed URL that includes **http://** in the URL address.

address *address*—For all other probe types, specify an IPv4 address for the target host.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring RPM Probes on page 1221
- Configuring Basic Real-Time Performance Monitoring (RPM) Probes in a VDCF

tcp

Syntax

```
tcp {
  destination-interface interface-name;
  port port;
}
```

Hierarchy Level [edit services rpm probe-server]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the port information for the TCP server.

The remaining statements are explained separately.

Usage Guidelines See "Configuring RPM Receiver Servers" on page 1224.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

test

Syntax

```
test test-name {
  data-fill data;
  data-size size;
  destination-interface interface-name;
  destination-port port;
  dscp-code-point dscp-bits;
  hardware-timestamp;
  history-size size;
  moving-average-size number;
  one-way-hardware-timestamp;
  probe-count count;
  probe-interval seconds;
  probe-type type;
  routing-instance instance-name;
  source-address address;
  target (url url | address address);
  test-interval interval;
  thresholds thresholds;
  traps traps;
}
```

Hierarchy Level [edit services rpm probe *owner*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the range of probes over which the standard deviation, average, and jitter are calculated. The test name combined with the owner name represent a single RPM configuration instance.

Options ***test-name***—Specify a test name. The name can be up to 32 characters in length.

The remaining statements are explained separately.

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

test-interval

Syntax test-interval *frequency*;

Hierarchy Level [edit services rpm bgp],
[edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the time to wait between tests, in seconds.

Options ***frequency***—Number of seconds, from 0 through 86400.

Usage Guidelines See "Configuring BGP Neighbor Discovery Through RPM" on page 1218 or "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

thresholds

Syntax thresholds *thresholds*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify thresholds used for the probes. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded.

Options *thresholds*—Specify one or more threshold measurements. The following options are supported:

- **egress-time**—Measures maximum source-to-destination time per probe.
- **ingress-time**—Measures maximum destination-to-source time per probe.
- **jitter-egress**—Measures maximum source-to-destination jitter per test.
- **jitter-ingress**—Measures maximum destination-to- source jitter per test.
- **jitter-rtt**—Measures maximum jitter per test, from 0 through 60,000,000 microseconds.
- **rtt**—Measures maximum round-trip time per probe, in microseconds.
- **std-dev-egress**—Measures maximum source-to-destination standard deviation per test.
- **std-dev-ingress**—Measures maximum destination-to-source standard deviation per test.
- **std-dev-rtt**—Measures maximum standard deviation per test, in microseconds.
- **successive-loss**—Measures successive probe loss count, indicating probe failure.
- **total-loss**—Measures total probe loss count indicating test failure, from 0 through 15.

Usage Guidelines See "Configuring RPM Probes" on page 1221.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

traps

Syntax traps *traps*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Set the trap bit to generate traps for probes. Traps are sent if the configured threshold is met or exceeded.

Options *traps*—Specify one or more traps. The following options are supported:

- **egress-jitter-exceeded**—Generates traps when the jitter in egress time threshold is met or exceeded.
- **egress-std-dev-exceeded**—Generates traps when the egress time standard deviation threshold is met or exceeded.
- **egress-time-exceeded**—Generates traps when the maximum egress time threshold is met or exceeded.
- **ingress-jitter-exceeded**—Generates traps when the jitter in ingress time threshold is met or exceeded.
- **ingress-std-dev-exceeded**—Generates traps when the ingress time standard deviation threshold is met or exceeded.
- **ingress-time-exceeded**—Generates traps when the maximum ingress time threshold is met or exceeded.
- **jitter-exceeded**—Generates traps when the jitter in round-trip time threshold is met or exceeded.
- **probe-failure**—Generates traps for successive probe loss thresholds crossed.
- **rtt-exceeded**—Generates traps when the maximum round-trip time threshold is met or exceeded.
- **std-dev-exceeded**—Generates traps when the round-trip time standard deviation threshold is met or exceeded.
- **test-completion**—Generates traps when a test is completed.
- **test-failure**—Generates traps when the total probe loss threshold is met or exceeded.



NOTE: For RPM traps to be generated, you must configure the **remote-operations** SNMP trap category by including the **categories** statement at the **[edit snmp trap-group** *trap-group-name* hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring RPM Probes on page 1221
- categories

twamp

Syntax

```
twamp {
  server {
    authentication-mode mode;
    client-list list-name {
      [ address address ];
    }
    inactivity-timeout seconds;
    max-connection-duration hours;
    maximum-connections count;
    maximum-connections-per-client count;
    maximum-sessions count;
    maximum-sessions-per-connection count;
    port number;
    server-inactivity-timeout minutes;
  }
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced in Junos OS Release 9.3.

Description Two-Way Active Measurement Protocol (TWAMP) configuration settings.

The remaining statements are described separately.

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level
system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

twamp-server

Syntax twamp-server;

Hierarchy Level [edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the service PIC logical interface to provide the TWAMP service.

Usage Guidelines See "Configuring TWAMP" on page 1228.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

udp

Syntax

```
udp {
  destination-interface interface-name;
  port port;
}
```

Hierarchy Level [edit services rpm probe-server]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify the port information for the UDP server.

The remaining statements are explained separately.

Usage Guidelines See "Configuring RPM Receiver Servers" on page 1224.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 8

Tunnel Services

- Tunnel Services Overview on page 1269
- Tunnel Interfaces Configuration Guidelines on page 1273
- Summary of Tunnel Services Configuration Statements on page 1295

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 60

Tunnel Services Overview

This chapter discusses the following topics:

- Tunnel Services Overview on page 1269
- GRE Keepalive Time Overview on page 1271

Tunnel Services Overview

By encapsulating arbitrary packets inside a transport protocol, tunneling provides a private, secure path through an otherwise public network. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and MPLS. If you have a Tunnel Physical Interface Card (PIC) installed in your M Series or T Series router, you can configure unicast, multicast, and logical tunnels.

You can configure two types of tunnels for VPNs: one to facilitate routing table lookups and another to facilitate VPN routing and forwarding instance (VRF) table lookups.

For information about encryption interfaces, see "Configuring Encryption Interfaces" on page 881 and the Junos OS System Basics Configuration Guide. For information about VPNs, see the Junos OS VPNs Configuration Guide. For information about MPLS, see the Junos OS MPLS Applications Configuration Guide.

On SRX Series and J Series devices, Generic Routing Encapsulation (GRE) and IP-IP tunnels use internal interfaces, gr-0/0/0 and ip-0/0/0, respectively. The Junos OS creates these interfaces at system bootup; they are not associated with physical interfaces.

The Juniper Networks Junos OS supports the tunnel types shown in Table 25 on page 1269.

Table 25: Tunnel Interface Types

Interface	Description
gr-0/0/0	Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol over another routing protocol. Within a router, packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then re-encapsulated with another protocol packet to complete the GRE. The GRE interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform GRE.

Copyright © 2013, Juniper Networks, Inc.

Table 25: Tunnel Interface Types *(continued)*

Interface	Description
gre	Internally generated GRE interface. This interface is generated by the Junos OS to handle GRE. NOTE: You can configure GRE interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. This type of interface does not require a Tunnel PIC. For more information about GMPLS, see the Junos OS MPLS Applications Configuration Guide and the Junos OS Feature Guides.
ip-0/0/0	Configurable IP-over-IP encapsulation (also called IP tunneling) interface. IP tunneling allows the encapsulation of one IP packet over another IP packet. Packets are routed to an internal interface where they are encapsulated with an IP packet and then forwarded to the encapsulating packet's destination address. The IP-IP interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform IP tunneling.
ipip	Internally generated IP-over-IP interface. This interface is generated by the Junos OS to handle IP-over-IP encapsulation. It is not a configurable interface.
lt-0/0/0	The **lt** interface on M Series and T Series routers supports configuration of logical systems—the capability to partition a single physical router into multiple logical devices that perform independent routing tasks. On SRX Series devices, the **lt** interface is a configurable logical tunnel interface that interconnects logical systems. See the *Junos OS Logical Systems Configuration Guide for Security Devices*. On J Series devices, the **lt** interface is used to provide class-of-service (CoS) support for real-time performance monitoring (RPM) probe packets. Packets are routed to this internal interface for services. The **lt** interface is an internal interface only; it is not associated with a physical interface. You must configure the interface for it to perform CoS for RPM services. See the *Junos OS Class of Service Configuration Guide for Security Devices*.
mt-0/0/0	Internally generated multicast tunnel interface. Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a **224/8**-or-greater prefix, the packet is dropped and a counter is incremented. Within a router, packets are routed to this internal interface for multicast filtering. The multicast tunnel interface is an internal interface only and is not associated with a physical interface. If your router has a Tunnel Services PIC, the Junos OS automatically configures one multicast tunnel interface (**mt-**) for each virtual private network (VPN) you configure. You do not need to configure multicast tunnel interfaces. However, you can configure properties on **mt-** interfaces, such as the **multicast-only** statement.
mtun	Internally generated multicast tunnel interface. This interface is generated by the Junos OS to handle multicast tunnel services. It is not a configurable interface.

Copyright © 2013, Juniper Networks, Inc.

Table 25: Tunnel Interface Types *(continued)*

Interface	Description
pd-0/0/0	Configurable Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree. Within a router, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform PIM de-encapsulation. NOTE: On SRX Series devices, this interface type is **ppd0**.
pe-0/0/0	Configurable PIM encapsulation interface. In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree. Within a router, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform PIM encapsulation. NOTE: On SRX Series devices, this interface type is **ppe0**.
pimd	Internally generated PIM de-encapsulation interface. This interface is generated by the Junos OS to handle PIM de-encapsulation. It is not a configurable interface.
pime	Internally generated PIM encapsulation interface. This interface is generated by the Junos OS to handle PIM encapsulation. It is not a configurable interface.
vt-0/0/0	Configurable virtual loopback tunnel interface. Facilitates VRF table lookup based on MPLS labels. This interface type is supported on M Series and T Series routers, but not on SRX Series or J Series devices. To configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels, you specify a virtual loopback tunnel interface name and associate it with a routing instance that belongs to a particular routing table. The packet loops back through the virtual loopback tunnel for route lookup.

GRE Keepalive Time Overview

Generic routing encapsulation (GRE) tunnel interfaces do not have a built-in mechanism for detecting when a tunnel is down. You can enable keepalive messages to serve as the detection mechanism.

Keepalives can be configured on the physical or on the logical interface. If configured on the physical interface, keepalives are sent on all logical interfaces that are part of the physical interface. If configured on a individual logical interface, keepalives are only sent

Copyright © 2013, Juniper Networks, Inc.

to that logical interface. In addition to configuring a keepalive, you must configure the hold time.

Related Documentation

- Configuring GRE Keepalive Time on page 1279
- Example: Configuring Keepalive for a GRE Interface on page 1293
- **keepalive-time on page 1302**
- **hold-time (OAM) on page 1301**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 61

Tunnel Interfaces Configuration Guidelines

This chapter includes the following tunnel interface configuration tasks and examples:

- Configuring Unicast Tunnels on page 1273
- Configuring GRE Keepalive Time on page 1279
- Restricting Tunnels to Multicast Traffic on page 1280
- Configuring Logical Tunnel Interfaces on page 1281
- Configuring Tunnel Interfaces for Routing Table Lookup on page 1282
- Configuring Virtual Loopback Tunnels for VRF Table Lookup on page 1283
- Configuring PIM Tunnels on page 1284
- Configuring IPv6-over-IPv4 Tunnels on page 1285
- Configuring IPv4-over-IPv6 Tunnels on page 1285
- Configuring Dynamic Tunnels on page 1286
- Configuring Tunnel Interfaces on MX Series Routers on page 1286
- Configuring Tunnel Interfaces on T4000 Routers on page 1287
- Examples: Configuring Unicast Tunnels on page 1288
- Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 1289
- Example: Configuring an IPv6-over-IPv4 Tunnel on page 1290
- Example: Configuring an IPv4-over-IPv6 Tunnel on page 1291
- Example: Configuring Logical Tunnels on page 1292
- Example: Configuring Keepalive for a GRE Interface on page 1293

Configuring Unicast Tunnels

To configure a unicast tunnel, you configure a **gr-** interface (to use GRE encapsulation) or an **ip-** interface (to use IP-IP encapsulation) and include the **tunnel** and **family** statements:

```
gr-fpc/pic/port or ip-fpc/pic/port {
  unit logical-unit-number {
    copy-tos-to-outer-ip-header;
    reassemble-packets;
    tunnel {
      allow-fragmentation;
```

Copyright © 2013, Juniper Networks, Inc.

```
        backup-destination address;
        destination destination-address;
        do-not-fragment;
        key number;
        routing-instance {
          destination routing-instance-name;
        }
        source address;
        ttl number;
      }
      family family {
        address address {
          destination address;
        }
      }
    }
  }
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

You can configure multiple logical units for each GRE or IP-IP interface, and you can configure only one tunnel per unit.



NOTE: On M Series and T Series routers, you can configure the interface on a service PIC or a tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC.

Each tunnel interface must be a point-to-point interface. Point to point is the default interface connection type, so you do not need to include the **point-to-point** statement in the logical interface configuration.

You must specify the tunnel's destination and source addresses. The remaining statements are optional.

Copyright © 2013, Juniper Networks, Inc.



NOTE: For transit packets exiting the tunnel, forwarding path features, such as reverse path forwarding (RPF), forwarding table filtering, source class usage, destination class usage, and stateless firewall filtering, are not supported on the interfaces you configure as tunnel sources, but are supported on tunnel-pic interfaces.

However, class-of-service (CoS) information obtained from the GRE or IP-IP header is carried over the tunnel and is used by the re-entering packets. For more information, see the Junos OS Class of Service Configuration Guide.

To prevent an invalid configuration, the Junos OS disallows setting the address specified by the **source** or **destination** statement at the **[edit interfaces** gr-*fpc/pic/port* unit *logical-unit-number* **tunnel]** hierarchy level to be the same as the interface's own subnet address, specified by the **address** statement at the **[edit interfaces** gr-*fpc/pic/port* unit *logical-unit-number* **family** *family-name***]** hierarchy level.

To set the time-to-live (TTL) field that is included in the encapsulating header, include the **ttl** statement. If you explicitly configure a TTL value for the tunnel, you must configure it to be one larger than the number of hops in the tunnel. For example, if the tunnel has seven hops, you must configure a TTL value of 8.

You must configure at least one family on the logical interface. To enable MPLS over GRE tunnel interfaces, you must include the **family mpls** statement in the GRE interface configuration. In addition, you must include the appropriate statements at the **[edit protocols]** hierarchy level to enable Resource Reservation Protocol (RSVP), MPLS, and label-switched paths (LSPs) over GRE tunnels. Unicast tunnels are bidirectional.

A configured tunnel cannot go through Network Address Translation (NAT) at any point along the way to the destination. For more information, see "Examples: Configuring Unicast Tunnels" on page 1288 and the Junos OS MPLS Applications Configuration Guide.

For a GRE tunnel, the default is to set the ToS bits in the outer IP header to all zeros. To have the Routing Engine copy the ToS bits from the inner IP header to the outer, include the **copy-tos-bits-to-outer-ip-header** statement. (This inner-to-outer ToS bits copying is already the default behavior for IP-IP tunnels.)

For GRE tunnel interfaces on Adaptive Services or Multiservices interfaces, you can configure additional tunnel attributes, as described in the following sections:

- Configuring a Key Number on GRE Tunnels on page 1276
- Enabling Fragmentation on GRE Tunnels on page 1276
- Specifying an MTU Setting for the Tunnel on page 1277
- Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 1277
- Configuring Packet Reassembly on page 1278

Copyright © 2013, Juniper Networks, Inc.

Configuring a Key Number on GRE Tunnels

For Adaptive Services and Multiservices interfaces on M Series and T Series routers, you can assign a key value to identify an individual traffic flow within a GRE tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*. However, only one key is allowed for each tunnel source and destination pair.

Each IP version 4 (IPv4) packet entering the tunnel is encapsulated with the GRE tunnel key value. Each IPv4 packet exiting the tunnel is verified by the GRE tunnel key value and de-encapsulated. The Adaptive Services or Multiservices PIC drops packets that do not match the configured key value.

To assign a key value to a GRE tunnel interface, include the **key** statement:

```
key number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* tunnel]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]**

The key number can be 0 through 4,294,967,295. You must configure the same GRE tunnel key value on tunnel endpoints.

The following example illustrates the use of the key statement in a GRE tunnel configuration:

```
interfaces {
    gr-1/2/0 {
        unit 0 {
            tunnel {
                source 10.58.255.193;
                destination 10.58.255.195;
                key 1234;
            }
            ...
            family inet {
                mtu 1500;
                address 10.200.0.1/30;
                ...
            }
        }
    }
}
```

Enabling Fragmentation on GRE Tunnels

For GRE tunnel interfaces on Adaptive Services and Multiservices interfaces only, you can enable fragmentation of IPv4 packets in GRE tunnels.

Copyright © 2013, Juniper Networks, Inc.

By default, IPv4 traffic transmitted over GRE tunnels is not fragmented. To enable fragmentation of IPv4 packets in GRE tunnels, include the **clear-dont-fragment-bit** statement:

clear-dont-fragment-bit;

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

When you include the **clear-dont-fragment-bit** statement in the configuration, the don't-fragment (DF) bit is cleared on all packets, even packets that do not exceed the tunnel maximum transmission unit (MTU). If the packet's size exceeds the tunnel's MTU value, the packet is fragmented before encapsulation. If the packet's size does not exceed the tunnel's MTU value, the packet is not fragmented.



NOTE: The Packet Forwarding Engine updates the IP identification field in the outer IP header of GRE-encapsulated packets, so that reassembly of the packets is possible after fragmentation. The previous CLI constraint check that required you to configure either the **clear-dont-fragment-bit** statement or a tunnel key with the **allow-fragmentation** statement is no longer enforced.

You can also clear the DF bit in packets transmitted over IP Security (IPsec) tunnels. For more information, see "Enabling IPsec Packet Fragmentation" on page 247.

Specifying an MTU Setting for the Tunnel

To enable key numbers and fragmentation on GRE tunnels (as described in "Configuring a Key Number on GRE Tunnels" on page 1276 and "Enabling Fragmentation on GRE Tunnels" on page 1276), you must also specify an MTU setting for the tunnel.

To specify an MTU setting for the tunnel, include the **mtu** statement:

mtu *bytes*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces gr-***fpc/pic/port* **unit** *logical-unit-number* **family inet]**
- **[edit logical-system** *logical-system-name* **interfaces gr-***fpc/pic/port* **unit** *logical-unit-number* **family inet]**

For more information about MTU settings, see the Junos® OS Network Interfaces.

Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

Unlike IP-IP tunnels, GRE tunnels do not copy the ToS bits to the outer IP header by default. To have the Routing Engine copy the inner ToS bits to the outer IP header (which is required for some tunneled routing protocols) on packets sent by the Routing Engine, include the **copy-tos-to-outer-ip-header** statement at the logical unit hierarchy level of

Copyright © 2013, Juniper Networks, Inc.

a GRE interface. This example copies the inner ToS bits to the outer IP header on a GRE tunnel:

```
[edit interfaces]
gr-0/0/0 {
  unit 0 {
    copy-tos-to-outer-ip-header;
    family inet;
  }
}
```

Configuring Packet Reassembly

On GRE tunnel interfaces only, you can enable reassembly of fragmented tunnel packets. To activate this capability, include the **reassemble-packets** statement:

```
reassemble-packets;
```

You can configure this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

For each tunnel you configure on the interface, you can enable or disable fragmentation of GRE packets by including the **allow-fragmentation** or **do-not-fragment** statement:

```
allow-fragmentation;
do-not-fragment;
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number* **tunnel]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **tunnel]**

If you configure **allow-fragmentation** on a tunnel, it clears the DF bit in the outer IP header, enabling post fragmentation of GRE-encapsulated packets if the packet size exceeds the maximum transmission unit (MTU) value for the egress interface. By default, packets that exceed the MTU size are dropped and post fragmentation of GRE packets is disabled.



NOTE: Whenever you configure **allow-fragmentation** on a tunnel, you must also include either the **tunnel key** or the **clear-dont-fragment-bit** statement. This configuration enables the router to send affected packets to the PIC so that the correct IP header can be placed in the fragments. Otherwise, on the reassembly side some packets might be lost when fragments arrive in the PIC out of sequence at high speeds.

Copyright © 2013, Juniper Networks, Inc.

Configuring GRE Keepalive Time

You can configure the keepalives on a GRE tunnel interface by including both the **keepalive-time** statement and the **hold-time** statement at the **[edit protocols oam gre-tunnel interface *interface-name*]** hierarchy level.



NOTE: For proper operation of keepalives on a GRE interface, you must also include the **family inet** statement at the **[edit interfaces *interface-name* unit *unit*]** hierarchy level. If you do not include this statement, the interface is marked as down.

To configure keepalive time for a GRE tunnel interface:

1. At the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level, set the **family** as **inet**.

   user@host# **set interfaces *interface-name* unit *unit-number* family *family-name***

2. Configure the Operation, Administration, and Maintenance (OAM) protocol:

   [edit]
   user@host# **edit protocols oam**

3. Configure the GRE tunnel interface:

   [edit protocols oam]
   user@host# **edit gre-tunnel interface *interface-name***

4. Configure the keepalive time:

   [edit protocols oam gre-tunnel interface *interface-name*]
   user@host# **set** keepalive-time ***seconds***

5. Configure the hold time, which must be at least twice the keepalive time.

   [edit protocols oam gre-tunnel interface *interface-name*]
   user@host# **set** hold-time (OAM) ***seconds***

When the keepalive hold time expires, the GRE tunnel will stay up even though the interface cannot send or receive traffic. To verify the GRE tunnel state, check the output for the following commands:

```
user@host> show interfaces gr-3/3/0.3 terse

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

user@host> show interfaces gr-3/3/0.3 extensive
Logical interface gr-3/3/0.3 (Index 73) (SNMP ifIndex 594) (Generation 900)
    Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
10.1.19.11:10.1.19.12:47:df:64:0000000000000000 Encapsulation: GRE-NULL
    Gre keepalives configured: On, Gre keepalives adjacency state: down
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
    Traffic statistics:
     Input  bytes  :             15629992
     Output bytes  :             15912273
```

Copyright © 2013, Juniper Networks, Inc.

```
     Input  packets:                243813
     Output packets:                179476
    Local statistics:
     Input  bytes  :             15322586
     Output bytes  :             15621359
     Input  packets:                238890
     Output packets:                174767
    Transit statistics:
     Input  bytes  :                307406                    0 bps
     Output bytes  :                290914                    0 bps
     Input  packets:                  4923                    0 pps
     Output packets:                  4709                    0 pps
    Protocol inet, MTU: 1476, Generation: 1564, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
        Destination: 200.1.3/24, Local: 200.1.3.1, Broadcast: 200.1.3.255,
Generation: 1366
    Protocol mpls, MTU: 1464, Maximum labels: 3, Generation: 1565, Route table:
0
```

NOTE: When the keepalive hold time has expired, the Link status will be Up and the Gre keepalives adjacency state will be Down.

Related Documentation

- GRE Keepalive Time Overview on page 1271
- Example: Configuring Keepalive for a GRE Interface on page 1293
- keepalive-time on page 1302
- hold-time (OAM) on page 1301

Restricting Tunnels to Multicast Traffic

For interfaces that carry IPv4 or IP version 6 (IPv6) traffic, you can configure a tunnel interface to allow multicast traffic only. To configure a multicast-only tunnel, include the **multicast-only** statement:

multicast-only;

You can configure this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family** *family***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family***]**

Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a 224/8 or greater prefix, the packet is dropped and a counter is incremented.

You can configure this property on GRE, IP-IP, PIM, and multicast tunnel (**mt**) interfaces only.

Copyright © 2013, Juniper Networks, Inc.

NOTE: If your router has a Tunnel Services PIC, the Junos OS automatically configures one multicast tunnel interface (**mt**) for each virtual private network (VPN) you configure. You do not need to configure multicast tunnel interfaces.

Configuring Logical Tunnel Interfaces

Logical tunnel (**lt-**) interfaces provide quite different services depending on the host router:

- On M Series, MX Series, and T Series routers, logical tunnel interfaces allow you to connect logical systems, virtual routers, or VPN instances. M Series and T Series routers must be equipped with a Tunnel Services PIC or an Adaptive Services Module (only available on M7i routers). MX Series routers must be equipped with a Trio MPC/MIC module. For more information about connecting these applications, see the Junos OS VPNs Configuration Guide.
- On SRX Series Services Gateways, the logical tunnel interface is used to interconnect logical systems. See the *Junos OS Logical Systems Configuration Guide for Security Devices*.
- On J Series Services Routers, the logical tunnel interface is used to provide class-of-service (CoS) support for real-time performance monitoring (RPM) probe packets. Packets are routed to this internal interface for services. See the *Junos OS Class of Service Configuration Guide for Security Devices*.

For M Series, MX Series, and T Series routers, see the following section:

- Connecting Logical Systems on page 1281

Connecting Logical Systems

To connect two logical systems, you configure a logical tunnel interface on both logical systems. Then you configure a peer relationship between the logical tunnel interfaces, thus creating a point-to-point connection.

To configure a point-to-point connection between two logical systems, configure the logical tunnel interface by including the **lt-*fpc/pic/port*** statement:

```
lt-fpc/pic/port {
  unit logical-unit-number {
    encapsulation encapsulation;
    peer-unit unit-number; # peering logical system unit number
    dlci dlci-number;
    family (inet | inet6 | iso | mpls);
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

Copyright © 2013, Juniper Networks, Inc.

When configuring logical tunnel interfaces, note the following:

- You can configure each logical tunnel interface with one of the following encapsulation types: Ethernet, Ethernet circuit cross-connect (CCC), Ethernet VPLS, Frame Relay, Frame Relay CCC, VLAN, VLAN CCC, or VLAN VPLS.
- You can configure the IP, IPv6, International Organization for Standardization (ISO), or MPLS protocol family.
- The peering logical interfaces must belong to the same logical tunnel interface derived from the Tunnel Services PIC or Adaptive Services Module.
- You can configure only one peer unit for each logical interface. For example, unit 0 cannot peer with both unit 1 and unit 2.
- To enable the logical tunnel interface, you must configure at least one physical interface statement.
- Logical tunnels are not supported with Adaptive Services, Multiservices, or Link Services PICs (but they are supported on the Adaptive Services Module on M7i routers, as noted above).
- On M Series routers other than the M40e router, logical tunnel interfaces require an Enhanced Flexible PIC Concentrator (FPC).
- On MX Series routers, logical tunnel interfaces require Trio MPC/MIC modules. They do not require a Tunnel Services PIC in the same system.

For more information about configuring logical systems, see the Junos OS Routing Protocols Configuration Guide.

Configuring Tunnel Interfaces for Routing Table Lookup

To configure tunnel interfaces to facilitate routing table lookups for VPNs, you specify a tunnel's endpoint IP addresses and associate them with a routing instance that belongs to a particular routing table. This enables the Junos OS to search in the appropriate routing table for the route prefix, because the same prefix can appear in multiple routing tables. To configure the destination VPN, include the **routing-instance** statement:

```
routing-instance {
  destination routing-instance-name;
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel]**
- **[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel]**

This configuration indicates that the tunnel's destination address is in routing instance ***routing-instance-name***. By default, the tunnel route prefixes are assumed to be in the default Internet routing table **inet.0**.

Copyright © 2013, Juniper Networks, Inc.

NOTE: If you configure a virtual loopback tunnel interface and the **vrf-table-label** statement on the same routing instance, the **vrf-table-label** statement takes precedence over the virtual loopback tunnel interface. For more information, see "Configuring Virtual Loopback Tunnels for VRF Table Lookup" on page 1283.

For more information about VPNs, see the Junos OS VPNs Configuration Guide.

Configuring Virtual Loopback Tunnels for VRF Table Lookup

To enable egress filtering, you can either configure filtering based on the IP header, or you can configure a virtual loopback tunnel on routers equipped with a Tunnel PIC. Table 26 on page 1283 describes each method.

Table 26: Methods for Configuring Egress Filtering

Method	Interface Type	Configuration Guidelines	Comments
Filter traffic based on the IP header	Nonchannelized Point-to-Point Protocol / High Level Data Link Control (PPP/HDLC) core-facing SONET/SDH interfaces	Include the **vrf-table-label** statement at the **[edit routing-instances *instance-name*]** hierarchy level. For more information, see the Junos OS VPNs Configuration Guide.	There is no restriction on customer-edge (CE) router-to-provider edge (PE) router interfaces.
Configure a virtual loopback tunnel on routers equipped with a Tunnel PIC	All interfaces	See the guidelines in this section.	Router must be equipped with a Tunnel PIC. There is no restriction on the type of core-facing interface used or CE router-to-PE router interface used. You cannot configure a virtual loopback tunnel and the **vrf-table-label** statement at the same time.

You can configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels. You might want to enable this functionality so you can do either of the following:

- Forward traffic on a PE router to CE device interface, in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch).

  The first lookup is done based on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared medium.

Copyright © 2013, Juniper Networks, Inc.

- Perform egress filtering at the egress PE router.

  The first lookup on the VPN label is done to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to filter and forward packets. You can enable this functionality by configuring output filters on the VRF interfaces.

To configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels, you specify a virtual loopback tunnel interface name and associate it with a routing instance that belongs to a particular routing table. The packet loops back through the virtual loopback tunnel for route lookup. To specify a virtual loopback tunnel interface name, you configure the virtual loopback tunnel interface at the **[edit interfaces]** hierarchy level and include the **family inet** and **family mpls** statements:

```
vt-fpc/pic/port {
    unit 0 {
        family inet;
        family mpls;
    }
    unit 1 {
        family inet;
    }
}
```

To associate the virtual loopback tunnel with a routing instance, include the virtual loopback tunnel interface name at the **[edit routing-instances]** hierarchy level:

```
interface vt-fpc/pic/port;
```



NOTE: On virtual loopback tunnel interfaces, none of the logical interface statements except the family statement is supported. Note that you can configure only inet and mpls families, and you cannot configure IPv4 or IPv6 addresses on virtual loopback tunnel interfaces. Also, virtual loopback tunnel interfaces do not support class-of-service (CoS) configurations.

Configuring PIM Tunnels

PIM tunnels are enabled automatically on routers that have a tunnel PIC and on which you enable PIM sparse mode. You do not need to configure the tunnel interface.

PIM tunnels are unidirectional.

In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point (RP) router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the RP. The RP then de-encapsulates the packets and transmits them through its multicast tree. To perform the encapsulation and de-encapsulation, the first-hop and RP routers must be equipped with Tunnel PICs.

The Junos OS creates two interfaces to handle PIM tunnels:

Copyright © 2013, Juniper Networks, Inc.

- **pe**—Encapsulates packets destined for the RP. This interface is present on the first-hop router.
- **pd**—De-encapsulates packets at the RP. This interface is present on the RP.



NOTE: **The pe and pd interfaces do not support class-of-service (CoS) configurations.**

Configuring IPv6-over-IPv4 Tunnels

If you have a Tunnel PIC installed in your M Series or T Series router, you can configure IPv6-over-IPv4 tunnels. To define a tunnel, you configure a unicast tunnel across an existing IPv4 network infrastructure. IPv6/IPv4 packets are encapsulated in IPv4 headers and sent across the IPv4 infrastructure through the configured tunnel. You manually configure configured tunnels on each end point.

On SRX Series and J Series devices, Generic Routing Encapsulation (GRE) and IP-IP tunnels use internal interfaces, gr-0/0/0 and ip-0/0/0, respectively. The Junos OS creates these interfaces at system bootup; they are not associated with a physical interface.

IPv6-over-IPv4 tunnels are defined in RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers*. For information about configuring a unicast tunnel, see "Configuring Unicast Tunnels" on page 1273. For an IPv6-over-IPv4 tunnel configuration example, see "Example: Configuring an IPv6-over-IPv4 Tunnel" on page 1290.

Configuring IPv4-over-IPv6 Tunnels

You can configure IPv6 IP-IP tunnels to carry IPv4 traffic across a network as specified in RFC 2473, *Generic Packet Tunneling in IPv6 Specification*. It is becoming common for networks to support IPv6 only. This feature provides a way to transition IPv4-based legacy networks to IPv6.

An IPv6 IP-IP tunnel is a virtual link between two IPv6 routers and is used to transmit data using IPv6 addressed packets. This feature allows you to encapsulate the IPv4 traffic into an IPv6 IP-IP tunnel across an IPv6 network. The tunnel is unidirectional and unicast addressed. For bi-directional connectivity, you need to configure a pair of unidirectional tunnels.

The following procedure outlines how to configure an IP-IP tunnel to carry IPv4 traffic between two IPv6 systems:

1. Configure the **family inet6** statement and the **family inet** statement at the **[edit interfaces *ip-interface-name* unit *number*]** hierarchy level.
2. Configure an IPv6 address for the **source** statement and the **destination** statement at the **[edit interfaces *ip-interface-name* tunnel]** hierarchy level.

Related Documentation

- Example: Configuring an IPv4-over-IPv6 Tunnel on page 1291

Copyright © 2013, Juniper Networks, Inc.

Configuring Dynamic Tunnels

A VPN that travels through a non-MPLS network requires a GRE tunnel. This tunnel can be either a static tunnel or a dynamic tunnel. A static tunnel is configured manually between two PE routers. A dynamic tunnel is configured using BGP route resolution.

When a router receives a VPN route that resolves over a BGP next hop that does not have an MPLS path, a GRE tunnel can be created dynamically, allowing the VPN traffic to be forwarded to that route. Only GRE IPv4 tunnels are supported.

To configure a dynamic tunnel between two PE routers, include the **dynamic-tunnels** statement:

```
dynamic-tunnels tunnel-name {
   destination-networks prefix;
   source-address address;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit routing-instances *routing-instance-name* routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]**

For more information about configuring routing options or BGP, see the Junos OS Routing Protocols Configuration Guide. For more information about VPNs, see the Junos OS VPNs Configuration Guide.

Configuring Tunnel Interfaces on MX Series Routers

Because the MX Series routers do not support Tunnel Services PICs, you create tunnel interfaces on MX Series routers by including the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
   pic number {
      tunnel-services {
         bandwidth (1g | 10g | 20g | 40g);
      }
   }
}
```

fpc *slot-number* is the slot number of the DPC, MPC, or MIC. On the MX80 router, the range is 0 through 1. On other MX series routers, if two SCBs are installed, the range is 0 through 11. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

The **pic *number*** On MX80 routers, if the FPC is 0, the PIC number can only be 0. If the FPC is 1, the PIC range is 0 through 3. For all other MX series routers, the range is 0 through 3.

Copyright © 2013, Juniper Networks, Inc.

bandwidth (**1g** | **10g** | **20g** | **40g**) is the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.



NOTE: When you use MPCs and MICs, tunnel interfaces are soft interfaces and allow as much traffic as the forwarding-path allows, so it is advantageous to setup tunnel services without artificially limiting traffic by use of the **bandwidth** option. However, you *must* specify **bandwidth** when configuring tunnel services for MX Series routers with DPCs or FPCs. The GRE key option is not supported on the tunnel interfaces for DPCs on MX960 routers.

Bandwidth rates of 20 gigabits per second and 40 gigabits per second require use of an MX Series router with the 100-Gigabit Ethernet Modular Port Concentrator (MPC) and the 100-Gigabit CFP MIC.

1g indicates that 1 gigabit per second of bandwidth is reserved for tunnel traffic.

10g indicates that 10 gigabits per second of bandwidth is reserved for tunnel traffic.

20g indicates that 20 gigabits per second of bandwidth is reserved for tunnel traffic.

40g indicates that 40 gigabits per second of bandwidth is reserved for tunnel traffic.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

To verify that the tunnel interfaces have been created, issue the **show interfaces terse** operational mode command. For more information, see the *Junos Interfaces Command Reference*.

Related Documentation

- Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC
- Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC
- Example: Configuring Tunnel Interfaces on the MPC3E
- bandwidth (Tunnel Services)
- tunnel-services (Chassis)
- [edit chassis] Hierarchy Level

Configuring Tunnel Interfaces on T4000 Routers

To create tunnel interfaces on a T4000 Core Router, include the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic number {
    tunnel-services {
      bandwidth bandwidth-value;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
}
```

fpc *slot-number* denotes the slot number of the FPC. On the T4000 router, the range is 0 through 7.



NOTE: This applies only to the T4000 Type 5 FPC. If any other type of FPC is configured in this slot, this configuration is ignored and no tunnel physical interface is created.

pic *number* on the T4000 router is 0 or 1.

bandwidth *bandwidth-value* is the amount of bandwidth to reserve for the tunnel traffic on each Packet Forwarding Engine. The bandwidth value accepted includes every multiple of 10g up to 100g.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 100-Gigabit Ethernet PIC with CFP.

To verify that the tunnel interfaces have been created, issue the **show interfaces terse** operational mode command. For more information, see the *Junos Interfaces Command Reference*.

Related Documentation

- bandwidth (Tunnel Services)
- tunnel-services (Chassis)
- [edit chassis] Hierarchy Level

Examples: Configuring Unicast Tunnels

Configure two unnumbered IP-IP tunnels:

```
[edit interfaces]
ip-0/3/0 {
  unit 0 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.253;
    }
    family inet;
  }
  unit 1 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.254;
    }
    family inet;
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

Configure numbered tunnel interfaces by including an address at the **[edit interfaces ip-0/3/0 unit (0 | 1) family inet]** hierarchy level:

```
[edit interfaces]
ip-0/3/0 {
    unit 0 {
        tunnel {
            source 192.168.4.18;
            destination 192.168.4.253;
        }
        family inet {
            address 10.5.5.1/30;
        }
    }
    unit 1 {
        tunnel {
            source 192.168.4.18;
            destination 192.168.4.254;
        }
        family inet {
            address 10.6.6.100/30;
        }
    }
}
```

Configure an MPLS over GRE tunnel by including the **family mpls** statement at the **[edit interfaces gr-1/2/0 unit 0]** hierarchy level:

```
[edit interfaces]
gr-1/2/0 {
    unit 0 {
        tunnel {
            source 192.168.1.1;
            destination 192.168.1.2;
        }
        family inet {
            address 10.1.1.1/30;
        }
        family mpls;
    }
}
```

Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup

Configure a virtual loopback tunnel for VRF table lookup:

```
[edit routing-instances]
routing-instance-1 {
    instance-type vrf;
    interface vt-1/0/0.0;
    interface so-0/2/2.0;
    route-distinguisher 2:3;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    routing-options {
        static {
```

Copyright © 2013, Juniper Networks, Inc.

```
        route 10.0.0.0/8 next-hop so-0/2/2.0;
      }
    }
  }
  routing-instance-2 {
    instance-type vrf;
    interface vt-1/0/0.1;
    interface so-0/3/2.0;
    route-distinguisher 4:5;
    vrf-import VPN-B-import;
    vrf-export VPN-B-export;
    routing-options {
      static {
        route 10.0.0.0/8 next-hop so-0/3/2.0;
      }
    }
  }
  [edit interfaces]
  vt-1/0/0 {
    unit 0 {
      family inet;
      family mpls;
    }
    unit 1 {
      family inet;
    }
  }
```

Example: Configuring an IPv6-over-IPv4 Tunnel

Configure a tunnel on both sides of the connection.

Configuration on Router 1

```
[edit]
interfaces {
  gr-1/0/0 {
    unit 0 {
      tunnel {
        source 10.19.2.1;
        destination 10.19.3.1;
      }
      family inet6 {
        address 2001:DB8:1:1/126;
      }
    }
  }
}
```

Configuration on Router 2

```
[edit]
interfaces {
  gr-1/0/0 {
    unit 0 {
      tunnel {
        source 10.19.3.1;
        destination 10.19.2.1;
      }
```


Copyright © 2013, Juniper Networks, Inc.

```
        family inet6 {
          address 2001:DB8:2:1/126;
        }
      }
    }
  }
```

Example: Configuring an IPv4-over-IPv6 Tunnel

You can configure IPv6 IP-IP tunnels to carry IPv4 traffic across a network as specified in RFC 2473, *Generic Packet Tunneling in IPv6 Specification*. It is becoming common for networks to support IPv6 only. This feature provides a way to transition IPv4-based legacy networks to IPv6.

Figure 17: IPv6 Tunnel Connecting Two IPv4 Networks Across an IPv6 Network

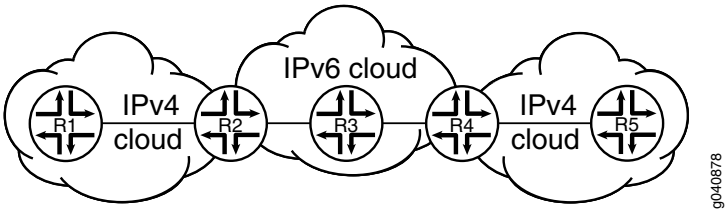

The following example is based on the topology show in Figure 17 on page 1291. Routers R2, R3, and R4 represent the IPv6 network. Routers R1 and R2 and R4 and R5 represent the IPv4 networks that need to be connected by an IPv6 tunnel. Routers R2 and R4 represent the IPv6 tunnel endpoints.

The following example illustrates the configuration for router R2 as shown in Figure 17 on page 1291. On router R2, you configure an IPv4 over IPv6 uni-directional IP-IP tunnel which includes the following elements:

- The tunnel source IPv6 address is 2001:DB8:2::1. It could match Router R2's loopback address.
- The tunnel destination IPv6 address is 2001:DB8:3::1 It could match Router R4's loopback address.
- On Router R4, the tunnel receiving traffic from the Router R1 to R2 IPv4 network needs to have an IPv4 address in the same subnet as 1.1.1.1/30 (for example, 1.1.1.2).

```
[edit]
  interfaces {
    ip-1/2/0 {
    unit 0 {
      tunnel {
        source 2001:DB8:2::1;
        destination 2001:DB8:3::1;
      }
      family inet {
        address 1.1.1.1/30;
      }
    }
```

```
    }
  }
```

The output from the **show interfaces ip-1/2/0** command displays the following:

```
user@host> show interfaces ip-1/2/0
Physical interface: ip-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 521
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-1/2/0.0 (Index 74) (SNMP ifIndex 540)
    Flags: Point-To-Point SNMP-Traps 0x4000
    IP-Header 2001:db8:2::1-2001:db8:3::1-41-64-00000000
    Encapsulation: IPIP-NULL
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: Unlimited
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.1.0/30, Local: 1.1.1.1

astatti@tp9>
```

When attempting to configure an IPv4 over IPv6 tunnel, be aware of the following:

- The IP-IP interface comes up only when the tunnel source and tunnel destination reachability information is populated in the routing table.
- To carry IPv6 traffic over an IPv6 IP-IP tunnel, the IP interface needs to be configured with an IPv6 address (using **set interfaces *ip-interface-name* unit 0 family inet6 address *address***).

Related Documentation

- Configuring IPv4-over-IPv6 Tunnels
- Example: Configuring an IPv6-over-IPv4 Tunnel on page 1290

Example: Configuring Logical Tunnels

Configure three logical tunnels:

```
[edit interfaces]
lt-4/2/0 {
  description “Logical tunnel interface connects three logical systems”;
}
[edit logical-systems]
lr1 {
  interfaces lt-4/2/0 {
    unit 12 {
      peer-unit 21; #Peering with lr2
      encapsulation frame-relay;
      dlci 612;
      family inet;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
      unit 13 {
        peer-unit 31; #Peering with lr3
        encapsulation frame-relay-ccc;
        dlci 613;
      }
    }
  }
  lr2 {
    interfaces lt-4/2/0 {
      unit 21 {
        peer-unit 12; #Peering with lr1
        encapsulation frame-relay-ccc;
        dlci 612;
      }
      unit 23 {
        peer-unit 32; #Peering with lr3
        encapsulation frame-relay;
        dlci 623;
      }
    }
  }
  lr3 {
    interfaces lt-4/2/0 {
      unit 31 {
        peer-unit 13; #Peering with lr1
        encapsulation frame-relay;
        dlci 613;
        family inet;
      }
      unit 32 {
        peer-unit 23; #Peering with lr2
        encapsulation frame-relay-ccc;
        dlci 623;
      }
    }
  }
```

Example: Configuring Keepalive for a GRE Interface

The following example illustrates the minimum configuration of a GRE tunnel interface:

```
[edit]
user@host# show interfaces gr-3/2/0 unit 0
  {
    family inet;
  }
[edit]
user@host# show
protocols {
  oam {
    gre-tunnel {
      interface gr-1/1/10.1 {
        keepalive-time 10;
          hold-time 30;
        }
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
    }
}
```

Related Documentation

- GRE Keepalive Time Overview on page 1271
- Configuring GRE Keepalive Time on page 1279
- **keepalive-time on page 1302**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 62

Summary of Tunnel Services Configuration Statements

The following sections explain each of the tunnel services statements. The statements are organized alphabetically.

allow-fragmentation

Syntax allow-fragmentation;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable fragmentation of generic routing encapsulation (GRE) encapsulated packets regardless of maximum transmission unit (MTU) value.

Default By default, the GRE-encapsulated packets are dropped if the packet size exceeds the MTU setting of the egress interface.

Usage Guidelines See "Configuring Packet Reassembly" on page 1278.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- reassemble-packets on page 1304

Copyright © 2013, Juniper Networks, Inc.

backup-destination

Syntax backup-destination *destination-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel],[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.

Description For tunnel interfaces, specify the remote address of the backup tunnel.

Options ***destination-address***—Address of the remote side of the connection.

Usage Guidelines See "Configuring IPsec Tunnel Redundancy" on page 889.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- destination (Interfaces) on page 893
- destination (Tunnel Remote End) on page 984

copy-tos-to-outer-ip-header

Syntax copy-tos-to-outer-ip-header;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For GRE tunnel interfaces only, enable the inner IP header's ToS bits to be copied to the outer IP packet header.

Default If you omit this statement, the ToS bits in the outer IP header are set to 0.

Usage Guidelines See "Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header" on page 1277.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

destination

See the following sections:

- destination (Tunnel Remote End) on page 1297
- destination (Routing Instance) on page 1297

destination (Tunnel Remote End)

Syntax destination *address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description For tunnel interfaces, specify the remote address of the tunnel.

Options ***destination-address***—Address of the remote side of the connection.

Usage Guidelines See "Configuring Unicast Tunnels" on page 1273, "Configuring Traffic Sampling" on page 912, and "Configuring Flow Monitoring" on page 920.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

destination (Routing Instance)

Syntax destination *routing-instance-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel routing-instance]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default The default Internet routing table **inet.0**.

Usage Guidelines See "Configuring Tunnel Interfaces for Routing Table Lookup" on page 1282.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

destination-networks

Syntax destination-networks *prefix*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name* rsvp-te *entry*],
[edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name* rsvp-te *entry*],
[edit routing-options dynamic-tunnels *tunnel-name*],
[edit routing-options dynamic-tunnels *tunnel-name* rsvp-te *entry*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.

Options ***prefix***—Destination prefix of the network.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring GRE Tunnels for Layer 3 VPNs
- Configuring Dynamic Tunnels on page 1286
- Configuring RSVP Automatic Mesh

Copyright © 2013, Juniper Networks, Inc.

do-not-fragment

Syntax do-not-fragment;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel]

Release Information Statement introduced in Junos OS Release 9.2.

Description Set the do-not-fragment (DF) bit on the packets entering the GRE tunnel so that they do not get fragmented anywhere in the path.

Default By default, fragmentation is disabled.

Usage Guidelines See "Configuring Packet Reassembly" on page 1278.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- reassemble-packets on page 1304

Copyright © 2013, Juniper Networks, Inc.

dynamic-tunnels

Syntax

```
dynamic-tunnels tunnel-name {
  destination-networks prefix;
  gre;
  rsvp-te entry-name {
    destination-networks network-prefix;
    label-switched-path-template {
      default-template;
      template-name;
    }
  }
  source-address address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a dynamic tunnel between two PE routers.

Options

tunnel-name—Name of the dynamic tunnel.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks
- Configuring GRE Tunnels for Layer 3 VPNs
- See Configuring Dynamic Tunnels on page 1286.

Copyright © 2013, Juniper Networks, Inc.

hold-time (OAM)

Syntax hold-time *seconds*;

Hierarchy Level [edit protocols oam],
[edit protocols oam gre-tunnel interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Length of time the originating end of a GRE tunnel waits for keepalive packets from the other end of the tunnel before marking the tunnel as operationally down.

Options *seconds*—Hold-time value.

Default: 5 seconds

Range: 5 through 250 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- GRE Keepalive Time Overview on page 1271
- Configuring GRE Keepalive Time on page 1279
- Example: Configuring Keepalive for a GRE Interface on page 1293
- **keepalive-time on page 1302**

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines See the Junos® OS Network Interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

keepalive-time

Syntax keepalive-time *seconds*;

Hierarchy Level [edit protocols oam],
[edit protocols oam gre-tunnel interface *interface-name*],
[edit protocols oam gre-tunnel interface *interface-name.unit-number*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Time difference between consecutive keepalive packets in a GRE tunnel.



NOTE: Support for GRE keepalive packets on MPC line cards became available as of Junos OS Release 11.4.

Options *seconds*—Keepalive time value.

Default: 1 second

Range: 1 through 50 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- GRE Keepalive Time Overview on page 1271
- Configuring GRE Keepalive Time on page 1279
- Example: Configuring Keepalive for a GRE Interface on page 1293
- **hold-time (OAM) on page 1301**

Copyright © 2013, Juniper Networks, Inc.

key

Syntax key *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.

Description Identify an individual traffic flow within a tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*. On M Series and T Series routers, you can configure the GRE interface on an Adaptive Services, Multiservices, or Tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC.

Options *number*—Value of the key.
Range: 0 through 4,294,967,295

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Key Number on GRE Tunnels on page 1276

multicast-only

Syntax multicast-only;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the unit and family so that the interface can transmit and receive multicast traffic only. You can configure this property on the IP family only.

Usage Guidelines See "Restricting Tunnels to Multicast Traffic" on page 1280.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- tunnel on page 1308

Copyright © 2013, Juniper Networks, Inc.

peer-unit

Syntax peer-unit *unit-number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a peer relationship between two logical systems.

Options ***unit-number***—Peering logical system unit number.

Usage Guidelines See "Configuring Logical Tunnel Interfaces" on page 1281.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

reassemble-packets

Syntax reassemble-packets;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable reassembly of fragmented tunnel packets on generic routing encapsulation (GRE) tunnel interfaces.

Usage Guidelines See "Configuring Packet Reassembly" on page 1278.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

routing-instance

Syntax
```
routing-instance {
  destination routing-instance-name;
}
```

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default The default Internet routing table **inet.0**.

Usage Guidelines See "Configuring Tunnel Interfaces for Routing Table Lookup" on page 1282.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

routing-instances

Syntax routing-instances *routing-instance-name* { ... }

Hierarchy Level
[edit],
[edit logical-systems *logical-system-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPF version 3 (OSPFv3), and RIP for a router.

Default Routing instances are disabled for the router.

Options ***routing-instance-name***—Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately.

Usage Guidelines See the Junos OS Routing Protocols Configuration Guide and the Routing Policy Configuration Guide.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

routing-options

Syntax routing-options { ... }

Hierarchy Level [edit],
[edit logical-systems *logical-system-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure protocol-independent routing properties.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Protocol-Independent Routing Properties Configuration Guide

source

Syntax source *source-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify the source address of the tunnel.

Default If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.

Options ***source-address***—Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)

Copyright © 2013, Juniper Networks, Inc.

source-address

Syntax source-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-instances *routing-instance-name* routing-options dynamic-tunnels *tunnel-name*],
[edit routing-options dynamic-tunnels *tunnel-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configure the tunnel source address.

Options ***address***—Name of the source address.

Usage Guidelines See "Configuring Dynamic Tunnels" on page 1286.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

ttl

Syntax ttl *value*;

Hierarchy Level [edit interfaces *interface-name* unit *number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Set the time-to-live value bit in the header of the outer IP packet.

Options ***value***—Time-to-live value.
Range: 0 through 255
Default: 64

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tunnel Properties
- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)

Copyright © 2013, Juniper Networks, Inc.

tunnel

Syntax

```
tunnel {
  allow-fragmentation;
  backup-destination address;
  destination destination-address;
  do-not-fragment;
  key number;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Encryption Interfaces on page 881
- Junos OS VPNs Configuration Guide

Copyright © 2013, Juniper Networks, Inc.

unit (Interfaces)

Syntax

```
unit logical-unit-number {
  peer-unit unit-number;
  reassemble-packets;
  tunnel {
    allow-fragmentation;
    backup-destination address;
    destination destination-address;
    do-not-fragment;
    key number;
    routing-instance {
      destination routing-instance-name;
    }
    source source-address;
    ttl number;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Junos® OS Network Interfaces for other statements that do not affect services interfaces.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 9

Index

- Index on page 1313
- Index of Statements and Commands on page 1337

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Index

Symbols

#, comments in configuration statements....l
( ), in syntax descriptions....l
< >, in syntax descriptions....xlix
[ ], in configuration statements....l
{ }, in configuration statements....l
| (pipe), in syntax descriptions....l

A

AACL
  action statements....842
  applications....841
  best-effort application identification....777
  example configuration....845
  logging flows....843
  match conditions....841
  rules....844
aacl-fields statement....868
aacl-statistics-profile statement....869
accelerations statement....667
accept
  action....912
accounting statement
  flow monitoring....978
  usage guidelines....967
acknowledge-retries statement....1189
  usage guidelines....1148
acknowledge-timer statement....1190
  usage guidelines....1148
action-red-differential-delay statement....1191
  usage guidelines....1149
actions statement....666
adaptive-services-pics statement....485
  usage guidelines....1058
address statement
  APPID
    usage guidelines....785
  application rule....800
  DFC....1103
    usage guidelines....1086
  encryption....891
    usage guidelines....881
  flow monitoring....979
    usage guidelines....912
  interfaces....527
    usage guidelines....516
  link services....1192
    usage guidelines....1135
  voice services....431
    usage guidelines....422
administrative statement
  BGF....554
admission-control statement....668, 669
aggregate-export-interval statement....979
  usage guidelines....967
aggregation statement....197
  flow monitoring....980
  usage guidelines....190, 929
alert (system logging severity level)....317, 479, 518
algorithm statement....556
ALGs
  application protocols....73
  configuring....74
  definition....73
allow-fragmentation statement....1295
  usage guidelines....1278
allow-ip-options statement....142
  usage guidelines....134
allow-multicast statement....486
  usage guidelines....480
allowed-destinations statement....1104
  usage guidelines....1087
AMS
  HA....167, 168
  NAT....167, 169
analyzer-address statement....1065
  usage guidelines....1055
analyzer-id statement....1066
  usage guidelines....1055
anomaly checklist....46
anti-replay-window-size statement....273, 487
  usage guidelines....248, 476
any (system logging severity level)....317, 479, 518
any-any match condition
  Ipsec....244
APPID
  best-effort application identification....777
  example configuration....796
application layer gateways *See* ALGs

application protocol
definition....73
application statement....121, 801, 802
APPID
usage guidelines....783
usage guidelines....74
application-aware-access-list-fields statement....870
application-data-inactivity-detection statement....556
application-group statement....802
APPID
usage guidelines....788
application-group-any statement....848
AACL
usage guidelines....841
PTSP....747
application-groups statement....803, 848
AACL
usage guidelines....841
APPID
usage guidelines....788
PTSP....747
application-profile statement....452
usage guidelines....446
application-protocol statement....122
usage guidelines....74
application-set statement....123
usage guidelines....85
application-sets statement
CoS....453
usage guidelines....444
IDS....198
usage guidelines....189
stateful firewall....143
usage guidelines....133
application-system-cache-timeout statement....803
APPID
usage guidelines....791
applications....189
example configuration....118
applications statement
AACL....847
usage guidelines....841
APPID
usage guidelines....788
application identification....804
application-level gateways....143
applications hierarchy....123
usage guidelines....73
CoS....453
usage guidelines....444
IDS....198
usage guidelines....189
PTSP....748
stateful firewall....143
usage guidelines....133
applying service set to interface....468
archive-sites statement....1066
usage guidelines....1057
AS PIC
multicast traffic....480
redundancy....320, 522, 975
asymmetrical routing support
APPID....793
attack detection....185
audit-observed-events-returns statement....557
authentication statement....274
usage guidelines....225
authentication-algorithm statement
IKE....275
usage guidelines....229
IPsec....275
usage guidelines....238
authentication-method statement....276
usage guidelines....229
authentication-mode statement
RPM....1237
automatic statement....804
APPID
usage guidelines....792
autonomous-system-type statement....981
usage guidelines....929
auxiliary-spi statement....276
usage guidelines....225
availability-check-profiles statement....670

B

backup AS PIC....522
backup Link Services IQ PIC....352
backup-destination statement....1296
usage guidelines....889
backup-interface statement....892
usage guidelines....888
backup-remote-gateway statement....277
usage guidelines....247

Copyright © 2013, Juniper Networks, Inc.

bandwidth
and delay buffer allocation....369
guaranteed....369, 374
base-root statement....558
best-effort application identification....777
bgf-core statement....559
bgp statement
RPM....1238
blacklist-period statement....671
braces, in configuration statements....l
brackets
angle, in syntax descriptions....xlix
square, in configuration statements....l
bundle statement....432, 1192
usage guidelines....426, 1135
by-destination statement....199
usage guidelines....190
by-pair statement....200
usage guidelines....190
by-source statement....201
usage guidelines....190
bypass-traffic-on-exceeding-flow-limits statement....488
bypass-traffic-on-pic-failure statement....488
usage guidelines....468

C

cancel-graceful statement....561
capture-group statement....1105
usage guidelines....1086
cflowd statement....982
usage guidelines....929
cgn-pic statement....528
chain-order statement
nested applications....805
CIR....374
cisco-interoperability statement....409
usage guidelines....351
cleanup-timeout statement....562
clear-don't-fragment-bit (NAT option)....498
clear-dont-fragment-bit statement
GRE tunnel....528
IPsec....277
usage guidelines....246
service-set....489
usage guidelines....247, 476, 519, 1276
clear-ike-sas-on-pic-restart statement....278
usage guidelines....228
clear-ipsec-sas-on-pic-restart statement....278
usage guidelines....228
client-list statement....1239
clusters statement....672
collector statement....1067
usage guidelines....1056
collector-pic statement
usage guidelines....1058
comments, in configuration statements....l
committed-burst-size statement....673
committed-information-rate statement....674
compression statement....432
usage guidelines....423, 424
compression-device statement....433
usage guidelines....426
configuration
dynamic flow capture interface....1092
flow collector interface....1058
flow-tap application....1101
content destinations
DFC....1083
flow-tap....1096
content-destination statement....1106
usage guidelines....1086
context statement
nested applications....805
context-indications statement....563
control source
DFC....1083
control-association-indications statement....564
control-cores statement....157
control-source statement....1107
usage guidelines....1087
controller-address statement....565
controller-failure statement....565
controller-port statement....566
conventions
text and syntax....xlix
copy-tos-to-outer-ip-header statement....1296
usage guidelines....1277
core-dump statement....984
usage guidelines....920
CoS
action statements....445
applications....444
example configuration....447
for tunnels
GRE TOS bits....1277
link services interfaces....365, 367, 1151

Copyright © 2013, Juniper Networks, Inc.

link services IQ interfaces....347
match conditions....444
rules....448
scheduler map
configuration example....1145
count-type statement
PTSP....748
critical (system logging severity level)....317, 479, 518
curly braces, in configuration statements....l
customer support....l
contacting JTAC....l

D

Data inactivity detection....630
data session identification
APPID....786
data statement....454
usage guidelines....446
data-cores statement....158
data-fill statement....1239
data-flow-affinity statement....158
data-format statement....1067
usage guidelines....1056
data-inactivity-detection statement....566, 674
data-size statement....1240
usage guidelines....1221
datastore statement....675
dead peer detection (DPD) protocol....247
default statement....567
default-media-realm statement....676
delay buffer
calculating....369, 374
shaping rate....369, 374
delay-buffer-rate statement
usage guidelines....369
delivery-function statement....568
demux statement
PTSP....749
description statement
IKE....279
usage guidelines....236
IPsec....279
usage guidelines....238, 241
destination statement....159, 893, 1193
APPID
usage guidelines....785
application identification rule....806
encryption
usage guidelines....881, 889
flow monitoring
usage guidelines....912
link services
usage guidelines....1135
tunnel....1297
usage guidelines....1273, 1282
destination-address statement
AACL....849
usage guidelines....841
BGF....568
CoS....454
usage guidelines....444
IDS....202
usage guidelines....189
IPsec....279
usage guidelines....244
stateful firewall....144
usage guidelines....133
destination-address-range statement
AACL....849
usage guidelines....841
IDS....202
usage guidelines....189
stateful firewall....144
usage guidelines....133
destination-interface statement
RPM....1241
usage guidelines....1221, 1225
destination-networks statement
tunnel....1298
usage guidelines....1286
destination-port statement
applications....123
BGF....569
RPM....124, 1242
usage guidelines....79
destination-prefix statement....203
usage guidelines....190
destination-prefix-ipv6 statement....203
usage guidelines....190
destination-prefix-list statement
AACL....850
usage guidelines....841
CoS....455
IDS....204
stateful firewall....145
usage guidelines....133

Copyright © 2013, Juniper Networks, Inc.

destinations statement
flow collection....1068
usage guidelines....1055
detect statement....569
DFC
architecture....1083
capture group....1086
control source configuration....1087
destination configuration....1086
example configuration....1092
interface configuration....1088
system logging....1090
threshold configuration....1090
dh-group statement....280
usage guidelines....230
dial-options statement....529
interfaces
usage guidelines....318
dialogs statement....677
diffserv statement....570
direction statement....281
nested applications....806
usage guidelines....224
disable statement
APPID
usage guidelines....783, 788
application....807
application group....807
flow monitoring....985
port mapping....808
traffic sampling
usage guidelines....915
disable-all-instances statement
flow monitoring....985
disable-global-timeout-override statement....808
usage guidelines....783
disable-mlppp-inner-ppp-pfc statement....1194
usage guidelines....1139
disable-session-mirroring statement....570
discard accounting
usage guidelines....967
disconnect statement....571
dlci statement....1194
usage guidelines....1143
DLCIs
multicast-capable connections....1143
point-to-point connections....1143
do-not-fragment statement
tunnel....1299
usage guidelines....1278
documentation
comments on....l
down statement....572
download statement
APPID....809
usage guidelines....792
drop-member-traffic statement
aggregated Multiservices....173
drop-timeout statement....1195
usage guidelines....1139
dscp statement....455
BGF....573
BSG....678
usage guidelines....445
dscp-code-point statement
RPM....1243
usage guidelines....1221
DTCP....1083, 1095
duplicates-dropped-periodicity statement....1107
usage guidelines....1091
dynamic authentication....250
dynamic flow capture *See* DFC
dynamic route insertion....251
dynamic rules....251
dynamic security associations
usage guidelines....227, 228
dynamic statement....282
usage guidelines....227
Dynamic Tasking Control Protocol *See* DTCP
dynamic tunnels....1300
destination....1298
source....1307
dynamic-flow-capture statement....1108
dynamic-tunnels statement....1300
usage guidelines....1286

E

egress-service-point statement....679
embedded-spdf statement....680
emergency (system logging severity level)....317, 479, 518
enable flow collection mode....1058
enable-asymmetic-traffic-processing statement....810
enable-heuristics statement....809, 810
usage guidelines....792

Copyright © 2013, Juniper Networks, Inc.

enable-rejoin statement
aggregated Multiservices....174
encapsulation statement....433
link services....1196
usage guidelines....1138
voice services
usage guidelines....425
encoding statement....573
encrypted traffic identification
APPID....792
encryption interface....881
applying inbound filter....887
example configuration....887
applying outbound filter....886
example configuration....885, 886
configuring inbound filter....886
example configuration....887
configuring MTU....882
encryption statement....283
usage guidelines....226
encryption-algorithm statement
IKE....284
usage guidelines....230
IPsec....284
usage guidelines....238
engine-id statement
flow monitoring....986
engine-type statement....987
error (system logging severity level)....317, 479, 518
ES interfaces
example configuration....882
ES PIC
apply inbound filter....887
PIC redundancy....888
redundancy
example configuration....889
tunnel redundancy....889
es-options statement....894
usage guidelines....888
event policy
all (tracing flag)....482
APPID....795
configuration (tracing flag)....482
database (tracing flag)....482
events (tracing flag)....482
policy (tracing flag)....482
event-timestamp-notification statement....574
examples
configuring IP fragment reassembly on l2tp....321
export-format statement....989
usage guidelines....923
extension-provider statement....160
extension-service statement....988

F

f-max-period statement....434
usage guidelines....423
facility-override statement....329, 490, 530
usage guidelines....479
failover statement....576
failover-cold statement....574
failover-warm statement....575
family statement
aggregated Multiservices....174
encryption....895
usage guidelines....881
flow monitoring....990
usage guidelines....912
interfaces....531
usage guidelines....516
link services....1198
usage guidelines....1135
voice services....435
fast-update-filters statement....577
file statement....995
BGF....578
border signaling gateway....681
L-PDF statistics....871
traffic sampling....995
traffic sampling output
usage guidelines....915, 917
file-specification statement....1069
usage guidelines....1056
filename statement....996
filename-prefix statement....1068
usage guidelines....1057
files
logging information output file....917
traffic sampling output files....915
var/log/sampled file....917
var/tmp/sampled.pkts file....915
files statement....996
usage guidelines....915

Copyright © 2013, Juniper Networks, Inc.

filter statement
encryption....896
usage guidelines....887
flow monitoring....997
usage guidelines....912
filters
used with services....468
firewall filters
actions....912
in traffic sampling....912
service filters....521
flag statement....579, 682
flow aggregation....928
multiple flow servers....947
flow collector
analyzer configuration....1055
destination configuration....1055
example configuration....1058
file format configuration....1056
interface mapping....1056
transfer log....1057
flow limiting....478
flow monitoring
example configuration
multiple port mirroring....957
next-hop groups....957
load balancing....963
overview....903
redundancy....975
flow server
replicating flows to multiple servers....947
flow-active-timeout statement....998
usage guidelines....923
flow-collector statement....1070
usage guidelines....1053, 1058
flow-control-options statement....999
flow-export-destination statement....1000
usage guidelines....923
flow-export-rate statement
flow monitoring....999
flow-inactive-timeout statement....1001
usage guidelines....923
flow-monitoring statement....1002
flow-server statement
flow monitoring....1003
flow-tap
application....1095
architecture....1096
interface....1097
permissions statement....1098
RADIUS configuration....1098
restrictions....1099
security....1098
flow-tap application
example configuration....1101
flow-tap statement....1109
flow-tap-dtcp statement....1098
font conventions....xlix
force-entry statement....204
usage guidelines....190
forward-manipulation statement....683
forward-rule statement
PTSP....751
forwarding instance....750
forwarding classes
fragmentation....365
forwarding-class statement....410, 456
usage guidelines....365, 445
forwarding-db-size statement....161
setting for stateful firewall....155
forwarding-options statement....1004
usage guidelines....908
fragment reassembly
configuring on l2tp....321
on l2tp....320
fragment-threshold statement
link services....1199
usage guidelines....1140
LSQ....411
usage guidelines....365
voice services....436
usage guidelines....424
fragmentation
forwarding classes....365
GRE tunnels....1276
multiclass MLPPP....367
fragmentation and reassembly....424, 1144
example configuration....1145
fragmentation-map statement....411
usage guidelines....365
fragmentation-maps statement....412
usage guidelines....365
Frame Relay connections
point-to-point connections....1143
Frame Relay encapsulation
multicast-capable connections....1143
framework statement....684

Copyright © 2013, Juniper Networks, Inc.

FRF.12....424
example configuration....398
LFI....1144
LSQ....395
FRF.15 and FRF.16....1127
FRF.16....385
configuration example....388
from statement
AACL....850
usage guidelines....840
border signaling gateway
new call usage policy....687
new transaction policy....688
service class....690
CoS....456
usage guidelines....443
HCM....825
usage guidelines....61
IDS....205
usage guidelines....187, 189
IPsec....285
usage guidelines....243, 244
PTSP....752
PTSP forward rule....751
stateful firewall....146
usage guidelines....132, 133
ftp statement....457
flow collection....1072
usage guidelines....446, 1055, 1057
FTP traffic, sampling....919

G

g-duplicates-dropped-periodicity statement....1110
usage guidelines....1091
g-max-duplicates statement....1111
usage guidelines....1091
gateway statement
BGF....580
border signaling gateway....691
gateway-address statement....584
gateway-controller statement....585
gateway-port statement....586
graceful statement....587
graceful-restart statement....588
GRE tunnels
fragmentation....1276
key number....1276
guaranteed rate....374
guaranteed-rate statement
usage guidelines....374

H

H.248
properties....592, 610, 611, 612, 613, 614, 615, 616, 617, 620, 621
BFG....570
BGF....573
h248-options statement....589
h248-profile statement....591
h248-properties statement....592
h248-stack statement....595
h248-timers statement....596
hanging-termination-detection statement....596
hard-limit statement....1111
usage guidelines....1086
hard-limit-target statement....1112
usage guidelines....1086
hardware requirements....3
hardware-timestamp statement....1244
hash-key statement....162
hcm statement....825
hello-interval statement
L2TP....330
usage guidelines....316
hello-timer statement
link services....1200
usage guidelines....1148
heuristics support
APPID....792
hide-avps statement....331
usage guidelines....316
high-availability-options statement
aggregated Multiservices....175
history-size statement....1244
usage guidelines....1219, 1221
hold-time statement
GRE tunnel interface....1301
host statement....491, 532
HCM....826
L2TP....331
usage guidelines....317, 479, 518
hot-standby statement....412

I

icmp-code statement....124
usage guidelines....78
icmp-type statement....125
usage guidelines....78

Copyright © 2013, Juniper Networks, Inc.

icons defined, notice....xlviii
idle-timeout statement....811
APPID
usage guidelines....783
IDS
action statements....190
applications....189
example configurations....193
match conditions....189
rules....187
ids-rule-sets statement
usage guidelines....472
ids-rules statement....491
usage guidelines....472
ignore-entry statement....204
usage guidelines....190
ignore-errors statement....811
usage guidelines....786
IKE....58, 228
authentication algorithm
usage guidelines....229
authentication-method statement
usage guidelines....229
DH (Diffie-Hellman) group
usage guidelines....230
dynamic SAs....228
encryption-algorithm statement
usage guidelines....230
lifetime
usage guidelines....231
mode statement
usage guidelines....233
policy....232
example....236
policy statement
usage guidelines....232
pre-shared-key statement
usage guidelines....234
proposals statement
usage guidelines....234
version statement
usage guidelines....233
IKE security associations
clearing....228
ike statement....286
usage guidelines....228
ike-access-profile statement....492
usage guidelines....253, 475
inactivity-delay statement....597
inactivity-duration statement....597, 696
inactivity-non-tcp-timeout statement....812
usage guidelines....783
inactivity-tcp-timeout statement....812
usage guidelines....783
inactivity-timeout statement....125
BGF....598
flow monitoring....532
RPM....1245
usage guidelines....82, 516
inactivity-timer statement....599
index statement....813
APPID
usage guidelines....783, 788
nested applications....813
info (system logging severity level)....317, 480, 519
initial-average-ack-delay statement....599
initiate-dead-peer-detection statement....287
usage guidelines....248
inline-jflow statement
flow monitoring....1005
usage guidelines....942, 946
input statement
flow monitoring....1006
interfaces....533
usage guidelines....468, 520
input-interface-index statement....1007
input-packet-rate-threshold statement....1112
usage guidelines....1090
inside and outside interfaces....471
inside-service-interface statement
usage guidelines....472
instance statement
port mirroring....1008
sampling....1009
usage guidelines....941
interchassis LSQ failover....350
interface preservation....356
interface statement
encryption
usage guidelines....881
flow monitoring....1011
usage guidelines....952
flow-tap....1113
usage guidelines....1097
service interface pool....657
interface style service sets....471
interface-map statement....1074
usage guidelines....1056

Copyright © 2013, Juniper Networks, Inc.

interface-service statement....492
usage guidelines....468
interfaces
naming....515
interfaces statement
aggregated Multiservices....176
DFC....1113
usage guidelines....1088
encryption....896
usage guidelines....881
flow monitoring....1013
usage guidelines....912
interfaces hierarchy....533
usage guidelines....513
link services....1200
usage guidelines....1127
tunnel....1301
usage guidelines....1273
voice services....436
interim-ah-scheme statement....600
interleave-fragments statement....1201
usage guidelines....1144
Internet Key Exchange *See* IKE
intrachassis LSQ failover....352
intrusion detection
example configurations....193
rule set....193
tasks....185
IP addresses
sampling traffic from single IP addresses....918
IP fragment reassembly
configuring on l2tp....321
on l2tp....320
ip statement
APPID
usage guidelines....785
application identification....814
ip-flow-stop-detection statement....600
ip-reassembly-rules statement
service-set....493
IPsec
action statements....246
authentication statement
usage guidelines....225
authentication-algorithm statement
usage guidelines....238
direction
usage guidelines....224
dynamic authentication....250
dynamic endpoints interface
configuration....254
dynamic rules....251
dynamic security associations
usage guidelines....227
encryption
usage guidelines....226
encryption-algorithm statement
usage guidelines....238
ES PIC....881
example configuration....258
inbound traffic....887
outbound traffic....885
IKE....58
lifetime of SA....239
lifetime-seconds statement....239
match conditions....244
minimum configurations
dynamic SA....221
manual SA....221
overview....57
perfect-forward-secrecy statement
usage guidelines....241
policy
overview....240
policy statement
usage guidelines....240
proposal statement
usage guidelines....237
proposals statement
usage guidelines....241
protocol statement (dynamic SA)
usage guidelines....240
protocol statement (manual SA)
usage guidelines....225
rule sets....249
security associations....57
security parameter index
usage guidelines....225
service set dynamic endpoints
configuration....253
traffic....883
IPSec
Services SDK
configuration....257
ipsec statement....287
usage guidelines....237
ipsec-inside-interface
usage guidelines....251

Copyright © 2013, Juniper Networks, Inc.

ipsec-inside-interface statement....288
usage guidelines....244
ipsec-interface-id statement
usage guidelines....254
ipsec-sa statement
encryption....897
usage guidelines....881
ipsec-transport-security-association statement....601
ipsec-vpn-options statement....494
usage guidelines....474
ipsec-vpn-rule-sets statement
usage guidelines....472
ipsec-vpn-rules statement....494
usage guidelines....472
IPv4 static source translation
AMS....169
example....169
ipv4-template statement....1013
IPv6
transition
configured tunnel....1285
IPv6-over-IPv4 tunnel
example configuration....1290
standards supported....1285
ipv6-template statement....1013

J

jservices-sfw package....153

K

keepalive-time statement
GRE tunnel interface....1302
key statement
tunnel....1303
usage guidelines....1276

L

L-PDF
best-effort application identification....777
L2TP
access profile....314, 315
attribute-value pairs....316
example configuration....324
redundancy....320
timers....316
l2tp
configuring IP fragment reassembly....321
IP fragment reassembly overview....320
L2TP LNS statements
service-interface....336
l2tp statement
usage guidelines....309
L2TP statements
traceoptions....340
l2tp-access-profile statement....332
usage guidelines....315
l2tp-interface-id statement
usage guidelines....318
l2tp-profile statement
usage guidelines....314
label-position statement....1014
latch-deadlock-delay statement....601
lawful intercept architecture....1096
learn-sip-register statement....126
LFI....390, 395, 424, 1144
example configuration....393, 398, 1145
lifetime-seconds statement
IKE....288
usage guidelines....231
IPsec....288
usage guidelines....239
limiting flows per service set....478
link fragmentation and interleaving *See* LFI
link PIC redundancy....356
link services interface
multilink bundles *See* multilink bundles
link services interfaces
CoS components....365, 367, 1151
example configuration....1152, 1159
interleave fragments....1144
example configuration....1145
link services IQ interfaces....393
CoS components....347
example configuration....383, 388
link state replication....356
link-layer overhead....367
link services protocols....1123
link state replication
LSQ PICs....356
link-layer overhead
link services IQ interfaces....367
link-layer-overhead statement....413
usage guidelines....363, 367, 377
lmi-type statement....1201
usage guidelines....1150
load balancing
on monitoring interfaces....963

Copyright © 2013, Juniper Networks, Inc.

load-balancing-options statement
aggregated Multiservices....177
local-address statement
PTSP....753
local-address-range statement
PTSP....754
local-certificate statement....289
usage guidelines....234
local-dump statement....1014
usage guidelines....949
local-gateway address statement....333
usage guidelines....315
local-gateway statement....495
usage guidelines....474
local-id statement....289
usage guidelines....236
local-policy-decision-function statement....872
local-port-range statement
PTSP....754
local-ports statement
PTSP....755
local-prefix-list statement
PTSP....755
log output
adaptive services....482
APPID....795
traffic sampling....917
log-prefix statement....495, 534
L2TP....333
usage guidelines....317, 479, 518
logging statement....205, 496
usage guidelines....190
logical interfaces
multicast-capable connections....1143
logical tunnels....1281
example configuration....1292
logical-system statement
RPM....1245
usage guidelines....1219
loopback tunnels....1283
LSQ bandwidth
oversubscribing....369
LSQ failover
interchassis....350
stateful intrachassis....352
stateless intrachassis....352
LSQ PICs....356
redundancy....352
lsq-failure-options statement....413
usage guidelines....350

M

manipulation-rule statement....697
manual security association....223
manual statement....290
usage guidelines....223
manuals
comments on....l
many-to-one statement
aggregated Multiservices....178
match direction usage in service sets....471
match statement....1015
match-direction statement
AACL....851
usage guidelines....841
CoS....457
usage guidelines....444
IDS....206
usage guidelines....187
IPsec....290
usage guidelines....243
PTSP....756
stateful firewall....146
usage guidelines....133
max-burst-size statement....603
max-checked-bytes statement....814
APPID
usage guidelines....791
max-concurrent-calls statement....604
max-connection-duration statement....1246
max-duplicates statement....1114
usage guidelines....1091
max-flows statement....496
usage guidelines....478
max-packets-per-second statement....1017
usage guidelines....913
maximum-age statement....1075
usage guidelines....1057
maximum-connections statement....1246
maximum-connections-per-client statement....1247
maximum-contexts statement....437
usage guidelines....423
maximum-fuf-percentage statement....605
maximum-inactivity-time statement....606
maximum-net-propagation-delay statement....607
maximum-packet-length statement....1016
maximum-records-in-cache statement....701

Copyright © 2013, Juniper Networks, Inc.

maximum-send-window statement....334
usage guidelines....316
maximum-sessions statement....1247
maximum-sessions-per-connection statement....1248
maximum-synchronization-mismatches statement....607
maximum-terms statement....608
maximum-time-in-cache statement....701
maximum-transactions statement
nested applications....815
maximum-waiting-delay statement....608
media statement....609
media-policy statement....698
media-type statement....699
mediation devices
flow-tap....1096
member statement
nested applications....815
member-failure-options statement
aggregated Multiservices....179
member-interface statement
aggregated Multiservices....181
message-manipulation statement....700
message-manipulation-rules statement....702
mg-maximum-pdu-size statement....610
mg-originated-pending-limit statement....611
mg-provisional-response-timer-value statement....612
mg-segmentation-timer statement....613
mgc-maximum-pdu-size statement....614
mgc-originated-pending-limit statement....615
mgc-provisional-response-timer-value statement....616
mgc-segmentation-timer statement....617
min-checked-bytes statement....816
APPID
usage guidelines....791
minimum links
link services interfaces....1141
multilink interfaces....1141
minimum statement
BGF....703
minimum-links statement....1202
usage guidelines....1141
minimum-priority statement....1114
usage guidelines....1087
MLFR and MLPPP....1127
mlfr-uni-nni-bundle-options statement....1203
usage guidelines....1147, 1150
MLPPP....380, 390
configuration example....383
example configuration....393
MLPPP (Multilink Point-to-Point Protocol)
sample topology....1169
mode statement....291
usage guidelines....233
monitor statement....618
monitoring statement....1018
usage guidelines....922
moving-average-size statement....1248
usage guidelines....1221
MPLS
packets
passive flow monitoring....970
mpls-ipv4-template statement....1019
mpls-template statement....1019
mrru statement....1204
usage guidelines....1141
mss statement....206
usage guidelines....190
mtu statement....1205
multicast traffic
AS PIC....480
multicast tunnels....1280
multicast-capable connections
Frame Relay encapsulation....1143
multicast-dlci statement....1205
usage guidelines....1143
multicast-only statement....1303
usage guidelines....1280
multiclass MLPPP
fragmentation....367
multilink bundles
fractional T1....390
example configuration....393, 395, 398
FRF.12....395
example configuration....398
MLPPP....390
example configuration....393
NxT1....380, 385
configuration example....383, 388
overview....1133
sample topology....1169
multilink interfaces
example configuration....1156
minimum links....1141

Copyright © 2013, Juniper Networks, Inc.

multilink-class statement....414
usage guidelines....367
multilink-max-classes statement....414
usage guidelines....367
multiservice-options statement....1020
MultiServices PIC
hardware requirements....37

N

n391 statement....1206
usage guidelines....1150
n392 statement....1206
usage guidelines....1150
n393 statement....1207
usage guidelines....1150
name-format statement....1076
usage guidelines....1056
name-resolution-cache statement....704
NAT
AMS....167
load balancing, example....169
overview....48
twice NAT
description....50
nat-options statement....498
nat-rule-sets statement
usage guidelines....472
nat-rules statement....498
usage guidelines....472
nested-application statement
APPID....817
usage guidelines....789
nested-application-settings statement
APPID....818
network-operator-id statement....618
networks
sample LFI and multilink bundle topology....1169
sample multilink bundle and LFI topology....1169
new-call-usage-input-policies statement....704
new-call-usage-output-policies statement....705
new-call-usage-policy statement....706
new-call-usage-policy-set statement....707
new-transaction-input-policies statement....707
new-transaction-output-policies statement....708
new-transaction-policy statement....709
new-transaction-policy-set statement....711
next-hop groups....950
next-hop statement....1020
border signaling gateway....712
next-hop groups
usage guidelines....952
usage guidelines....950
next-hop style service sets....472
next-hop-group statement
forwarding-options....1021
port mirroring....1022
usage guidelines....950, 952
next-hop-service statement....499
usage guidelines....470
no-anti-replay statement....291, 500
usage guidelines....248, 476
no-application-identification statement....818
APPID
usage guidelines....791
no-application-system-cache statement....819
APPID
usage guidelines....791
no-clear-application-system-cache statement....819
APPID
usage guidelines....791
no-core-dump statement....984
usage guidelines....920
no-dscp-bit-mirroring statement....619
no-filter-check statement....1022
usage guidelines....950
no-fragmentation statement....415
usage guidelines....365
no-ipsec-tunnel-in-traceroute statement....292
usage guidelines....255
no-local-dump statement....1014
usage guidelines....949
no-nested-application statement....820
usage guidelines....790
no-per-unit-scheduler statement....415
no-protocol-method statement....820
APPID
usage guidelines....791
no-remote-trace statement
flow monitoring....1023
no-rtcp-check statement....619
no-signature-based statement....821
APPID
usage guidelines....791
no-stamp statement....1041
usage guidelines....915

Copyright © 2013, Juniper Networks, Inc.

no-syslog statement
DFC....1115
flow monitoring....1042
usage guidelines....1090
no-termination-request statement....416
usage guidelines....350
no-world-readable statement
flow monitoring....1051
usage guidelines....915
normal-mg-execution-time statement....620
normal-mgc-execution-time statement....621
notice (system logging severity level)....317, 479, 519
notice icons defined....xlviii
Notification behavior....622
notification-behavior statement....622
notification-rate-limit statement....622
notification-regulation statement....623
notification-targets statement....1115
usage guidelines....1087
NxT1 bundles
FRF.16....385
configuration example....388
MLPPP....380
configuration example....383

O

object-cache-size statement....163
setting for stateful firewall....155
offloading flows
configuring....525
on-3xx-response statement....713
one-way-hardware-timestamp statement....1249
usage guidelines....1226
open-timeout statement....535
usage guidelines....516
option-refresh-rate statement....1024
order statement....821
APPID
usage guidelines....785
output files
logging information output file....917
traffic sampling output files....915
output statement....535
discard accounting....1025
flow monitoring....1026
port mirroring....1026
sampling....1027
usage guidelines....468, 520
output-interface-index statement....1028
outside-service-interface statement
usage guidelines....472
overload-control statement....623
oversubscription....369

P

package statement
loading on PIC....163
packages
jservices-sfw....153
packet-based IPsec....244
parentheses, in syntax descriptions....l
passive flow monitoring....903
MPLS packets....970
passive-mode-tunneling statement....500
usage guidelines....477
passive-monitor-mode statement....1029
usage guidelines....968
password statement
flow collection....1078
usage guidelines....1055, 1057
pattern statement
nested applications....822
peak-data-rate statement....624, 625
peer-unit statement
tunnel....1304
usage guidelines....1281
per-unit-scheduler statement....417
usage guidelines....369, 374, 380, 385
perfect-forward-secrecy statement....292
usage guidelines....241
performance, monitoring....1220
pgcp-rules statement
service-set....501
PIC types for services....3
pic-memory-threshold statement....1116
usage guidelines....1090
PIM
tunnels....1284
PIR....369
platform statement....626
platforms, supported....4
point-to-point connections
Frame Relay encapsulation....1143

Copyright © 2013, Juniper Networks, Inc.

policy statement
IKE....293
usage guidelines....232
IPsec....294
usage guidelines....240
policy-db-size statement....164
setting for stateful firewall....155
policy-decision-statistics-profile statement....873
pool statement
service interface pool....658
pop-all-labels statement....1030
usage guidelines....970
port mirroring....950
disabling....985
disabling all instances....985
port statement
cflowd
usage guidelines....929
flow monitoring....1031
RPM....1250
TWAMP....1250
voice services....437
usage guidelines....423
port-mapping statement....822
port-mirroring statement....1032
usage guidelines....950
port-range statement....823
APPID
usage guidelines....785
post-service-filter statement....536
usage guidelines....468
ppp-access-profile statement....334
usage guidelines....315
ppp-profile statement
usage guidelines....314
pre-shared-key statement....294
usage guidelines....234
preserve-interface statement....418
usage guidelines....356
primary statement
link services....418
usage guidelines....353
services PIC....536
usage guidelines....522
probe statement
RPM....1251
usage guidelines....1221
probe-count statement....1252
usage guidelines....1221
probe-interval statement....1252
usage guidelines....1221
probe-limit statement....1253
usage guidelines....1225
probe-server statement....1253
usage guidelines....1224
probe-type statement....1254
usage guidelines....1221
probes, for monitoring traffic....1221
procedural overview....43
profile statement
APPID
usage guidelines....788
application identification....823
profile-name statement....627
profile-version statement....627
proposal statement
IKE....295
usage guidelines....228
IPsec....296
usage guidelines....237
proposals statement
IKE....296
usage guidelines....234
IPsec....296
usage guidelines....241
protocol statement
applications....127
usage guidelines....76
IPsec....297
usage guidelines....225, 240
nested applications....824
PTSP....756
PTSP statements
application-group-any....747
application-groups....747
applications....748
count-type....748
demux....749
forward-rule
....751
forwarding instance....750
local-address....753
local-address-range....754
local-port-range....754
local-ports....755
local-prefix-list....755
match-direction....756
protocol....756

Copyright © 2013, Juniper Networks, Inc.

remote-address....757
remote-address-range....758
remote-port-range....758
remote-ports....759
remote-prefix-list....759
rule-set....761
services....762
term
forward rule....763
rule....764
then
forward rule....765
rule....766
ptsp-rule-sets statement
usage guidelines....472
ptsp-rules statement....502
usage guidelines....472

Q

queue-limit-percentage statement....628
queues statement....438
usage guidelines....423

R

RADIUS servers
configuration example....169
rate statement....537, 1033
usage guidelines....913, 950
Real-Time Performance Monitoring *See* RPM
reassemble-packets statement....1304
usage guidelines....1278
receive-options-packets statement....1033
usage guidelines....968
receive-ttl-exceeded statement....1034
usage guidelines....968
receive-window statement....335
usage guidelines....316
reconnect statement....629
red-differential-delay statement....1207
usage guidelines....1149
redistribute-all-traffic statement
aggregated Multiservices....182
redundancy
AS PIC....522
flow monitoring....975
L2TP....320
redundancy-options statement....419, 537
usage guidelines....522
reflexive | reverse statement....458
usage guidelines....446
reject-all-commands-threshold statement....629
reject-new-calls-threshold statement....630
rejoin-timeout statement
aggregated Multiservices....182
remote-address statement
PTSP....757
remote-address-range statement
PTSP....758
remote-gateway statement....297
usage guidelines....247
remote-id statement....298
usage guidelines....236
remote-port-range statement
PTSP....758
remote-ports statement
PTSP....759
remote-prefix-list statement
PTSP....759
report-service-change statement....630
request-timestamp statement....631
request-uri statement....714
request-url statement....826
required-depth statement....1034
usage guidelines....970
retransmit-interval statement....335
usage guidelines....316
retry statement....1079
usage guidelines....1057
retry-delay statement....1079
usage guidelines....1057
reverse-manipulation statement....715
RFC 2890....1276
route statement....716
route-record statement
usage guidelines....928
routing-destinations statement....717
routing-instance statement
BGF....631
RPM....1255
tunnel....1305
usage guidelines....1282
routing-instances statement
RPM....1255
usage guidelines....1220
routing-options statement....1306
rpc-program-number statement....128
usage guidelines....84

Copyright © 2013, Juniper Networks, Inc.

RPM....1215, 1217
example configuration....1232
rpm statement....1256
usage guidelines....1225
rtp statement....438, 632
usage guidelines....423
rule statement
AACL....852
usage guidelines....840
APPID
usage guidelines....785
application identification....830
BGF....633
CoS....459
usage guidelines....443
IDS....207
usage guidelines....187
IPsec....299
usage guidelines....243
PTSP....760, 761
stateful firewall....147
usage guidelines....132
rule-set statement
AACL....853
usage guidelines....844
APPID
usage guidelines....785
application identification....831
BGF....633
CoS....460
usage guidelines....448
IDS....208
usage guidelines....193
IPsec....300
usage guidelines....249
NAT....829
PTSP....761
stateful firewall....148
usage guidelines....136
run-length statement....1035
usage guidelines....913, 950

S

sample (firewall filter action)....912
sample-once statement
flow monitoring....1035
usage guidelines....915
sampled file....917
sampled.pkts file....915
sampling
logical interface....913
monitoring interface....920
sampling rate....913
sampling statement....1039
usage guidelines....912
sbc-utils statement....634, 718
scheduler map
CoS
configuration example....1145
secondary statement
link services....419
usage guidelines....353
services PIC....538
usage guidelines....522
security associations
clearing....228
segmentation statement....635
senable-asymmetic-traffic-processing statement
usage guidelines....793
send cflowd records to flow collector....1058
send-notification-on-delay statement....636
server statement....1257
server-inactivity-timeout statement....1257
servers statement....719
service filters....521
service interface configuration....468
service packages....39
service rules configuration....472
service sets
example configuration....483
overview....38
service statement....538
usage guidelines....520
service-change statement....637
service-change-type statement....638
service-class statement....720
service-domain statement....539
usage guidelines....470
service-filter statement....539
firewall
usage guidelines....521
interfaces
usage guidelines....468

Copyright © 2013, Juniper Networks, Inc.

service-interface statement....336, 502
BGF....638
border signaling gateway
gateway....721
service point....721
usage guidelines....315, 468
service-interface-pools statement....658
service-point statement....722
service-point-type statement....723
service-policies statement....723
service-port statement....1116
usage guidelines....1087
service-set statement....503, 540
usage guidelines....467, 520
service-state statement
virtual BGF....639
virtual interface in BGF....640
services configuration overview....43
services PICs....3
services statement
AACL
usage guidelines....839
APPID
usage guidelines....781
BGF....640
border signaling gateway....724
CoS....460
usage guidelines....441
DFC....1117
usage guidelines....1085
flow monitoring
usage guidelines....911
flow-monitoring....1040
IDS....208
usage guidelines....185
interfaces....541
usage guidelines....518
IPsec....300
usage guidelines....219
L2TP....337
usage guidelines....317
PTSP....762
RPM....1258
usage guidelines....1217
service sets....505
usage guidelines....479
stateful firewall....148
usage guidelines....131
services-options statement....542
usage guidelines....516, 518
session-limit statement....209, 543
usage guidelines....190
session-mirroring statement....641
session-timeout statement....833
usage guidelines....783
session-trace statement....725
shaping-rate statement
usage guidelines....363, 369, 374
shared-key statement....1117
usage guidelines....1087
short-sequence statement....1208
usage guidelines....1142
signaling statement....726
signaling-realms statement
border signaling gateway
new transaction policy....727
signature statement
nested applications....834
signature-method-all-ports statement....834
AAPID
usage guidelines....791
sip statement....461, 728
usage guidelines....446
sip-call-hold-timeout statement....128
sip-header statement....731
sip-stack statement....733
size statement....1040
usage guidelines....917
snmp-command statement....129
usage guidelines....84
soft-limit statement....1118
usage guidelines....1086
soft-limit-clear statement....1118
usage guidelines....1086
softwire-rules statement
usage guidelines....473
SONET interfaces
sampling SONET interfaces....917
source statement
APPID
usage guidelines....785
application identification rule....835
encryption....897
tunnel....1306
usage guidelines....889, 1273

Copyright © 2013, Juniper Networks, Inc.

source-address statement
AACL....854
usage guidelines....841
BGF....641
CoS....461
usage guidelines....444
flow monitoring....1041
usage guidelines....922
IDS....210
usage guidelines....189
IPsec....301
usage guidelines....244
RPM....1258
usage guidelines....1221
service-set system log....543
stateful firewall....149
usage guidelines....133
tunnel....1307
tunnel services
usage guidelines....1286
source-address-range statement
AACL....854
usage guidelines....841
IDS....210
usage guidelines....189
stateful firewall....149
usage guidelines....133
source-addresses statement
DFC....1119
usage guidelines....1087
source-port statement
BGF....642
RPM....129
usage guidelines....79
source-prefix statement....211
usage guidelines....190
source-prefix-ipv6 statement....211
usage guidelines....190
source-prefix-list statement
AACL....855
usage guidelines....841
CoS....462
IDS....212
stateful firewall....150
usage guidelines....133
spi statement....301
usage guidelines....225
stamp option....916
stamp statement....1041
usage guidelines....915
state-loss statement....643
stateful firewall
action statements....134
anomalies....46
applications....133
example configuration....136
match conditions....133
restrictions....155
rules....136
stateful firewall plug-in
configuring memory for....155
stateful firewall use with APPID....786
stateful firewalls
jservices-sfw package....153
SDK Kerberos-enabled, configuring....155
SDK plug-in for, loading....153
stateful-firewall-rule-sets statement
usage guidelines....473
stateful-firewall-rules statement....507
usage guidelines....473
statement
flow monitoring
usage guidelines....917
IPsec
usage guidelines....255
L2TP
usage guidelines....322
services
usage guidelines....1058
statistics statement
L-PDF....874
stop-detection-on-drop statement....643
support, technical *See* technical support
support-uni-directional-traffic statement....835
usage guidelines....793
sustained-data-rate statement....644
gate in packet gateway....645
syn-cookie statement....212
usage guidelines....190
syntax conventions....xlix
syslog statement....165
CoS....462
usage guidelines....445
flow monitoring....1042
IDS....213
usage guidelines....190

Copyright © 2013, Juniper Networks, Inc.

interfaces....544
usage guidelines....518
IPsec....302
usage guidelines....246, 249
L2TP....339
usage guidelines....317
service sets....507
usage guidelines....479
stateful firewall....150
usage guidelines....134

T

t391 statement....1208
usage guidelines....1150
t392 statement....1209
usage guidelines....1150
target statement....1259
RPM....1259
target-url statement
usage guidelines....1221
tcp statement
RPM....1259
tcp-mss statement....508
tcp-tickles statement....544, 545
technical support
contacting JTAC....I
template statement
flow monitoring....1043
template-refresh-rate statement....1045
term statement
AACL....856
usage guidelines....840
border signaling gateway
new call usage policy....735
new transaction policy....736
service-class....737
CoS....463
usage guidelines....443
HCM....827
IDS....214
usage guidelines....187
IPsec....303
usage guidelines....243
PTSP
forward rule....763
rule....764
stateful firewall....151
usage guidelines....132
test statement
RPM....1260
usage guidelines....1221
test-interval statement....1261
usage guidelines....1221
then statement
AACL....857
usage guidelines....840
border signaling gateway
new call usage policy....738
new transaction policy....739
service-class....740
CoS....464
usage guidelines....443
HCM....827
IDS....216
usage guidelines....187
IPsec....304
usage guidelines....243
PTSP
forward rule....765
rule....766
stateful firewall....152
usage guidelines....132, 134
threshold statement....217
usage guidelines....190
thresholds statement
RPM....1262
usage guidelines....1221
time-to-live threshold....84
timer-c statement....741
timers statement....741
timerx statement....646
timestamp option....916
tmax-retransmission-delay statement....646
topology
sample LFI and multilink bundle network....1169
sample multilink bundle and LFI network....1169
trace-options
server (tracing flag)....482
timer-events (tracing flag)....482
traceoptions statement....742
application identification....836
BGF....647
flow monitoring....1045
IPsec....305
L-PDF....875
L2TP....340

Copyright © 2013, Juniper Networks, Inc.

security....307
services....509
tracing flags
event policy
all....482, 795
configuration....482
database....482
events....482
policy....482
server....482
timer-events....482
tracing operations
adaptive services....480
APPID....794
traffic....883
inbound (decryption)....887
IPsec, configuring....883
monitoring....1220
outbound (encryption)....885
traffic sampling
configuring....912
disabling....915, 985
example configurations....917
flow aggregation....928
FTP traffic....919
logging information output file....917
output files....915
SONET interfaces....917
traffic from single IP addresses....918
traffic-control-profiles statement
usage guidelines....369, 374
traffic-management statement....648
transactions statement....743
transfer statement....1080
usage guidelines....1056
transfer-log-archive statement....1080
usage guidelines....1057
transport-details statement....744
traps statement....1263
usage guidelines....1221
trigger-link-failure statement....420
usage guidelines....350
trusted-ca statement....510
usage guidelines....475
ttl statement
DFC....1120
usage guidelines....1086
tunnel....1307
usage guidelines....1273
ttl-threshold statement....130
usage guidelines....84
tunnel interfaces
configuration statements....1273, 1280, 1283
dynamic tunnels....1286
example configuration....1288
logical tunnels....1281
loopback tunnels....1283
multicast tunnels....1280
PIM tunnels....1284
unicast tunnels....1273
tunnel statement....1308
encryption....898
usage guidelines....881
redundancy
usage guidelines....889
unicast
usage guidelines....1273
tunnel-group statement....344
usage guidelines....314
tunnel-mtu statement....308, 511
usage guidelines....249, 477
tunnel-timeout statement....345
usage guidelines....316
tunnels
definition....1269
GRE
fragmentation of....1276
key number....1276
interface types....1269
IPv6-over-IPv4....1285, 1290
twamp statement....1264
twamp-server statement....1265
twice NAT....50
type statement....837
APPID
usage guidelines....783
type-of-service statement....837
APPID
usage guidelines....783

U

udp statement
RPM....1265
undirectional traffic support
APPID....793
unicast tunnels....1273

Copyright © 2013, Juniper Networks, Inc.

unit statement
aggregated Multiservices....183
encryption....899
usage guidelines....881
flow monitoring....1046
usage guidelines....912
interfaces....546
usage guidelines....513
link services....439, 1210
usage guidelines....1127
tunnel....1309
usage guidelines....1273
Universal Unique Identifier....84
up statement
BGF....649
url statement....838
APPID
usage guidelines....792
url-rule statement....829
url_identifier, statement....828
use-lower-case statement....649
use-wildcard-response statement....650
username statement
flow collection....1081
usage guidelines....1057
uuid statement....130
usage guidelines....84

V

var/log/sampled file....917
var/tmp/sampled.pkts file....915
variant statement....1081
usage guidelines....1056
version statement
flow monitoring....1047
IKE....308
usage guidelines....233, 929
version-ipfix statement....1050
usage guidelines....942, 946
version9 statement....1048
video statement....464
usage guidelines....446
virtual loopback tunnel
configuration guidelines....1282
VRF table lookup
example configuration....1289
virtual-interface statement....651
virtual-interface-down statement....652
virtual-interface-indications statement....653
virtual-interface-up statement....653
voice services
bundles....426
configuration....421
encapsulation....425
example configuration....426
interface type....422
voice services interfaces
interleave fragments....424
voice statement....465
usage guidelines....446

W

warm standby
AS PIC....522
LSQ PIC....352
warm statement....654
warm-standby statement....420
warning (system logging severity level)....317, 479, 518
wired-process-mem-size statement....166
world-readable statement
flow monitoring....1051
usage guidelines....915

Y

yellow-differential-delay statement....1211
usage guidelines....1149

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Index of Statements and Commands

A

aacl-fields statement....868
aacl-statistics-profile statement....869
accelerations statement....667
accounting statement
  flow monitoring....978
acknowledge-retries statement....1189
acknowledge-timer statement....1190
action-red-differential-delay statement....1191
actions statement....666
adaptive-services-pics statement....485
address statement
  application rule....800
  DFC....1103
  encryption....891
  flow monitoring....979
  interfaces....527
  link services....1192
  voice services....431
administrative statement
  BGF....554
admission-control statement....668, 669
aggregate-export-interval statement....979
aggregation statement....197
  flow monitoring....980
algorithm statement....556
allow-fragmentation statement....1295
allow-ip-options statement....142
allow-multicast statement....486
allowed-destinations statement....1104
analyzer-address statement....1065
analyzer-id statement....1066
anti-replay-window-size statement....273, 487
application statement....121, 801, 802
application-aware-access-list-fields
  statement....870
application-data-inactivity-detection
  statement....556
application-group statement....802
application-group-any statement....848
  PTSP....747
application-groups statement....803, 848
  PTSP....747
application-profile statement....452
application-protocol statement....122
application-set statement....123
application-sets statement
  CoS....453
  IDS....198
  stateful firewall....143
application-system-cache-timeout
  statement....803
applications statement
  AACL....847
  application identification....804
  application-level gateways....143
  applications hierarchy....123
  CoS....453
  IDS....198
  PTSP....748
  stateful firewall....143
archive-sites statement....1066
audit-observed-events-returns statement....557
authentication statement....274
authentication-algorithm statement
  IKE....275
  IPsec....275
authentication-method statement....276
authentication-mode statement
  RPM....1237
automatic statement....804
autonomous-system-type statement....981
auxiliary-spi statement....276
availability-check-profiles statement....670

B

backup-destination statement....1296
backup-interface statement....892
backup-remote-gateway statement....277
base-root statement....558
bgf-core statement....559
bgp statement
  RPM....1238
blacklist-period statement....671
bundle statement....432, 1192

Copyright © 2013, Juniper Networks, Inc.

by-destination statement....199
by-pair statement....200
by-source statement....201
bypass-traffic-on-pic-failure statement....488

C

cancel-graceful statement....561
capture-group statement....1105
cflowd statement....982
cgn-pic statement....528
chain-order statement
    nested applications....805
cisco-interoperability statement....409
cleanup-timeout statement....562
clear-dont-fragment-bit statement
    GRE tunnel....528
    IPsec....277
    service-set....489
clear-ike-sas-on-pic-restart statement....278
clear-ipsec-sas-on-pic-restart statement....278
client-list statement....1239
clusters statement....672
collector statement....1067
committed-burst-size statement....673
committed-information-rate statement....674
compression statement....432
compression-device statement....433
content-destination statement....1106
context statement
    nested applications....805
context-indications statement....563
control-association-indications statement....564
control-cores statement....157
control-source statement....1107
controller-address statement....565
controller-failure statement....565
controller-port statement....566
copy-tos-to-outer-ip-header statement....1296
core-dump statement....984
count-type statement
    PTSP....748

D

data statement....454
data-cores statement....158
data-fill statement....1239
data-flow-affinity statement....158
data-format statement....1067
data-inactivity-detection statement....566, 674
data-size statement....1240
datastore statement....675
default statement....567
default-media-realm statement....676
delivery-function statement....568
demux statement
    PTSP....749
description statement
    IKE....279
    IPsec....279
destination statement....159, 893, 1193
    application identification rule....806
    tunnel....1297
destination-address statement
    AACL....849
    BGF....568
    CoS....454
    IDS....202
    IPsec....279
    stateful firewall....144
destination-address-range statement
    AACL....849
    IDS....202
    stateful firewall....144
destination-networks statement
    tunnel....1298
destination-port statement
    applications....123
    BGF....569
    RPM....124, 1242
destination-prefix statement....203
destination-prefix-ipv6 statement....203
destination-prefix-list statement
    AACL....850
    CoS....455
    IDS....204
    stateful firewall....145
destinations statement
    flow collection....1068
detect statement....569
dh-group statement....280
dial-options statement....529
dialogs statement....677
diffserv statement....570
direction statement....281
    nested applications....806
disable statement
    application....807
    application group....807

Copyright © 2013, Juniper Networks, Inc.

flow monitoring....985
port mapping....808
disable-all-instances statement
flow monitoring....985
disable-global-timeout-override statement....808
disable-mlppp-inner-ppp-pfc statement....1194
disable-session-mirroring statement....570
disconnect statement....571
dlci statement....1194
do-not-fragment statement
tunnel....1299
down statement....572
download statement
APPID....809
drop-member-traffic statement
aggregated Multiservices....173
drop-timeout statement....1195
dscp statement....455
BGF....573
BSG....678
dscp-code-point statement
RPM....1243
duplicates-dropped-periodicity statement....1107
dynamic route insertion....251
dynamic statement....282
dynamic-flow-capture statement....1108
dynamic-tunnels statement....1300

E

egress-service-point statement....679
embedded-spdf statement....680
enable-asymmetic-traffic-processing
statement....810
enable-heuristics statement....809, 810
enable-rejoin statement
aggregated Multiservices....174
encapsulation statement....433
link services....1196
encoding statement....573
encryption statement....283
encryption-algorithm statement
IKE....284
IPsec....284
engine-id statement
flow monitoring....986
engine-type statement....987
es-options statement....894
event-timestamp-notification statement....574
export-format statement....989
extension-provider statement....160
extension-service statement....988

F

f-max-period statement....434
facility-override statement....329, 490, 530
failover statement....576
failover-cold statement....574
failover-warm statement....575
family statement
aggregated Multiservices....174
encryption....895
flow monitoring....990
interfaces....531
link services....1198
voice services....435
fast-update-filters statement....577
file statement....995
BGF....578
border signaling gateway....681
L-PDF statistics....871
traffic sampling....995
file-specification statement....1069
filename statement....996
filename-prefix statement....1068
files statement....996
filter statement
encryption....896
flow monitoring....997
flag statement....579, 682
flow-active-timeout statement....998
flow-collector statement....1070
flow-control-options statement....999
flow-export-destination statement....1000
flow-export-rate statement
flow monitoring....999
flow-inactive-timeout statement....1001
flow-monitoring statement....1002
flow-server statement
flow monitoring....1003
flow-tap statement....1109
force-entry statement....204
forward-manipulation statement....683
forward-rule statement
PTSP....751
forwarding instance....750
forwarding-class statement....410, 456
forwarding-db-size statement....161
forwarding-options statement....1004

Copyright © 2013, Juniper Networks, Inc.

fragment-threshold statement
link services....1199
LSQ....411
voice services....436
fragmentation-map statement....411
fragmentation-maps statement....412
framework statement....684
from statement
AACL....850
border signaling gateway
new call usage policy....687
new transaction policy....688
service class....690
CoS....456
HCM....825
IDS....205
IPsec....285
PTSP....752
PTSP forward rule....751
stateful firewall....146
ftp statement....457
flow collection....1072

G

g-duplicates-dropped-periodicity statement....1110
g-max-duplicates statement....1111
gateway statement
BGF....580
border signaling gateway....691
gateway-address statement....584
gateway-controller statement....585
gateway-port statement....586
graceful statement....587
graceful-restart statement....588

H

h248-options statement....589
h248-profile statement....591
h248-properties statement....592
h248-stack statement....595
h248-timers statement....596
hanging-termination-detection statement....596
hard-limit statement....1111
hard-limit-target statement....1112
hardware-timestamp statement....1244
hash-key statement....162
hcm statement....825
hello-interval statement
L2TP....330
hello-timer statement
link services....1200
hide-avps statement....331
high-availability-options statement
aggregated Multiservices....175
history-size statement....1244
hold-time statement
GRE tunnel interface....1301
host statement....491, 532
HCM....826
L2TP....331
hot-standby statement....412

I

icmp-code statement....124
icmp-type statement....125
idle-timeout statement....811
ids-rules statement....491
ignore-entry statement....204
ignore-errors statement....811
ike statement....286
ike-access-profile statement....492
inactivity-delay statement....597
inactivity-duration statement....597, 696
inactivity-non-tcp-timeout statement....812
inactivity-tcp-timeout statement....812
inactivity-timeout statement....125
BGF....598
flow monitoring....532
RPM....1245
inactivity-timer statement....599
index statement....813
nested applications....813
initial-average-ack-delay statement....599
initiate-dead-peer-detection statement....287
inline-jflow statement
flow monitoring....1005
input statement
flow monitoring....1006
interfaces....533
input-interface-index statement....1007
input-packet-rate-threshold statement....1112
instance statement
port mirroring....1008
sampling....1009
interface statement
flow monitoring....1011
flow-tap....1113
service interface pool....657

Copyright © 2013, Juniper Networks, Inc.

interface-map statement....1074
interface-service statement....492
interfaces statement
aggregated Multiservices....176
DFC....1113
encryption....896
flow monitoring....1013
interfaces hierarchy....533
link services....1200
tunnel....1301
voice services....436
interim-ah-scheme statement....600
interleave-fragments statement....1201
ip statement
application identification....814
ip-flow-stop-detection statement....600
ip-reassembly-rules statement
service-set....493
ipsec statement....287
ipsec-inside-interface statement....288
ipsec-sa statement
encryption....897
ipsec-transport-security-association statement....601
ipsec-vpn-options statement....494
ipsec-vpn-rules statement....494
ipv4-template statement....1013
ipv6-template statement....1013

K

keepalive-time statement
GRE tunnel interface....1302
key statement
tunnel....1303

L

L2TP statements
traceoptions....340
l2tp-access-profile statement....332
label-position statement....1014
latch-deadlock-delay statement....601
learn-sip-register statement....126
lifetime-seconds statement
IKE....288
IPsec....288
link-layer-overhead statement....413
lmi-type statement....1201
load-balancing-options statement
aggregated Multiservices....177
local-address statement
PTSP....753
local-address-range statement
PTSP....754
local-certificate statement....289
local-dump statement....1014
local-gateway address statement....333
local-gateway statement....495
local-id statement....289
local-policy-decision-function statement....872
local-port-range statement
PTSP....754
local-ports statement
PTSP....755
local-prefix-list statement
PTSP....755
log-prefix statement....495, 534
L2TP....333
logging statement....205, 496
logical-system statement
RPM....1245
lsq-failure-options statement....413

M

manipulation-rule statement....697
manual statement....290
many-to-one statement
aggregated Multiservices....178
match statement....1015
match-direction statement
AACL....851
CoS....457
IDS....206
IPsec....290
PTSP....756
stateful firewall....146
max-burst-size statement....603
max-checked-bytes statement....814
max-concurrent-calls statement....604
max-connection-duration statement....1246
max-duplicates statement....1114
max-flows statement....496
max-packets-per-second statement....1017
maximum-age statement....1075
maximum-connections statement....1246
maximum-connections-per-client statement....1247
maximum-contexts statement....437
maximum-fuf-percentage statement....605
maximum-inactivity-time statement....606

Copyright © 2013, Juniper Networks, Inc.

maximum-net-propagation-delay statement.....607
maximum-records-in-cache statement.....701
maximum-send-window statement.....334
maximum-sessions statement.....1247
maximum-sessions-per-connection statement.....1248
maximum-synchronization-mismatches statement.....607
maximum-terms statement.....608
maximum-time-in-cache statement.....701
maximum-transactions statement
nested applications.....815
maximum-waiting-delay statement.....608
media statement.....609
media-policy statement.....698
media-type statement.....699
member statement
nested applications.....815
member-failure-options statement
aggregated Multiservices.....179
member-interface statement
aggregated Multiservices.....181
message-manipulation statement.....700
message-manipulation-rules statement.....702
mg-maximum-pdu-size statement.....610
mg-originated-pending-limit statement.....611
mg-provisional-response-timer-value statement.....612
mg-segmentation-timer statement.....613
mgc-maximum-pdu-size statement.....614
mgc-originated-pending-limit statement.....615
mgc-provisional-response-timer-value statement.....616
mgc-segmentation-timer statement.....617
min-checked-bytes statement.....816
minimum statement
BGF.....703
minimum-links statement.....1202
minimum-priority statement.....1114
mlfr-uni-nni-bundle-options statement.....1203
mode statement.....291
monitor statement.....618
monitoring statement.....1018
moving-average-size statement.....1248
mpls-ipv4-template statement.....1019
mpls-template statement.....1019
mrru statement.....1204
mss statement.....206
mtu statement.....1205
multicast-dlci statement.....1205
multicast-only statement.....1303
multilink-class statement.....414
multilink-max-classes statement.....414
multiservice-options statement.....1020

N

n391 statement.....1206
n392 statement.....1206
n393 statement.....1207
name-format statement.....1076
name-resolution-cache statement.....704
nat-options statement.....498
nat-rules statement.....498
nested-application statement
APPID.....817
nested-application-settings statement
APPID.....818
network-operator-id statement.....618
new-call-usage-input-policies statement.....704
new-call-usage-output-policies statement.....705
new-call-usage-policy statement.....706
new-call-usage-policy-set statement.....707
new-transaction-input-policies statement.....707
new-transaction-output-policies statement.....708
new-transaction-policy statement.....709
new-transaction-policy-set statement.....711
next-hop statement.....1020
border signaling gateway.....712
next-hop-group statement
forwarding-options.....1021
port mirroring.....1022
next-hop-service statement.....499
no-anti-replay statement.....291, 500
no-application-identification statement.....818
no-application-system-cache statement.....819
no-clear-application-system-cache statement.....819
no-core-dump statement.....984
no-dscp-bit-mirroring statement.....619
no-filter-check statement.....1022
no-fragmentation statement.....415
no-ipsec-tunnel-in-traceroute statement.....292
no-local-dump statement.....1014
no-nested-application statement.....820
no-per-unit-scheduler statement.....415
no-protocol-method statement.....820
no-remote-trace statement
flow monitoring.....1023

Copyright © 2013, Juniper Networks, Inc.

no-rtcp-check statement....619
no-signature-based statement....821
no-stamp statement....1041
no-syslog statement
DFC....1115
flow monitoring....1042
no-termination-request statement....416
no-world-readable statement
flow monitoring....1051
normal-mg-execution-time statement....620
normal-mgc-execution-time statement....621
notification-behavior statement....622
notification-rate-limit statement....622
notification-regulation statement....623
notification-targets statement....1115

O

object-cache-size statement....163
on-3xx-response statement....713
one-way-hardware-timestamp statement....1249
open-timeout statement....535
option-refresh-rate statement....1024
order statement....821
output statement....535
discard accounting....1025
flow monitoring....1026
port mirroring....1026
sampling....1027
output-interface-index statement....1028
overload-control statement....623

P

package statement
loading on PIC....163
passive-mode-tunneling statement....500
passive-monitor-mode statement....1029
password statement
flow collection....1078
pattern statement
nested applications....822
peak-data-rate statement....624, 625
peer-unit statement
tunnel....1304
per-unit-scheduler statement....417
perfect-forward-secrecy statement....292
pgcp-rules statement
service-set....501
pic-memory-threshold statement....1116
platform statement....626
policy statement
IKE....293
policy-db-size statement....164
policy-decision-statistics-profile statement....873
pool statement
service interface pool....658
pop-all-labels statement....1030
port statement
flow monitoring....1031
RPM....1250
TWAMP....1250
voice services....437
port-mapping statement....822
port-mirroring statement....1032
port-range statement....823
post-service-filter statement....536
ppp-access-profile statement....334
pre-shared-key statement....294
preserve-interface statement....418
primary statement
link services....418
services PIC....536
probe statement
RPM....1251
probe-count statement....1252
probe-interval statement....1252
probe-limit statement....1253
probe-server statement....1253
probe-type statement....1254
profile statement
application identification....823
profile-name statement....627
profile-version statement....627
proposal statement
IKE....295
IPsec....296
proposals statement
IKE....296
IPsec....296
protocol statement
applications....127
IPsec....297
nested applications....824
PTSP....756
PTSP statements
application-group-any....747
application-groups....747
applications....748
count-type....748

Copyright © 2013, Juniper Networks, Inc.

demux....749
forward-rule
....751
forwarding instance....750
local-address....753
local-address-range....754
local-port-range....754
local-ports....755
local-prefix-list....755
match-direction....756
protocol....756
remote-address....757
remote-address-range....758
remote-port-range....758
remote-ports....759
remote-prefix-list....759
rule-set....761
services....762
term
forward rule....763
rule....764
then
forward rule....765
rule....766
ptsp-rules statement....502

Q

queue-limit-percentage statement....628
queues statement....438

R

rate statement....537, 1033
reassemble-packets statement....1304
receive-options-packets statement....1033
receive-ttl-exceeded statement....1034
receive-window statement....335
reconnect statement....629
red-differential-delay statement....1207
redistribute-all-traffic statement
aggregated Multiservices....182
redundancy-options statement....419, 537
reflexive | reverse statement....458
reject-all-commands-threshold statement....629
reject-new-calls-threshold statement....630
rejoin-timeout statement
aggregated Multiservices....182
remote-address statement
PTSP....757
remote-address-range statement
PTSP....758
remote-gateway statement....297
remote-id statement....298
remote-port-range statement
PTSP....758
remote-ports statement
PTSP....759
remote-prefix-list statement
PTSP....759
report-service-change statement....630
request-timestamp statement....631
request-uri statement....714
request-url statement....826
required-depth statement....1034
retransmit-interval statement....335
retry statement....1079
retry-delay statement....1079
reverse-manipulation statement....715
route statement....716
routing-destinations statement....717
routing-instance statement
BGF....631
RPM....1255
tunnel....1305
routing-instances statement
RPM....1255
routing-options statement....1306
rpc-program-number statement....128
rpm statement....1256
rtp statement....438, 632
rule statement
AACL....852
application identification....830
BGF....633
CoS....459
IDS....207
IPsec....299
PTSP....760, 761
stateful firewall....147
rule-set statement
AACL....853
application identification....831
BGF....633
CoS....460
IDS....208
IPsec....300
NAT....829

Copyright © 2013, Juniper Networks, Inc.

PTSP....761
stateful firewall....148
run-length statement....1035

S

sample-once statement
flow monitoring....1035
sampling statement....1039
sbc-utils statement....634, 718
secondary statement
link services....419
services PIC....538
segmentation statement....635
send-notification-on-delay statement....636
server statement....1257
server-inactivity-timeout statement....1257
servers statement....719
service statement....538
service-change statement....637
service-change-type statement....638
service-class statement....720
service-domain statement....539
service-filter statement....539
service-interface statement....336, 502
BGF....638
border signaling gateway
gateway....721
service point....721
service-interface-pools statement....658
service-point statement....722
service-point-type statement....723
service-policies statement....723
service-port statement....1116
service-set statement....503, 540
service-state statement
virtual BGF....639
virtual interface in BGF....640
services statement
BGF....640
border signaling gateway....724
CoS....460
DFC....1117
flow-monitoring....1040
IDS....208
interfaces....541
IPsec....300
L2TP....337
PTSP....762
RPM....1258
service sets....505
stateful firewall....148
services-options statement....542
session-limit statement....209, 543
session-mirroring statement....641
session-timeout statement....833
session-trace statement....725
shared-key statement....1117
short-sequence statement....1208
signaling statement....726
signaling-realms statement
border signaling gateway
new transaction policy....727
signature statement
nested applications....834
sip statement....461, 728
sip-call-hold-timeout statement....128
sip-header statement....731
sip-stack statement....733
size statement....1040
snmp-command statement....129
soft-limit statement....1118
soft-limit-clear statement....1118
source statement
application identification rule....835
encryption....897
tunnel....1306
source-address statement
AACL....854
BGF....641
CoS....461
flow monitoring....1041
IDS....210
IPsec....301
RPM....1258
stateful firewall....149
tunnel....1307
source-address-range statement
AACL....854
IDS....210
stateful firewall....149
source-addresses statement
DFC....1119
source-port statement
BGF....642
RPM....129
source-prefix statement....211
source-prefix-ipv6 statement....211

Copyright © 2013, Juniper Networks, Inc.

source-prefix-list statement
AACL....855
CoS....462
IDS....212
stateful firewall....150
spi statement....301
stamp statement....1041
state-loss statement....643
stateful-firewall-rules statement....507
statistics statement
L-PDF....874
stop-detection-on-drop statement....643
support-uni-directional-traffic statement....835
sustained-data-rate statement....644
gate in packet gateway....645
syn-cookie statement....212
syslog statement....165
CoS....462
flow monitoring....1042
IDS....213
interfaces....544
IPsec....302
L2TP....339
service sets....507
stateful firewall....150

T

t391 statement....1208
t392 statement....1209
target statement....1259
RPM....1259
tcp statement
RPM....1259
tcp-mss statement....508
tcp-tickles statement....544, 545
template statement
flow monitoring....1043
template-refresh-rate statement....1045
term statement
AACL....856
border signaling gateway
new call usage policy....735
new transaction policy....736
service-class....737
CoS....463
HCM....827
IDS....214
IPsec....303
PTSP
forward rule....763
rule....764
stateful firewall....151
test statement
RPM....1260
test-interval statement....1261
then statement
AACL....857
border signaling gateway
new call usage policy....738
new transaction policy....739
service-class....740
CoS....464
HCM....827
IDS....216
IPsec....304
PTSP
forward rule....765
rule....766
stateful firewall....152
threshold statement....217
thresholds statement
RPM....1262
timer-c statement....741
timers statement....741
timerx statement....646
tmax-retransmission-delay statement....646
traceoptions statement....742
application identification....836
BGF....647
flow monitoring....1045
IPsec....305
L-PDF....875
L2TP....340
security....307
services....509
traffic-management statement....648
transactions statement....743
transfer statement....1080
transfer-log-archive statement....1080
transport-details statement....744
traps statement....1263
trigger-link-failure statement....420
trusted-ca statement....510
ttl statement
DFC....1120
tunnel....1307
ttl-threshold statement....130

Copyright © 2013, Juniper Networks, Inc.

tunnel statement....1308
encryption....898
tunnel-group statement....344
tunnel-mtu statement....308, 511
tunnel-timeout statement....345
twamp statement....1264
twamp-server statement....1265
type statement....837
type-of-service statement....837

U

udp statement
RPM....1265
unit statement
aggregated Multiservices....183
encryption....899
flow monitoring....1046
interfaces....546
link services....439, 1210
tunnel....1309
up statement
BGF....649
url statement....838
url-rule statement....829
url_identifier, statement....828
use-lower-case statement....649
use-wildcard-response statement....650
username statement
flow collection....1081
uuid statement....130

V

variant statement....1081
version statement
flow monitoring....1047
IKE....308
version-ipfix statement....1050
version9 statement....1048
video statement....464
virtual-interface statement....651
virtual-interface-down statement....652
virtual-interface-indications statement....653
virtual-interface-up statement....653
voice statement....465

W

warm statement....654
warm-standby statement....420
wired-process-mem-size statement....166
world-readable statement
flow monitoring....1051

Y

yellow-differential-delay statement....1211

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.