Juniper Networks JUNOS® Software

Services Interfaces Configuration Guide

Release 10.0

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

Published: 2009-10-07

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, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark 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.

Juniper Networks JUNOS® Software Services Interfaces Configuration Guide
Release 10.0
Copyright © 2009, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Alan Twhigg, Justine Kangas, Myron Weintraub
Editing: Benjamin Mann, Joanne McClintock
Illustration: Nathaniel Woodward
Cover Design: Edmonds Design

Revision History
October 2009—R1 JUNOS 10.0

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at http://www.gnu.org/licenses/gpl.html, and a copy of the LGPL at http://www.gnu.org/licenses/lgpl.html.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

About This Guide xlix

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 33
Chapter 4 Applications Configuration Guidelines 59
Chapter 5 Summary of Applications Configuration Statements 97
Chapter 6 Stateful Firewall Services Configuration Guidelines 107
Chapter 7 Summary of Stateful Firewall Configuration Statements 117
Chapter 8 Network Address Translation Services Configuration Guidelines 129
Chapter 9 Summary of Network Address Translation Configuration Statements 155
Chapter 10 Intrusion Detection Service Configuration Guidelines 179
Chapter 11 Summary of Intrusion Detection Service Configuration Statements 191
Chapter 12 IPsec Services Configuration Guidelines 213
Chapter 13 Summary of IPsec Services Configuration Statements 257
Chapter 14 Layer 2 Tunneling Protocol Services Configuration Guidelines 291
Chapter 15 Summary of Layer 2 Tunneling Protocol Configuration Statements 309
Chapter 16 Link Services IQ Interfaces Configuration Guidelines 323
Chapter 17 Summary of Link Services IQ Configuration Statements 383
Chapter 18 Voice Services Configuration Guidelines 395
Chapter 19 Summary of Voice Services Configuration Statements 405
Chapter 20 Class-of-Service Configuration Guidelines 415
Chapter 21 Summary of Class-of-Service Configuration Statements 425
Chapter 22 Service Set Configuration Guidelines 439
Chapter 23 Summary of Service Set Configuration Statements 457
Chapter 24 Service Interface Configuration Guidelines 479
Chapter 25 Summary of Service Interface Configuration Statements 493
Chapter 26 PGCP Configuration Guidelines for the BGF Feature 511
Chapter 27 Summary of PGCP Configuration Statements 517
Chapter 28 Service Interface Pools Configuration Guidelines 613
Chapter 29 Summary of Service Interface Pools Statements 615
Chapter 30 Border Signaling Gateway Configuration Guidelines 617

Chapter 31 Summary of Border Signaling Gateway Configuration Statements 621

Part 3 Dynamic Application Awareness for JUNOS Software

Chapter 32 Dynamic Application Awareness for JUNOS Software Overview 687
Chapter 33 Application Identification Configuration Guidelines 691
Chapter 34 Summary of Application Identification Configuration Statements 701
Chapter 35 Application-Aware Access List Configuration Guidelines 725
Chapter 36 Summary of AACL Configuration Statements 731
Chapter 37 Local Policy Decision Function Configuration Guidelines 743
Chapter 38 Summary of L-PDF Configuration Statements 749

Part 4 Encryption Services

Chapter 39 Encryption Overview 761
Chapter 40 Encryption Interfaces Configuration Guidelines 763
Chapter 41 Summary of Encryption Configuration Statements 773

Part 5 Flow Monitoring and Discard Accounting Services

Chapter 42 Flow Monitoring and Discard Accounting Overview 783
Chapter 43 Flow Monitoring and Discard Accounting Configuration Guidelines 787
Chapter 44 Summary of Flow-Monitoring Configuration Statements 841
Chapter 45 Flow Collection Configuration Guidelines 899
Chapter 46 Summary of Flow Collection Configuration Statements 911
Chapter 47 Dynamic Flow Capture Configuration Guidelines 929
Chapter 48 Flow-Tap Configuration Guidelines 941
Chapter 49 Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements 949

Part 6 Link and Multilink Services

Chapter 50 Link and Multilink Services Overview 969
Chapter 51 Link and Multilink Services Configuration Guidelines 973
Chapter 52 Summary of Multilink and Link Services Configuration Statements 1011

Part 7 Real-Time Performance Monitoring Services

Chapter 53 Real-Time Performance Monitoring Services Overview 1035
Chapter 54 Real-Time Performance Monitoring Configuration Guidelines 1037
Chapter 55 Summary of Real-Time Performance Monitoring Configuration Statements 1055

Part 8 Tunnel Services

Chapter 56 Tunnel Services Overview 1083

Chapter 57 Tunnel Interfaces Configuration Guidelines 1087
Chapter 58 Summary of Tunnel Services Configuration Statements 1103

Part 9 Index

Index 1119
Index of Statements and Commands 1139

Table of Contents

About This Guide **xlix**

JUNOS Documentation and Release Notes xlix
Objectives l
Audience l
Supported Platforms l
Using the Indexes li
Using the Examples in This Manual li
    Merging a Full Example li
    Merging a Snippet lii
Documentation Conventions lii
Documentation Feedback liv
Requesting Technical Support liv
    Self-Help Online Tools and Resources lv
    Opening a Case with JTAC lv

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 11
[edit services] Hierarchy Level 11

Part 2 Adaptive Services

Chapter 3 Adaptive Services Overview 33

Enabling Service Packages 35
Layer 2 Service Package Capabilities and Interfaces 38
Services Configuration Procedure 39
Packet Flow Through the Adaptive Services or Multiservices PIC 40
Stateful Firewall Overview 41
Stateful Firewall Support for Application Protocols 42
Stateful Firewall Anomaly Checking 42
Network Address Translation Overview 44
Traditional NAT 44
Twice NAT 45
IPsec Overview 45
IPsec 46
Security Associations 46
IKE 46
Comparison of IPsec Services and ES Interface Configuration 47
Layer 2 Tunneling Protocol Overview 48
Voice Services Overview 48
Class of Service Overview 49
Examples: Services Interfaces Configuration 49

Chapter 4 Applications Configuration Guidelines 59

Configuring Application Protocol Properties 60
Configuring an Application Protocol 60
Configuring the Network Protocol 62
Configuring the ICMP Code and Type 63
Configuring Source and Destination Ports 65
Configuring the Inactivity Timeout Period 68
Configuring SIP 68
Configuring an SNMP Command for Packet Matching 69
Configuring an RPC Program Number 69
Configuring the TTL Threshold 69
Configuring a Universal Unique Identifier 70
Configuring Application Sets 70
ALG Descriptions 70
Basic TCP ALG 71
Basic UDP ALG 71
BOOTP 72
DCE RPC Services 72
FTP 72
H323 73
ICMP 73
IIOP 74
NetShow 74
RealAudio 74

RPC and RPC Portmap Services 74
RTSP 76
SMB 76
SNMP 76
SQLNet 77
TFTP 77
Traceroute 77
UNIX Remote-Shell Services 77
WinFrame 78
Verifying the Output of ALG Sessions 78
FTP Example 78
Sample Output 78
FTP System Log Messages 79
Analysis 80
Troubleshooting Questions 80
RTSP ALG Example 81
Sample Output 81
Analysis 81
Troubleshooting Questions 82
System Log Messages 83
System Log Configuration 84
System Log Output 85
JUNOS Default Groups 85
Examples: Referencing the Preset Statement from the JUNOS Default Group 91
Examples: Configuring Application Protocols 93

Chapter 5 Summary of Applications Configuration Statements 97

application 97
application-protocol 98
application-set 99
applications 99
destination-port 100
icmp-code 100
icmp-type 101
inactivity-timeout 101
learn-sip-register 102
protocol 103
rpc-program-number 104
sip-call-hold-timeout 104
snmp-command 105
source-port 105
ttl-threshold 106
uuid 106

Chapter 6 Stateful Firewall Services Configuration Guidelines 107

Configuring Stateful Firewall Rules 108
Configuring Match Direction for Stateful Firewall Rules 108
Configuring Match Conditions in Stateful Firewall Rules 109
Configuring Actions in Stateful Firewall Rules 110
Configuring IP Option Handling 111
Configuring Stateful Firewall Rule Sets 112
Examples: Configuring Stateful Firewall Rules 112

Chapter 7 Summary of Stateful Firewall Configuration Statements 117

allow-ip-options 118
application-sets 119
applications 119
destination-address 120
destination-address-range 120
destination-prefix-list 121
from 121
match-direction 122
rule 123
rule-set 124
services 124
source-address 125
source-address-range 125
source-prefix-list 126
syslog 126
term 127
then 128

Chapter 8 Network Address Translation Services Configuration Guidelines 129

Configuring Addresses and Ports for Use in NAT Rules 130
Configuring Pools of Addresses and Ports 130
Specifying Destination and Source Prefixes when Pools Are Not Used 132
Requirements for NAT Addresses 132
Configuring IPv6 Multicast Filters 133
Configuring NAT Rules 133
Configuring Match Direction for NAT Rules 134
Configuring NAT Type for Terms in NAT Rules 135
Configuring Match Conditions in NAT Rules 136
Configuring Actions in NAT Rules 137
Configuring NAT Rule Sets 139
Examples: Configuring NAT Rules 139
Example: Configuring Dynamic Source Translation 140
Example: Configuring Static Source Translation 140
Example: Configuring Dynamic and Static Source Translation 141

Example: Configuring an Oversubscribed Pool with No Fallback 141
Example: Configuring an Oversubscribed Pool with Fallback to NAPT 142
Example: Configuring Static Source Translation with Multiple Prefixes and Address Ranges 142
Example: Assigning Addresses from a Dynamic Pool for Static Use 143
Example: Configuring NAT Rules Without Defining a Pool 144
Example: Preventing Translation of Specific Addresses 144
Example: Configuring NAT for Multicast Traffic 145
Rendezvous Point Configuration 145
Router 1 Configuration 148
Example: Configuring Twice NAT 149
Example: Configuring Full-Cone NAT 150
Example: Configuring NAT in Mixed IPv4 and IPv6 Networks 150

Chapter 9 Summary of Network Address Translation Configuration Statements 155

address 155
address-range 156
application-sets 156
applications 157
destination-address 157
destination-address-range 158
destination-pool 158
destination-prefix 159
destination-prefix-list 159
from 160
hint 161
ipv6-multicast-interfaces 161
match-direction 162
nat-type 162
no-translation 163
overload-pool 163
overload-prefix 164
pgcp 164
pool 165
port 166
ports-per-session 166
remotely-controlled 167
rule 168
rule-set 169
services 169
source-address 170
source-address-range 170
source-pool 171
source-prefix 171
source-prefix-list 172
syslog 172
term 173

then 174
translated 175
translation-type 176
translation-type (Traditional NAT) 176
translation-type (Twice NAT) 176
transport 177

Chapter 10 Intrusion Detection Service Configuration Guidelines 179

Configuring IDS Rules 181
Configuring Match Direction for IDS Rules 182
Configuring Match Conditions in IDS Rules 183
Configuring Actions in IDS Rules 184
Configuring IDS Rule Sets 187
Examples: Configuring IDS Rules 188

Chapter 11 Summary of Intrusion Detection Service Configuration Statements 191

aggregation 191
application-sets 192
applications 192
by-destination 193
by-pair 194
by-source 195
destination-address 196
destination-address-range 196
destination-prefix 197
destination-prefix-ipv6 197
destination-prefix-list 198
force-entry 198
from 199
ignore-entry 199
logging 200
match-direction 200
mss 201
rule 202
rule-set 203
services 203
session-limit 204
source-address 205
source-address-range 205
source-prefix 206
source-prefix-ipv6 206
source-prefix-list 207
syn-cookie 207
syslog 208
term 209
then 211
threshold 212

Chapter 12 IPsec Services Configuration Guidelines 213

Minimum Security Association Configurations 215
Minimum Manual SA Configuration 215
Minimum Dynamic SA Configuration 215
Configuring Security Associations 216
Configuring Manual Security Associations 217
Configuring the Direction for IPsec Processing 217
Configuring the Protocol for a Manual IPsec SA 218
Configuring the Security Parameter Index 219
Configuring the Auxiliary Security Parameter Index 219
Configuring Authentication for a Manual IPsec SA 219
Configuring Encryption for a Manual IPSec SA 220
Configuring Dynamic Security Associations 221
Clearing Security Associations 222
Configuring IKE Proposals 222
Configuring the Authentication Algorithm for an IKE Proposal 223
Configuring the Authentication Method for an IKE Proposal 223
Configuring the Diffie-Hellman Group for an IKE Proposal 224
Configuring the Encryption Algorithm for an IKE Proposal 224
Configuring the Lifetime for an IKE SA 225
Example: Configuring an IKE Proposal 225
Configuring IKE Policies 225
Configuring the Mode for an IKE Policy 227
Configuring the Proposals in an IKE Policy 227
Configuring the Preshared Key for an IKE Policy 227
Configuring the Local Certificate for an IKE Policy 228
Configuring a Certificate Revocation List 228
Configuring the Description for an IKE Policy 229
Configuring Local and Remote IDs for IKE Phase 1 Negotiation 229
Example: Configuring an IKE Policy 230
Configuring IPsec Proposals 231
Configuring the Authentication Algorithm for an IPsec Proposal 231
Configuring the Description for an IPsec Proposal 231
Configuring the Encryption Algorithm for an IPsec Proposal 232
Configuring the Lifetime for an IPsec SA 232
Configuring the Protocol for a Dynamic SA 233
Configuring IPsec Policies 233
Configuring the Description for an IPsec Policy 234
Configuring Perfect Forward Secrecy 234
Configuring the Proposals in an IPsec Policy 234
Example: Configuring an IPsec Policy 235
Configuring IPsec Rules 235
Configuring Match Direction for IPsec Rules 236
Configuring Match Conditions in IPsec Rules 237
Configuring Actions in IPsec Rules 239
Enabling IPsec Packet Fragmentation 240
Configuring Destination Addresses for Dead Peer Detection 240
Configuring or Disabling IPSec Anti-Replay 241

Enabling System Log Messages 242
Specifying the MTU for IPsec Tunnels 242
Configuring IPsec Rule Sets 242
Configuring Dynamic Endpoints for IPsec Tunnels 242
Authentication Process 243
Implicit Dynamic Rules 244
Reverse Route Insertion 244
Configuring an IKE Access Profile 245
Referencing the IKE Access Profile in a Service Set 246
Configuring the Interface Identifier 247
Default IKE and IPsec Proposals 247
Tracing IPsec Operations 248
Disabling IPsec Tunnel Endpoint in Traceroute 249
Examples: Configuring IPsec Services 249
Example: Configuring Statically Assigned Tunnels 249
Example: Configuring Dynamically Assigned Tunnels 252

Chapter 13 Summary of IPsec Services Configuration Statements 257

anti-replay-window-size 257
authentication 258
authentication-algorithm 259
authentication-algorithm (IKE) 259
authentication-algorithm (IPsec) 259
authentication-method 260
auxiliary-spi 260
backup-remote-gateway 261
clear-dont-fragment-bit 261
clear-ike-sas-on-pic-restart 262
clear-ipsec-sas-on-pic-restart 262
description 263
destination-address 263
dh-group 264
direction 265
dynamic 266
encryption 267
encryption-algorithm 268
from 269
ike 270
initiate-dead-peer-detection 271
ipsec 271
ipsec-inside-interface 272
lifetime-seconds 272
local-certificate 273
local-id 273
manual 274
match-direction 274
mode 275
no-anti-replay 275
no-ipsec-tunnel-in-traceroute 276

perfect-forward-secrecy 276
policy 277
policy (IKE) 277
policy (IPsec) 278
pre-shared-key 278
proposal 279
proposal (IKE) 279
proposal (IPsec) 280
proposals 280
protocol 281
remote-gateway 281
remote-id 282
rule 283
rule-set 284
services 284
source-address 285
spi 285
syslog 286
term 287
then 288
traceoptions 289
tunnel-mtu 290

Chapter 14 Layer 2 Tunneling Protocol Services Configuration Guidelines 291

L2TP Services Configuration Overview 293
L2TP Minimum Configuration 294
Configuring L2TP Tunnel Groups 296
Configuring Access Profiles for L2TP Tunnel Groups 297
Configuring the Local Gateway Address and PIC 297
Configuring Window Size for L2TP Tunnels 298
Configuring Timers for L2TP Tunnels 298
Hiding Attribute-Value Pairs for L2TP Tunnels 299
Configuring System Logging of L2TP Tunnel Activity 299
Configuring the Identifier for Logical Interfaces that Provide L2TP Services 300
Example: Configuring Multilink PPP on a Shared Logical Interface 301
AS PIC Redundancy for L2TP Services 302
Tracing L2TP Operations 303
Examples: Configuring L2TP Services 304

Chapter 15 Summary of Layer 2 Tunneling Protocol Configuration Statements 309

facility-override 309
hello-interval 310
hide-avps 310
host 311
l2tp-access-profile 311
local-gateway address 312

log-prefix 312
maximum-send-window 313
ppp-access-profile 313
receive-window 314
retransmit-interval 314
service-interface 315
services 316
services (Hierarchy) 316
services (L2TP System Logging) 317
syslog 318
traceoptions 319
tunnel-group 321
tunnel-timeout 322

Chapter 16 Link Services IQ Interfaces Configuration Guidelines 323

Layer 2 Service Package Capabilities and Interfaces 324
Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS 326
Configuring the Association between LSQ and SONET Interfaces 326
Configuring SONET APS Interoperability with Cisco Systems FRF.16 327
Restrictions on APS Redundancy for LSQ Interfaces 328
Configuring LSQ Interface Redundancy in a Single Router Using SONET APS 328
Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces 328
Configuring Redundant Paired LSQ Interfaces 329
Restrictions on Redundant LSQ Interfaces 330
Configuring Link State Replication for Redundant Link PICs 331
Examples: Configuring Redundant LSQ Interfaces for Failure Recovery 333
Configuring CoS Scheduling Queues on Logical LSQ Interfaces 337
Configuring Scheduler Buffer Size 338
Configuring Scheduler Priority 339
Configuring Scheduler Shaping Rate 339
Configuring Drop Profiles 339
Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces 341
Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces 343
Configuring Multiclass MLPPP on LSQ Interfaces 343
Oversubscribing Interface Bandwidth on LSQ Interfaces 344
Example: Oversubscribing an LSQ Interface 347
Configuring Guaranteed Minimum Rate on LSQ Interfaces 348
Example: Configuring Guaranteed Minimum Rate 351
Configuring Link Services and CoS on Services PICs 351
Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP 354
Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP 357
Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 359
Example: Configuring an LSQ Interface as an NxT1 Bundle Using FRF.16 363

Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI 365
Example: Configuring an LSQ Interface for a Fractional T1 Interface Using MLPPP and LFI 368
Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 369
Examples: Configuring an LSQ Interface for a Fractional T1 Interface Using FRF.12 372
Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 377
Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP 378
Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 379
Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP 381

Chapter 17 Summary of Link Services IQ Configuration Statements 383

cisco-interoperability 383
forwarding-class 384
fragment-threshold 385
fragmentation-map 385
fragmentation-maps 386
hot-standby 386
link-layer-overhead 387
lsq-failure-options 387
multilink-class 388
multilink-max-classes 388
no-fragmentation 389
no-termination-request 389
per-unit-scheduler 390
preserve-interface 390
primary 391
redundancy-options 391
secondary 392
trigger-link-failure 392
warm-standby 393

Chapter 18 Voice Services Configuration Guidelines 395

Configuring Services Interfaces for Voice Services 396
Configuring the Logical Interface Address for the MLPPP Bundle 396
Configuring Compression of Voice Traffic 397
Configuring Delay-Sensitive Packet Interleaving 398
Example: Configuring Compression of Voice Traffic 398
Configuring Encapsulation for Voice Services 399
Configuring Network Interfaces for Voice Services 399
Configuring Voice Services Bundles with MLPPP Encapsulation 400
Configuring the Compression Interface with PPP Encapsulation 400
Examples: Configuring Voice Services 400

Chapter 19 Summary of Voice Services Configuration Statements 405

address 405
bundle 406
compression 406
compression-device 407
encapsulation 407
f-max-period 408
family 409
fragment-threshold 410
interfaces 410
maximum-contexts 411
port 411
queues 412
rtp 412
unit 413

Chapter 20 Class-of-Service Configuration Guidelines 415

Restrictions and Cautions for CoS Configuration on Services Interfaces 416
Configuring CoS Rules 416
Configuring Match Direction for CoS Rules 417
Configuring Match Conditions In CoS Rules 418
Configuring Actions in CoS Rules 419
Configuring Application Profiles for Use as CoS Rule Actions 420
Configuring Reflexive and Reverse CoS Rule Actions 420
Example: Configuring CoS Rules 421
Configuring CoS Rule Sets 421
Examples: Configuring CoS on Services Interfaces 422

Chapter 21 Summary of Class-of-Service Configuration Statements 425

application-profile 425
application-sets 426
applications 426
destination-address 427
destination-prefix-list 427
dscp 428
forwarding-class 428
from 429
match-direction 429
(reflexive | reverse) 430
rule 431
rule-set 432
services 432
sip-text 433
sip-video 433
sip-voice 434

source-address 434
source-prefix-list 435
syslog 435
term 436
then 437

Chapter 22 Service Set Configuration Guidelines 439

Configuring Service Sets to be Applied to Services Interfaces 440
Configuring Interface Service Sets 440
Configuring Next-Hop Service Sets 441
Determining Traffic Direction 442
Interface Style Service Sets 443
Next-Hop Style Service Sets 443
Configuring Service Rules 444
Configuring IPsec Service Sets 445
Configuring the Local Gateway Address for IPsec Service Sets 445
IKE Addresses in VRF Instances 446
Configuring IKE Access Profiles for IPsec Service Sets 446
Configuring Certification Authorities for IPsec Service Sets 447
Configuring or Disabling Antireplay Service 447
Clearing the Don't-Fragment Bit 448
Configuring Passive-Mode Tunneling 449
Configuring the Tunnel MTU Value 449
Configuring Service Set Limitations 450
Configuring System Logging for Service Sets 450
Enabling Services PICs to Accept Multicast Traffic 452
Tracing Services PIC Operations 452
Configuring the Adaptive Services Log Filename 453
Configuring the Number and Size of Adaptive Services Log Files 453
Configuring Access to the Log File 453
Configuring a Regular Expression for Lines to Be Logged 454
Configuring the Trace Operations 454
Example: Configuring Service Sets 455

Chapter 23 Summary of Service Set Configuration Statements 457

adaptive-services-pics 457
allow-multicast 458
anti-replay-window-size 459
clear-dont-fragment-bit 460
facility-override 461
host 461
ids-rules 462
ike-access-profile 462
interface-service 463
ipsec-vpn-options 463
ipsec-vpn-rules 464
local-gateway 464
log-prefix 465

logging 465
max-flows 466
nat-rules 466
next-hop-service 467
no-anti-replay 468
passive-mode-tunneling 468
pgcp-rules 469
service-interface 469
service-set 470
services 472
    services (Hierarchy) 472
    services (System Logging) 473
stateful-firewall-rules 474
syslog 474
tcp-mss 475
traceoptions 476
trusted-ca 477
tunnel-mtu 478

Chapter 24 Service Interface Configuration Guidelines 479

Services Interface Naming Overview 480
Configuring the Address and Domain for Services Interfaces 481
Configuring Default Timeout Settings for Services Interfaces 482
Configuring System Logging for Services Interfaces 483
Enabling Fragmentation on GRE Tunnels 484
Applying Filters and Services to Interfaces 485
    Configuring Service Filters 486
Configuring AS or Multiservices PIC Redundancy 487
Examples: Configuring Services Interfaces 490

Chapter 25 Summary of Service Interface Configuration Statements 493

address 493
clear-dont-fragment-bit 494
dial-options 495
facility-override 496
family 497
host 498
inactivity-timeout 498
input 499
interfaces 499
log-prefix 500
open-timeout 500
output 501
post-service-filter 501
primary 502
redundancy-options 502
secondary 503
service 503

service-domain 504
service-filter 504
service-set 505
services 506
services-options 507
syslog 508
unit 509

Chapter 26 PGCP Configuration Guidelines for the BGF Feature 511

Chapter 27 Summary of PGCP Configuration Statements 517

administrative 518
administrative (Control Association) 518
administrative (Virtual Interface) 519
algorithm 519
application-data-inactivity-detection 520
audit-observed-events-returns 520
base-root 521
bgf-core 522
cancel-graceful 523
cancel-graceful (Control Association) 523
cancel-graceful (Virtual Interface) 524
cleanup-timeout 524
context-indications 525
control-association-indications 526
controller-address 526
controller-failure 527
controller-port 527
data-inactivity-detection 528
default 529
delivery-function 530
destination-address 530
destination-port 531
detect 531
diffserv 532
disable-session-mirroring 532
disconnect 533
down 533
dscp 534
encoding 534
event-timestamp-notification 535
failover-cold 535
failover-warm 536
failure 537
fast-update-filters 538
file 539
flag 540
gateway 541
gateway-address 545

gateway-controller 546
gateway-port 547
graceful 548
    graceful (Control Association) 548
    graceful (Virtual Interface) 549
graceful-restart 549
h248-options 550
h248-profile 551
h248-properties 552
h248-stack 555
h248-timers 556
hanging-termination-detection 556
inactivity-delay 557
inactivity-duration 557
inactivity-timeout 558
inactivity-timer 558
initial-average-ack-delay 559
interim-ah-scheme 559
ip-flow-stop-detection 560
ipsec-transport-security-association 560
latch-deadlock-delay 561
max-burst-size 562
    max-burst-size (All Streams) 562
    max-burst-size (RTCP Streams) 563
max-concurrent-calls 564
maximum-fuf-percentage 565
maximum-inactivity-time 566
maximum-net-propagation-delay 567
maximum-synchronization-mismatches 567
maximum-terms 568
maximum-waiting-delay 568
media 569
media-service 569
mg-maximum-pdu-size 570
mg-originated-pending-limit 571
mg-provisional-response-timer-value 572
mg-segmentation-timer 573
mgc-maximum-pdu-size 574
mgc-originated-pending-limit 575
mgc-provisional-response-timer-value 576
mgc-segmentation-timer 577
monitor 578
nat-pool 578
network-operator-id 579
no-dscp-bit-mirroring 579
no-rtcp-check 580
normal-mg-execution-time 580
normal-mgc-execution-time 581
notification-behavior 582
notification-rate-limit 582
notification-regulation 583

overload-control 583
peak-data-rate 584
peak-data-rate (All Streams) 584
peak-data-rate (RTCP) 585
platform 586
profile-name 586
profile-version 587
queue-limit-percentage 587
reconnect 588
reject-all-commands-threshold 588
reject-new-calls-threshold 589
report-service-change 589
request-timestamp 590
routing-instance 590
rtcp 591
rtp 591
rule 592
rule-set 592
sbc-utils 593
segmentation 594
send-notification-on-delay 595
service-change 596
service-change-type 597
service-interface 597
service-state 598
service-state (Virtual BGF) 598
service-state (Virtual Interface) 599
services 599
session-mirroring 600
source-address 600
source-port 601
state-loss 601
stop-detection-on-drop 602
sustained-data-rate 603
sustained-data-rate (All Streams) 603
sustained-data-rate (RTCP Streams) 604
timerx 605
tmax-retransmission-delay 605
traceoptions 606
traffic-management 607
up 608
use-lower-case 608
use-wildcard-response 609
virtual-interface 609
virtual-interface-down 610
virtual-interface-indications 610
virtual-interface-up 611
warm 611

Chapter 28 Service Interface Pools Configuration Guidelines 613

Configuring Service Interface Pools 613

Chapter 29 Summary of Service Interface Pools Statements 615

interface 615
pool 616
service-interface-pools 616

Chapter 30 Border Signaling Gateway Configuration Guidelines 617

Chapter 31 Summary of Border Signaling Gateway Configuration Statements 621

actions 621
accelerations 622
admission-control 623
    admission-control (Border Signaling Gateway) 623
    admission-control (New Transaction Policy) 624
blacklist-period 624
committed-burst-size 625
committed-information-rate 626
data-inactivity-detection 626
datastore 627
default-media-realm 628
dialogs 629
dscp 630
egress-service-point 630
embedded-spdf 631
file 632
flag 633
forward-manipulation 634
framework 635
from 637
    from (New Call Usage Policy) 638
    from (New Transaction Policy) 639
    from (Service Class) 640
gateway 641
inactivity-duration 645
manipulation-rule 646
media-policy 647
media-type 647
message-manipulation 648
maximum-records-in-cache 648
maximum-time-in-cache 649
message-manipulation-rules 650
minimum 651

name-resolution-cache 652
new-call-usage-policies 652
new-call-usage-policy 653
new-call-usage-policy-set 654
new-transaction-policies 654
new-transaction-policy 655
new-transaction-policy-set 656
next-hop 657
request-uri 658
reverse-manipulation 659
route 660
sbc-utils 661
service-class 662
service-interface 663
    service-interface (Gateway) 663
    service-interface (Service Point) 663
service-point 664
service-point-type 664
service-policies 665
services 665
session-trace 666
signaling 667
sip 669
sip-header 671
sip-stack 673
term 675
    term (New Call Usage Policy) 675
    term (New Transaction Policy) 676
    term (Service Class) 677
then 678
    then (New Call Usage Policy) 678
    then (New Transaction Policy) 679
    then (Service Class) 680
timer-c 680
timers 681
traceoptions 682
transactions 683
transport-details 684

Part 3 Dynamic Application Awareness for JUNOS Software

Chapter 32 Dynamic Application Awareness for JUNOS Software Overview 687

IDP Overview 687
APPID Overview 688
AACL Overview 688
L-PDF Overview 689

Chapter 33 Application Identification Configuration Guidelines 691

Defining an Application Identification 692
Configuring APPID Rules 694
Configuring Application Profiles 695
Configuring Application Groups 695
Configuring Global APPID Properties 696
Configuring Automatic Download of Application Package Updates 697
Tracing APPID Operations 697
Configuring the APPID Log Filename 698
Configuring the Number and Size of APPID Log Files 698
Configuring Access to the Log File 699
Configuring a Regular Expression for Lines to Be Logged 699
Configuring the Tracing Flags 699
Examples: Configuring Application Identification Properties 699

Chapter 34 Summary of Application Identification Configuration Statements 701

address 701
application 702
application (Defining) 702
application (Including in Rule) 703
application-group 703
application-groups 704
application-system-cache-timeout 704
applications 705
automatic 705
destination 706
disable 707
disable (APPID Application) 707
disable (APPID Application Group) 707
disable (APPID Port Mapping) 708
disable-global-timeout-override 708
download 709
idle-timeout 709
inactivity-non-tcp-timeout 710
inactivity-tcp-timeout 710
index 711
ip 711
max-checked-bytes 712
min-checked-bytes 712
no-application-identification 713
no-application-system-cache 713
no-clear-application-system-cache 713
no-signature-based 714
order 714
port-mapping 715
port-range 715
profile 716

rule 717
rule (Configuring) 717
rule (Including in Rule Set) 718
rule-set 718
services 719
session-timeout 720
session-timeout (interfaces) 720
session-timeout (application-identification) 721
....721
source 721
traceoptions 722
type 723
type-of-service 723
url 724

Chapter 35 Application-Aware Access List Configuration Guidelines 725

Configuring AACL Rules 726
Configuring Match Direction for AACL Rules 726
Configuring Match Conditions in AACL Rules 727
Configuring Actions in AACL Rules 728
Configuring AACL Rule Sets 729
Example: Configuring AACL Rules 729

Chapter 36 Summary of AACL Configuration Statements 731

applications 731
application-groups 732
application-group-any 732
destination-address 733
destination-address-range 733
destination-prefix-list 734
from 734
match-direction 735
rule 736
rule-set 737
services 737
source-address 738
source-address-range 738
source-prefix-list 739
term 740
then 741

Chapter 37 Local Policy Decision Function Configuration Guidelines 743

Configuring Statistics Profiles 743
Configuring an L-PDF Statistics Profile 743
Configuring an AACL Statistics Profile 744
Applying L-PDF Profiles to Service Sets 746
Tracing L-PDF Operations 747

Chapter 38 Summary of L-PDF Configuration Statements 749

aacl-fields 750
aacl-statistics-profile 751
application-aware-access-list-fields 752
file 753
local-policy-decision-function 754
policy-decision-statistics-profile 755
statistics 756
traceoptions 757

Part 4 Encryption Services

Chapter 39 Encryption Overview 761

Chapter 40 Encryption Interfaces Configuration Guidelines 763

Configuring Encryption Interfaces 763
Specifying the Security Association Name for Encryption Interfaces 764
Configuring the MTU for Encryption Interfaces 764
Example: Configuring an Encryption Interface 764
Configuring Filters for Traffic Transiting the ES PIC 765
Traffic Overview 765
Configuring the Security Association 766
Configuring an Outbound Traffic Filter 767
Example: Configuring an Outbound Traffic Filter 767
Applying the Outbound Traffic Filter 768
Example: Applying the Outbound Traffic Filter 768
Configuring an Inbound Traffic Filter 768
Example: Configuring an Inbound Traffic Filter 769
Applying the Inbound Traffic Filter to the Encryption Interface 769
Example: Applying the Inbound Traffic Filter to the Encryption Interface 769
Configuring an ES Tunnel Interface for a Layer 3 VPN 770
Configuring ES PIC Redundancy 770
Example: Configuring ES PIC Redundancy 771
Configuring IPsec Tunnel Redundancy 771

Chapter 41 Summary of Encryption Configuration Statements 773

address 773
backup-destination 774
backup-interface 774
destination 775
es-options 775
family 776
filter 777
interfaces 777
ipsec-sa 778
source 778
tunnel 779
unit 780

Part 5 Flow Monitoring and Discard Accounting Services

Chapter 42 Flow Monitoring and Discard Accounting Overview 783

Passive Flow Monitoring 783
Active Flow Monitoring 784

Chapter 43 Flow Monitoring and Discard Accounting Configuration Guidelines 787

Configuring Traffic Sampling 791
Minimum Configuration for Traffic Sampling 791
Configuring Traffic Sampling 792
Disabling Traffic Sampling 793
Sampling Once 794
Configuring Traffic Sampling Output 794
Traffic Sampling Output Format 795
Tracing Traffic Sampling Operations 796
Traffic Sampling Examples 796
Example: Sampling a Single SONET/SDH Interface 796
Example: Sampling All Traffic from a Single IP Address 797
Example: Sampling All FTP Traffic 798
Configuring Flow Monitoring 799
Configuring Flow-Monitoring Interfaces 799
Configuring Flow-Monitoring Properties 800
Directing Traffic to Flow-Monitoring Interfaces 801
Exporting Flows 801
Configuring Time Periods when Flow Monitoring is Active and Inactive 802
Example: Configuring Flow Monitoring 803
Enabling Flow Aggregation 804

Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd 804
Configuring Flow Aggregation to Use Version 9 Flow Templates 807
Configuring the Traffic to be Sampled 808
Configuring the Version 9 Template Properties 808
Restrictions 809
Fields Included in Each Template Type 810
MPLS Sampling Behavior 811
Verification 812
Examples: Configuring Version 9 Flow Templates 812
Configuring Sampling Instances 814
Directing Replicated Flows to Multiple Flow Servers 815
Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers 815
Directing Replicated Version 9 Flow Aggregates to Multiple Servers 816
Logging cflowd Flows Before Export 817
Configuring Port Mirroring 818
Configuring Tunnels 819
Port Mirroring with Next-Hop Groups 820
Filter-Based Forwarding with Multiple Monitoring Interfaces 821
Restrictions 821
Configuring Port Mirroring on Services Interfaces 822
Examples: Configuring Port Mirroring 823
Load Balancing Among Multiple Monitoring Interfaces 828
Configuring Discard Accounting 831
Enabling Passive Flow Monitoring 833
Passive Flow Monitoring for MPLS Encapsulated Packets 834
Removing MPLS Labels from Incoming Packets 834
Example: Enabling Passive Flow Monitoring 835
Configuring Services Interface Redundancy with Flow Monitoring 837

Chapter 44 Summary of Flow-Monitoring Configuration Statements 841

accounting 842
address 843
aggregate-export-interval 843
aggregation 844
autonomous-system-type 845
cflowd 846
cflowd (Discard Accounting and Sampling) 846
cflowd (Flow Monitoring) 847
core-dump 847
destination 848
disable 848
disable-all-instances 849
engine-id 849
engine-type 850
export-format 850
family 851
family (Interfaces) 851
family (Monitoring) 852

family (Port Mirroring) 853
family (Sampling) 853
file 854
file (Sampling) 854
file (Trace Options) 854
filename 855
files 855
filter 856
flow-active-timeout 857
flow-export-destination 858
flow-inactive-timeout 858
flow-monitoring 859
forwarding-options 860
input 861
input (Port Mirroring) 861
input (Sampling) 861
input-interface-index 862
instance 863
instance (Port Mirroring) 863
instance (Sampling) 864
interface 865
interface (Accounting or Sampling) 865
interface (Monitoring) 866
interface (Port Mirroring) 866
interfaces 867
ipv4-template 867
ipv6-template 867
label-position 868
local-dump 868
max-packets-per-second 869
monitoring 870
mpls-ipv4-template 871
mpls-template 871
multiservice-options 872
next-hop 872
next-hop-group 873
next-hop-group (Forwarding Options) 873
next-hop-group (Port Mirroring) 874
no-core-dump 874
no-filter-check 874
no-local-dump 874
no-stamp 875
no-syslog 875
no-world-readable 875
option-refresh-rate 875
output 876
output (Accounting) 876
output (Monitoring) 877
output (Port Mirroring) 877
output (Sampling) 878
output-interface-index 879

passive-monitor-mode 879
pop-all-labels 880
port 880
port-mirroring 881
rate 882
receive-options-packets 882
receive-ttl-exceeded 883
required-depth 883
run-length 884
sample-once 884
sampling 885
sampling (Forwarding Options) 886
sampling (Interfaces) 887
services 887
size 888
source-address 889
stamp 889
syslog 890
template 891
template (Forwarding Options) 891
template (Services) 892
template-refresh-rate 893
traceoptions 893
unit 894
version 895
version9 896
version9 (Forwarding Options) 896
version9 (Services) 897
world-readable 898

Chapter 45 Flow Collection Configuration Guidelines 899

Configuring Flow Collection 901
Configuring Destination FTP Servers for Flow Records 901
Configuring a Packet Analyzer 902
Configuring File Formats 902
Configuring Interface Mappings 903
Configuring Transfer Logs 903
Configuring Retry Attempts 904
Sending cflowd Records to Flow Collector Interfaces 904
Configuring Flow Collection Mode and Interfaces on Services PICs 904
Example: Configuring Flow Collection 905

Chapter 46 Summary of Flow Collection Configuration Statements 911

analyzer-address 911
analyzer-id 911
archive-sites 912
collector 912
data-format 913

destinations 913
filename-prefix 914
file-specification 915
file-specification (File Format) 915
file-specification (Interface Mapping) 915
flow-collector 916
ftp 918
ftp (Flow Collector Files) 919
ftp (Transfer Log Files) 920
interface-map 920
maximum-age 921
name-format 922
password 924
password (Flow Collector File Servers) 924
password (Transfer Log File Servers) 924
retry 925
retry-delay 925
transfer 926
transfer-log-archive 926
username 927
variant 927

Chapter 47 Dynamic Flow Capture Configuration Guidelines 929

Dynamic Flow Capture Architecture 929
Liberal Sequence Windowing 930
Configuring Dynamic Flow Capture 930
Configuring the Capture Group 931
Configuring the Content Destination 932
Configuring the Control Source 933
Configuring the DFC PIC Interface 934
Configuring System Logging 935
Configuring Thresholds 935
Limiting the Number of Duplicates of a Packet 936
Example: Configuring Dynamic Flow Capture 937

Chapter 48 Flow-Tap Configuration Guidelines 941

Flow-Tap Architecture 942
Configuring the Flow-Tap Service 943
Configuring the Flow-Tap Interface 943
Strengthening Flow-Tap Security 944
Restrictions on Flow-Tap Services 944
Configuring FlowTapLite 945
Examples: Configuring Flow-Tap Services 946

Chapter 49 **Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements** **949**

address 949
allowed-destinations 950
capture-group 951
content-destination 952
control-source 953
duplicates-dropped-periodicity 953
dynamic-flow-capture 954
flow-tap 955
flow-tap-lite 955
g-duplicates-dropped-periodicity 956
g-max-duplicates 956
hard-limit 957
hard-limit-target 957
input-packet-rate-threshold 958
interface 958
interfaces 959
max-duplicates 959
minimum-priority 960
no-syslog 960
notification-targets 961
pic-memory-threshold 961
service-port 962
services 962
shared-key 963
soft-limit 963
soft-limit-clear 964
source-addresses 964
ttl 965

Part 6 **Link and Multilink Services**

Chapter 50 **Link and Multilink Services Overview** **969**

Chapter 51 **Link and Multilink Services Configuration Guidelines** **973**

Multilink and Link Services PICs Overview 974
Configuring the Number of Bundles on Link Services PICs 975
Configuring the Links in a Multilink or Link Services Bundle 976
Multilink and Link Services Logical Interface Configuration Overview 977
Default Settings for Multilink and Link Services Logical Interfaces 978
Configuring Encapsulation for Multilink and Link Services Logical Interfaces 979

Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces 979
Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces 981
Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces 982
Configuring MRRU on Multilink and Link Services Logical Interfaces 982
Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces 983
Configuring DLCIs on Link Services Logical Interfaces 984
Configuring Point-to-Point DLCIs for MLFR FRF.16 and MLPPP Bundles 984
Configuring Multicast-Capable DLCIs for MLFR FRF.16 Bundles 984
Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces 985
Configuring LFI with DLCI Scheduling 986
Example: Configuring LFI with DLCI Scheduling 986
Configuring Link Services Physical Interfaces 988
Default Settings for Link Services Interfaces 988
Configuring Encapsulation for Link Services Physical Interfaces 989
Configuring Acknowledgment Timers on Link Services Physical Interfaces 989
Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 990
Configuring Keepalives on Link Services Physical Interfaces 991
Configuring CoS on Link Services Interfaces 992
CoS for Link Services Interfaces on M Series and T Series Routers 992
Example: Configuring CoS on Link Services Interfaces 993
Examples: Configuring Multilink Interfaces 997
Example: Configuring a Multilink Interface with MLPPP 997
Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces 998
Example: Configuring a Multilink Interface with MLFR FRF.15 999
Examples: Configuring Link Interfaces 1000
Example: Configuring a Link Services Interface with Two Links 1001
Example: Configuring a Link Services Interface with MLPPP 1002
Example: Configuring a Link Services Interface with MLFR FRF.15 1003
Example: Configuring a Link Services PIC with MLFR FRF.16 1003
Example: Configuring Link and Voice Services Interfaces with a Combination of Bundle Types 1004

Chapter 52 Summary of Multilink and Link Services Configuration Statements 1011

acknowledge-retries 1011
acknowledge-timer 1012
action-red-differential-delay 1012
address 1013
bundle 1013
destination 1014

disable-mlppp-inner-ppp-pfc 1014
dlci 1015
drop-timeout 1016
encapsulation 1017
encapsulation (Logical Interface) 1017
encapsulation (Physical Interface) 1018
family 1019
fragment-threshold 1020
hello-timer 1020
interfaces 1021
interleave-fragments 1021
lmi-type 1022
minimum-links 1022
mlfr-uni-nni-bundle-options 1023
mrru 1024
mtu 1025
multicast-dlci 1025
n391 1026
n392 1026
n393 1027
red-differential-delay 1027
short-sequence 1028
t391 1028
t392 1029
unit 1030
yellow-differential-delay 1031

Part 7 Real-Time Performance Monitoring Services

Chapter 53 Real-Time Performance Monitoring Services Overview 1035

Chapter 54 Real-Time Performance Monitoring Configuration Guidelines 1037

Configuring BGP Neighbor Discovery Through RPM 1038
Configuring Real-Time Performance Monitoring 1040
Configuring RPM Probes 1040
Configuring RPM Receiver Servers 1044
Limiting the Number of Concurrent RPM Probes 1045
Configuring RPM Timestamping 1045
Configuring TWAMP 1048
Configuring TWAMP Interfaces 1049
Configuring TWAMP Servers 1049
Examples: Configuring BGP Neighbor Discovery Through RPM 1050
Examples: Configuring Real-Time Performance Monitoring 1051

Chapter 55 Summary of Real-Time Performance Monitoring Configuration Statements 1055

authentication-mode 1055
client-list 1056
data-fill 1056
data-size 1057
destination-interface 1058
destination-port 1058
dscp-code-point 1059
hardware-timestamp 1060
history-size 1060
inactivity-timeout 1061
logical-system 1061
maximum-connections 1062
maximum-connections-per-client 1062
maximum-sessions 1063
maximum-sessions-per-connection 1063
moving-average-size 1064
one-way-hardware-timestamp 1064
port 1065
    port (RPM) 1065
    port (TWAMP) 1065
probe 1066
probe-count 1067
probe-interval 1067
probe-limit 1068
probe-server 1068
probe-type 1069
routing-instance 1069
routing-instances 1070
rpm 1070
server 1071
services 1071
source-address 1072
target 1072
tcp 1073
test 1074
test-interval 1075
thresholds 1076
traps 1077
twamp 1078
twamp-server 1078
udp 1079

Part 8 Tunnel Services

Chapter 56 Tunnel Services Overview 1083

Chapter 57 Tunnel Interfaces Configuration Guidelines 1087

Configuring Unicast Tunnels 1087
Configuring a Key Number on GRE Tunnels 1089
Enabling Fragmentation on GRE Tunnels 1090
Specifying an MTU Setting for the Tunnel 1090
Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header 1091
Configuring Packet Reassembly 1091
Restricting Tunnels to Multicast Traffic 1092
Configuring Logical Tunnel Interfaces 1092
Connecting Logical Systems 1093
Configuring Tunnel Interfaces for Routing Table Lookup 1094
Configuring Virtual Loopback Tunnels for VRF Table Lookup 1094
Configuring PIM Tunnels 1096
Configuring IPv6-over-IPv4 Tunnels 1096
Configuring Dynamic Tunnels 1097
Configuring Tunnel Interfaces on MX Series Routers 1097
Examples: Configuring Unicast Tunnels 1098
Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup 1099
Example: Configuring an IPv6-over-IPv4 Tunnel 1100
Example: Configuring Logical Tunnels 1100

Chapter 58 Summary of Tunnel Services Configuration Statements 1103

allow-fragmentation 1103
backup-destination 1104
copy-tos-to-outer-ip-header 1104
destination 1105
destination (Tunnel Remote End) 1105
destination (Routing Instance) 1105
destination-networks 1106
do-not-fragment 1106
dynamic-tunnels 1107
interfaces 1107
key 1108
multicast-only 1108
peer-unit 1109
reassemble-packets 1109
routing-instance 1110
routing-instances 1110
routing-options 1111
source 1111
source-address 1112
ttl 1112
tunnel 1113

tunnel-type 1114
unit 1115

Part 9 Index

Index 1119
Index of Statements and Commands 1139

List of Figures

Part 2 Adaptive Services

Chapter 3 **Adaptive Services Overview** 33
Figure 1: Packet Flow Through the Adaptive Services or Multiservices PIC 41
Chapter 8 **Network Address Translation Services Configuration Guidelines** 129
Figure 2: Configuring NAT for Multicast Traffic 145
Chapter 12 **IPsec Services Configuration Guidelines** 213
Figure 3: IPsec Dynamic Endpoint Tunneling Topology 253

Part 4 Encryption Services

Chapter 40 **Encryption Interfaces Configuration Guidelines** 763
Figure 4: Example: IPsec Tunnel Connecting Security Gateways 765
Figure 5: IPsec Tunnel Redundancy 771

Part 5 Flow Monitoring and Discard Accounting Services

Chapter 42 **Flow Monitoring and Discard Accounting Overview** 783
Figure 6: Passive Monitoring Application Topology 784
Figure 7: Active Monitoring Configuration Topology 786
Chapter 43 **Flow Monitoring and Discard Accounting Configuration Guidelines** 787
Figure 8: Configure Sampling Rate 793
Chapter 45 **Flow Collection Configuration Guidelines** 899
Figure 9: Flow Collector Interface Topology Diagram 905
Chapter 47 **Dynamic Flow Capture Configuration Guidelines** 929
Figure 10: Dynamic Flow Capture Topology 930
Chapter 48 **Flow-Tap Configuration Guidelines** 941
Figure 11: Flow-Tap Topology 943

Part 6 Link and Multilink Services

Chapter 51 **Link and Multilink Services Configuration Guidelines** 973
Figure 12: Multilink Interface Configuration 976

List of Tables

About This Guide **xlix**

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

Part 2 Adaptive Services

Chapter 3 Adaptive Services Overview **33**

Table 3: AS and Multiservices PIC Services by Service Package, PIC, and Platform 36
Table 4: Statement Equivalents for ES and AS Interfaces 47

Chapter 4 Applications Configuration Guidelines **59**

Table 5: Application Protocols Supported by Services Interfaces 61
Table 6: Network Protocols Supported by Services Interfaces 63
Table 7: ICMP Codes and Types Supported by Services Interfaces 64
Table 8: Port Names Supported by Services Interfaces 65
Table 9: Supported RPC Services 74

Chapter 6 Stateful Firewall Services Configuration Guidelines **107**

Table 10: IP Option Values 111

Chapter 12 IPsec Services Configuration Guidelines **213**

Table 11: Default IKE and IPsec Proposals for Dynamic Negotiations 247

Chapter 14 Layer 2 Tunneling Protocol Services Configuration Guidelines **291**

Table 12: System Log Message Severity Levels 299

Chapter 22 Service Set Configuration Guidelines **439**

Table 13: System Log Message Severity Levels 451
Table 14: Adaptive Services Tracing Flags 454

Chapter 24 Service Interface Configuration Guidelines **479**

Table 15: System Log Message Severity Levels 483

Part 6 Link and Multilink Services

Chapter 51 Link and Multilink Services Configuration Guidelines **973**

Table 16: Multilink Services PIC Capacities 975
Table 17: Multilink and Link Services Logical Interface Statements 978
Table 18: Link Services Physical Interface Statements for MLFR FRF.16 988
Table 19: Link Services CoS Queues 992
Table 20: Link Services Bundle 1001

Part 8 Tunnel Services

Chapter 56 Tunnel Services Overview 1083
Table 21: Tunnel Interface Types 1083
Chapter 57 Tunnel Interfaces Configuration Guidelines 1087
Table 22: Methods for Configuring Egress Filtering 1095

About This Guide

This preface provides the following guidelines for using the *Juniper Networks JUNOS® Software Services Interfaces Configuration Guide*:

- JUNOS Documentation and Release Notes on page xlix
- Objectives on page l
- Audience on page l
- Supported Platforms on page l
- Using the Indexes on page li
- Using the Examples in This Manual on page li
- Documentation Conventions on page lii
- Documentation Feedback on page liv
- Requesting Technical Support on page liv

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 Software 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 JUNOS Software 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 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 .

Objectives

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

NOTE: For additional information about JUNOS Software—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, JUNOS Software currently supports the following platforms:

- J Series
- M Series

- MX 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. 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;
        }
    }
}
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 *JUNOS CLI User Guide*.

Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page liii 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 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 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*
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; 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***>;**

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, 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/customers/support/downloads/710059.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/
- Search for known bugs: http://www2.juniper.net/kb/
- 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

Part 1
Overview

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

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 Network Interfaces Configuration Guide*.
- 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 Software 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 33.

NOTE: On Juniper Networks MX Series Ethernet Services 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, and support for digital certificates. For more information about encryption interfaces, see "Encryption Interfaces Configuration Guidelines" on page 763.

- 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 and Discard Accounting Configuration Guidelines" on page 787.

- 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 Software 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 Services Configuration Guidelines" on page 973.
- 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 Interfaces Configuration Guidelines" on page 1087.

Supported Platforms

For information about which platforms support Adaptive Services and Multiservices PICs and their features, see Table 3 on page 36.

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 J Series Services Routers, see the J Series documentation, in particular the *JUNOS Software Interfaces and Routing Configuration Guide*.

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 Software, 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 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 11
- [edit services] Hierarchy Level on page 11

[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;
}
```

```
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 *JUNOS Policy Framework 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;
```

```
            input-interface-index number;
            output-interface-index number;
            source-address address;
        }
    }
  }
next-hop-group (Forwarding Options) group-name {
  interface interface-name {
    next-hop address;
  }
}
port-mirroring {
  input {
    rate rate;
    run-length number;
  }
  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;
  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;
        }
        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 {
    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 Network Interfaces Configuration Guide*. 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;
            }
        }
    }
```

```
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;
            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 client;
    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);
}
services-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    syslog {
        host hostname {
            services severity-level;
            facility-override facility-name;
            log-prefix prefix-value;
        }
    }
```

```
    }
}
rlsq number {
    redundancy-options {
        hot-standby | warm-standby;
        primary lsq-fpc/pic/port;
        secondary lsq-fpc/pic/port;
    }
}
rsp number {
    redundancy-options {
        primary sp-fpc/pic/port;
        secondary sp-fpc/pic/port;
    }
}
so-fpc/pic/port {
    unit logical-unit-number {
        passive-monitor-mode;
    }
}
```

[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 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 Hierarchy and RFC Reference*. This listing includes only the statements documented in this manual; additional statements are documented in the *JUNOS Subscriber Access Configuration Guide*.

```
aacl {
    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;
            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 {
  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;
    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;
  max-checked-bytes bytes;
  min-checked-bytes bytes;
```

```
    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 {
      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;
    }
    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-policies [ policy-and-policy-set-names ];
    new-transaction-policies [ policy-and-policy-set-names ];
  }
  transport-details <port port-number> <ip-address ip-address> <tcp> <udp>;
}
sip {
  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 {
            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 [ contact-fields ];
        method {
          method-invite;
          method-message;
          method-options;
          method-publish;
```

```
                method-refer;
                method-register;
                method-subscribe;
            }
            request-uri [ uri-fields ];
            source-address [ ip-addresses ];
        }
        then {
            (accept | reject);
            admission-control admission-control-profile;
            route {
                egress-service-point service-point-name;
                next-hop (request-uri | address ipv4-address <port port-number>
                    <transport-protocol (udp | tcp)>);
            }
            trace;
        }
    }
  }
  new-transaction-policy-set policy-set-name {
    policy-name [ policy-names ];
  }
  timers {
    inactive-call seconds;
    inactive-call seconds;
    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;
```

```
                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 {
        sip-text {
            dscp (alias | bits);
            forwarding-class class-name;
        }
        sip-video {
            dscp (alias | bits);
            forwarding-class class-name;
        }
        sip-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 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;
            (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;
}
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-tapflow-tap {
    interface interface-name;
}
flow-tap-lite {
    interface interface-name;
}
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;
    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);
        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);
        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;
            }
            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;
  }
  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 {
            nat-type(full-cone | symmetric);
            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;
                    overload-pool overload-pool-name;
                    overload-prefix overload-prefix;
                    source-pool nat-pool-name;
                    source-prefix source-prefix;
                    translation-type (destination type | source type);
```

```
            }
          }
        }
      }
      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 {
              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;
    }
}
traffic-management {
    max-burst-size {
        default bytes-per-scond;
        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;
            }
         }
      }
   }
   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);
               [Unresolved xref] (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);
```

```
                [Unresolved xref] (forced-904 | forced-906 | none);
                graceful (graceful-905 | none);
                link-loss (forced-906 | 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 source-port;
        }
        disable-session-mirroring;
    }
}
media-service media-service-name {
    nat-pool nat-pool-name;
}
rule rule-name {
    gateway gateway-name;
    media-service media-service-name;
}
rule-set rule-set-name {
    rule rule-name1;
    rule rule-name2;
    rule rule-name3;
}
traceoptions {
    file <filename filename> <files number> <match regexg> <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 {
    media-service media-service-name;
    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 source-port;
    }
    disable-session-mirroring;
}
}
rpm {
    bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name <routing-instances routing-instance-name>;
        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 | address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
        }
    }
    probe-limit limit;
    probe-server {
        tcp {
            destination-interface interface-name;
            port (RPM) number;
        }
        udp {
            destination-interface interface-name;
            port (RPM) 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);
    (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;
    }
    service-order {
        forward-flow [ service-name1 service-name2 ];
        reverse-flow [ service-name1 service-name2 ];
    }
    syslog {
        host hostname {
            services severity-level;
            facility-override facility-name;
            log-prefix prefix-value;
        }
    }
}
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;
    }
  }
}
```

Part 2

Adaptive Services

- Adaptive Services Overview on page 33
- Applications Configuration Guidelines on page 59
- Summary of Applications Configuration Statements on page 97
- Stateful Firewall Services Configuration Guidelines on page 107
- Summary of Stateful Firewall Configuration Statements on page 117
- Network Address Translation Services Configuration Guidelines on page 129
- Summary of Network Address Translation Configuration Statements on page 155
- Intrusion Detection Service Configuration Guidelines on page 179
- Summary of Intrusion Detection Service Configuration Statements on page 191
- IPsec Services Configuration Guidelines on page 213
- Summary of IPsec Services Configuration Statements on page 257
- Layer 2 Tunneling Protocol Services Configuration Guidelines on page 291
- Summary of Layer 2 Tunneling Protocol Configuration Statements on page 309
- Link Services IQ Interfaces Configuration Guidelines on page 323
- Summary of Link Services IQ Configuration Statements on page 383
- Voice Services Configuration Guidelines on page 395
- Summary of Voice Services Configuration Statements on page 405
- Class-of-Service Configuration Guidelines on page 415
- Summary of Class-of-Service Configuration Statements on page 425
- Service Set Configuration Guidelines on page 439
- Summary of Service Set Configuration Statements on page 457
- Service Interface Configuration Guidelines on page 479
- Summary of Service Interface Configuration Statements on page 493
- PGCP Configuration Guidelines for the BGF Feature on page 511
- Summary of PGCP Configuration Statements on page 517
- Service Interface Pools Configuration Guidelines on page 613
- Summary of Service Interface Pools Statements on page 615
- Border Signaling Gateway Configuration Guidelines on page 617
- Summary of Border Signaling Gateway Configuration Statements on page 621

Chapter 3

Adaptive Services Overview

The Adaptive Services (AS) and Multiservices PICs provide *adaptive services interfaces*, which allow you to coordinate multiple services on a single PIC by configuring a set of services and applications. The AS and Multiservices PICs offers a special range of services you configure in one or more service sets.

The AS PIC is available in two versions that differ in memory size:

- The Adaptive Services II PIC with 512 MB of memory is supported on all Juniper Networks M Series and T Series routers, including the M320 router.
- The Adaptive Services PIC with 256 megabytes (MB) of memory is supported on all M Series routers except the M320 router.

The M7i router includes the Adaptive Services Module (ASM), an integrated version of the AS PIC as an optional component, which offers all the features of the standalone version at a reduced bandwidth.

NOTE: To take advantage of the features available on the AS PIC, you must install it in an Enhanced Flexible PIC Concentrator (FPC) in an M Series router equipped with an Internet Processor II application-specific integrated circuit (ASIC), or a similarly equipped T Series router. To find out whether your router hardware is suitably equipped, use the show chassis hardware command. For more information, see the *JUNOS System Basics and Services Command Reference*.

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 fragmentation); it also supports graceful Routing Engine switchover (GRES) and Dynamic Applicaton Awareness for JUNOS Software. For more information about supported packages, see “Enabling Service Packages” on page 35.

NOTE: The Adaptive Services and 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 AS and Multiservices PICs, 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 Software 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).
- 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 Software 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.

This chapter includes the following topics:

- Enabling Service Packages on page 35
- Services Configuration Procedure on page 39
- Packet Flow Through the Adaptive Services or Multiservices PIC on page 40
- Stateful Firewall Overview on page 41
- Network Address Translation Overview on page 44
- IPsec Overview on page 45
- Layer 2 Tunneling Protocol Overview on page 48
- Voice Services Overview on page 48
- Class of Service Overview on page 49
- Examples: Services Interfaces Configuration on page 49

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 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 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 36 lists the services supported within each service package for each PIC and platform. Some of these services are also available on J Series routers, but they are implemented in a different way and do not use the Layer 2 and Layer 3 service packages. For more information, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

On the AS and Multiservices PICs, *link services* support includes JUNOS Software 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 38 and “Link Services IQ Interfaces Configuration Guidelines” on page 323.

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 Feature Guide*.

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:					

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
■ Link services	Yes	Yes	Yes	Yes	No
■ Multiclass MLPPP	Yes	Yes	Yes	Yes	No
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:					

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—"Link Services IQ Interfaces Configuration Guidelines" on page 323 describes how the JUNOS CoS components work on link services IQ (lsq) interfaces. For detailed information about JUNOS CoS components, see the *JUNOS 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.

- 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)

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 36.

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 "Link Services IQ Interfaces Configuration Guidelines" on page 323 and "Link and Multilink Services Configuration Guidelines" on page 973.

NOTE: Interface type sp is created because it is needed by the JUNOS Software. 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.

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 41. (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 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

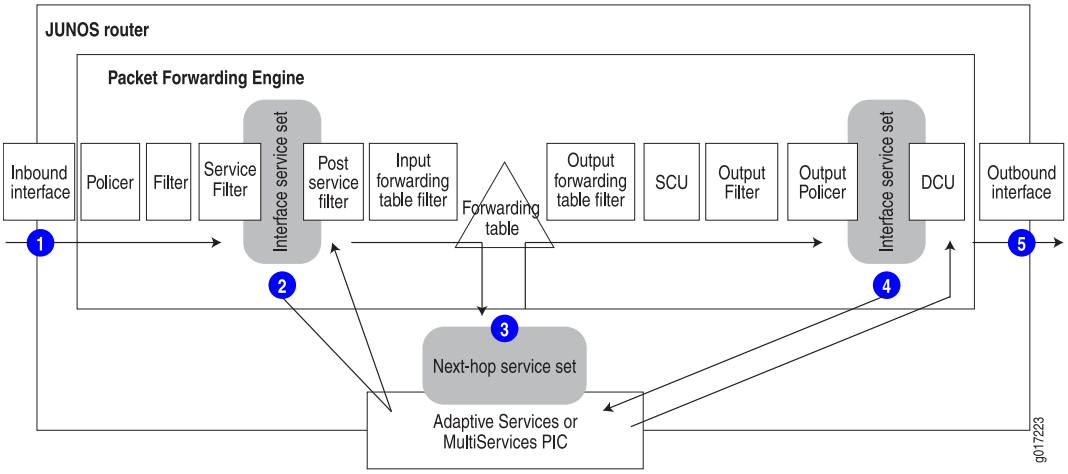

NOTE: When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG_ERR is generated. This mechanism applies to both Layer 2 and Layer 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, 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 about configuring stateful firewalls, see "Stateful Firewall Services Configuration Guidelines" on page 107.

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.

  - 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

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

Network Address Translation Overview

NAT is a mechanism for concealing a set of host addresses on a private network behind a pool of public addresses. It can be used as a security measure to protect the host addresses from direct targeting in network attacks. The Juniper Networks JUNOS Software supports NAT on IPv4 and IPv6 networks, and also on traffic transiting between the two.

You can configure NAT using traditional NAT or twice NAT, as described in the following sections:

- Traditional NAT on page 44
- Twice NAT on page 45

Traditional NAT

Traditional NAT, specified in RFC 3022, *Traditional IP Network Address Translator*, is fully supported by the JUNOS Software. In addition, network address port translation (NAPT) is supported for source addresses.

The AS and Multiservices PIC interfaces support three types of NAT processing:

- Static-source NAT hides a private network without using NAPT.
- Dynamic-source NAT hides a private network using NAPT.
- Static-destination NAT makes selected private servers accessible.

You can implement NAT to hide one or many hosts on a private network behind a pool of public IP addresses. The pool can be as small as one IP address, or it can be a set of contiguous IP addresses. You can specify a port range to restrict port translation when NAT is configured in dynamic-source mode.

Private address to public address binding can be either static or dynamic. In the basic NAT mode, a NAT rule can force a private IP address to be always bound to a public address; in the NAPT mode, a NAT rule can force a paired private address and private TCP or UDP port to be mapped to a public IP and public TCP or UDP port. However, when the address binding is not statically forced by the NAT rules, NAT can dynamically pick an available address or address and TCP or UDP port pairing when a new session starts. You can specify multiple prefixes and address ranges in a dynamic or static source NAT pool.

The option to assign NAT addresses statically from a dynamic NAT pool enables you to advertise one subnet that represents the NAT pool and use an address within that subnet for static rules. Statically assigned addresses are not reused for dynamic assignment and can only be used for static-source NAT (not for static-destination NAT).

You can configure an overload (fallback) pool to be used when the source pool of addresses is exhausted. The overload pool must be configured with NAPT.

You can also configure NAT rules without configuring a pool by directly specifying the address prefix to be translated within the rule. And, within the rule, you can assign particular addresses that you do not want to be translated.

Like most traditional NAT implementations, the JUNOS implementation of NAT supports sessions initiated from the private side only. Sessions initiated from the public side are supported only when you configure static address binding.

You are not required to configure a stateful firewall rule to allow NAT traffic. By default, NAT traffic is allowed unless it is explicitly configured to be dropped. If only NAT is configured in a service set, all traffic is accepted.

For more information about configuring NAT rules, see "Network Address Translation Services Configuration Guidelines" on page 129.

Twice NAT

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

In twice NAT, both the source and destination addresses are subject to translation as packets traverse the NAT router. 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 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 other ALGs. By default, the twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages.

For more information about configuring NAT rules, see "Network Address Translation Services Configuration Guidelines" on page 129.

IPsec Overview

The Juniper Networks JUNOS Software supports IPsec. This section discusses the following topics, which provide background information about configuring IPsec:

- IPsec on page 46
- Security Associations on page 46
- IKE on page 46
- Comparison of IPsec Services and ES Interface Configuration on page 47

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

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 Software 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.

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).

IKE consists of two phases. In the first phase, it negotiates security attributes and establishes shared secrets to form the bidirectional IKE SA. In the second phase, inbound and outbound IPsec SAs are established and the IKE SA secures the exchanges. IKE also generates keying material, provides Perfect Forward Secrecy, and exchanges identities.

Comparison of IPsec Services and ES Interface Configuration

Table 4 on page 47 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 {...}
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 Services Configuration Guidelines" on page 213. For more information about configuring encryption services on an ES PIC, see "Encryption Interfaces Configuration Guidelines" on page 763.

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 about configuring L2TP services, see “Layer 2 Tunneling Protocol Services Configuration Guidelines” on page 291.

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.

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 “Voice Services Configuration Guidelines” on page 395.

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 351.

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 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 Configuration Guidelines” on page 415.

Examples: Services 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;
          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 {
        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;
    }
}
```

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

```
[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;
      }
      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;
      }
   }
}
```

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;
      }
   }
}
[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;
      }
    }
```

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 {
    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;
   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 source dynamic;
            }
        }
    }
}
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 {
    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;
        }
    }
```

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

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

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 60
- Configuring Application Sets on page 70
- ALG Descriptions on page 70
- Verifying the Output of ALG Sessions on page 78
- JUNOS Default Groups on page 85
- Examples: Configuring Application Protocols on page 93

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;
    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;
}
```

You can group application objects by configuring the application-set statement; for more information, see “Configuring Application Sets” on page 70.

This section includes the following tasks for configuring applications:

- Configuring an Application Protocol on page 60
- Configuring the Network Protocol on page 62
- Configuring the ICMP Code and Type on page 63
- Configuring Source and Destination Ports on page 65
- Configuring the Inactivity Timeout Period on page 68
- Configuring SIP on page 68
- Configuring an SNMP Command for Packet Matching on page 69
- Configuring an RPC Program Number on page 69
- Configuring the TTL Threshold on page 69
- Configuring a Universal Unique Identifier on page 70

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 61 shows the list of supported protocols. For more information about specific protocols, see "ALG Descriptions" on page 70.

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	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
Internet Control Message Protocol (ICMP)	icmp	Requires the protocol statement to have the value icmp or to be unspecified.
Internet Inter-ORB Protocol (IIOP) Transmission Control Protocol (TCP)	iiop	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
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.
RealAudio	realaudio	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
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.
Session Initiation Protocol (SIP)	sip	For more information, see "Configuring SIP" on page 68.

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.
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.
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	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.

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 Services Configuration Guidelines" on page 129.

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 63 shows the list of the supported protocols.

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 Services Configuration Guidelines" on page 129.

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 64 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 *JUNOS Policy Framework 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 *JUNOS Policy Framework 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).

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 61.

You can specify either a numeric value or one of the text synonyms listed in Table 8 on page 65.

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
domain	53

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

Port Name	Corresponding Port Number
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
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518

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

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

For more information about matching criteria, see the *JUNOS Policy Framework 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 482.

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 60. 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.

  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 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).

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 60.

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 60.

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 60.

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]
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 60. 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 93.

ALG Descriptions

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

- Basic TCP ALG on page 71
- Basic UDP ALG on page 71
- BOOTP on page 72
- DCE RPC Services on page 72
- FTP on page 72
- H323 on page 73
- ICMP on page 73
- IIOP on page 74
- NetShow on page 74
- RealAudio on page 74
- RPC and RPC Portmap Services on page 74

- RTSP on page 76
- SMB on page 76
- SNMP on page 76
- SQLNet on page 77
- TFTP on page 77
- Traceroute on page 77
- UNIX Remote-Shell Services on page 77
- WinFrame on page 78

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 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.

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 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 Bootstrap Protocol 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. 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

DCE 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.

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 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 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-video conferencing and collaboration applications. H323 consists of h.225 call signaling protocols and h.245, the control protocol for media communication. During h.225 negotiation, 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.

The JUNOS stateful firewall service monitors the h.225 control channel to open the h.245 control channel. Once the h.245 channel is created, the stateful firewall service also monitors this channel for media channel information and allows the media traffic through the firewall. The H323 ALG supports static destination, static source, 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 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 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

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

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

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.

RealAudio

The Real Networks PNA protocol RealAudio is not a separate service. RealAudio was the original protocol used by RealPlayer. Newer versions of RealPlayer use RTSP.

Support for stateful firewall and NAT services requires that you configure the RealAudio ALG on TCP port 7070. The stateful firewall monitors the traffic on the TCP control channel and dynamically opens the ports for data channels.

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 9 on page 74:

Table 9: 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).

Table 9: Supported RPC Services *(continued)*

Name	Description	Comments
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.
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 may use RTP, but it is not required. Media may 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.

SMB

Server message block (SMB) is a popular PC protocol that allows sharing of files, disks, directories, printers, and in some cases, COM ports across a network. SMB is a client/server, request-response-based protocol. Though there are some exceptions to this, most of the communication takes place using the request reply paradigm. Servers make file systems and resources available to clients on the network. Clients can send commands (smbs) to the server that allow them to access these shared resources. SMB can run over multiple protocols, including TCP/IP, NetBEUI, and IPX/SPX. In almost all cases, the NetBIOS interface is used. Microsoft is trying to rename SMB-based networking to Windows Networking and the protocol to CIFS. The SMB protocol is undocumented, although there is a public CIFS group. For more information, refer to the following link on CIFS: ftp://ftp.microsoft.com/developr/drg/CIFS/.

The SMB name service uses well-known UDP and TCP port 137, without requiring a special ALG. For NetBIOS data tunneled through UDP port 138 or TCP port 139, you must configure the NetBIOS ALG. Support for stateful firewall and NAT services requires that you configure the NetBIOS ALG on UDP port 138 and TCP port 139. For SMB name services, both TCP and UDP port 137 must be opened, without a special ALG.

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 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 TTL field to trigger ICMP time-exceeded messages from routers or 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 need to 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 the Citrix Windows application. Support of stateful firewall and NAT services requires that you configure the WinFrame ALG for TCP destination port 1494 and UDP port 1604.

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 78
- RTSP ALG Example on page 81
- System Log Messages on page 83

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 78
- FTP System Log Messages on page 79
- Analysis on page 80
- Troubleshooting Questions on page 80

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: sp-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.
- 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 83.

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-3/3/0:1.1.1.2:4450 -> 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
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
```

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 RealAudio 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 81
- Analysis on page 81
- Troubleshooting Questions on page 82

Sample Output

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

```
user@host# show services stateful-firewall conversations
Interface: sp-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
TCP            2.2.2.2:554   ->        1.1.1.3:58795 Watch    O             5
UDP            2.2.2.2:1028  ->        1.1.1.3:1028  Forward  O             6
UDP            2.2.2.2:1029  ->        1.1.1.3:1029  Forward  O             0
UDP            2.2.2.2:1030  ->        1.1.1.3:1030  Forward  O             3
UDP            2.2.2.2:1031  ->        1.1.1.3:1031  Forward  O             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: sp-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
      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
      H323: 0, ICMP: 0, IIOP: 0
      Login: 0, NetBIOS: 0, NetShow: 0
      Real Audio: 0, RPC: 0, RPC portmap: 0
      RTSP: 0, Shell: 0, SIP: 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 84
- System Log Output on page 85

System Log Configuration

You can configure the enabling of system log messages at a number of different levels in the JUNOS 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 System Basics Configuration Guide* (system level) or the *JUNOS Services Interfaces Configuration Guide* (all other levels).

1. At the topmost global level:

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

2. At the AS PIC level:

```
user@host# show interfaces sp-3/2/0
services-options {
   syslog {
      host local {
         services any;
      }
   }
   open-timeout 3600;
   inactivity-timeout 3600;
}
unit 0 {
   family inet;
}
```

3. 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 sp-3/2/0;
}
```

4. 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
```

The following system log message indicates that the AS PIC created a valid flow:

```
Oct 25 16:11:37 (FPC Slot 3, PIC Slot 2) {svc_set}[FWNAT]:
   ASP_SFW_CREATE_ACCEPT_FLOW: proto 6 (TCP) application: rtsp,
   ge-2/0/1.0:1.1.1.2:35595 -> 2.2.2.2:554, creating forward or watch flow
```

The following system log message indicates that the packet did not match any rule on the AS PIC:

```
Oct 25 16:03:13 (FPC Slot 3, PIC Slot 2) {svc_set}[FWNAT]: ASP_SFW_NO_RULE_DROP:
   proto 17 (UDP), ge-2/0/3.0:2.2.2.2:1334 -> 1.1.1.78:1334, No matching SFW rule
```

If the AS PIC does not create data flows and starts receiving data, this message is logged. It represents an error condition for ALG processing.

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

JUNOS Default Groups

The JUNOS Software 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-.

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;
}
# 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
application junos-h323 {
    application-protocol h323;
    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.
application junos-iiop-java {
    application-protocol iiop;
    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.
application junos-iiop-orbix {
    application-protocol iiop;
    protocol tcp;
    destination-port 3075;
}
# RealPlayer uses RealAudio for real-time streaming.
# This was the original RealPlayer protocol.
# RTSP is more widely used by RealPlayer,
# but it still supports RealAudio.
application junos-realaudio {
    application-protocol realaudio;
    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.
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;
    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;
}
application junos-https {
    protocol tcp;
    destination-port 443;
}
# SIP control session for VoIP
application junos-sip {
    application-protocol sip;
    protocol udp;
    destination-port 5060;
    learn-sip-register;
}
# " 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-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-dce-rpc-portmap;
```

```
        application junos-dcerpc-msexchange-directory-rfr;
        application junos-dcerpc-msexchange-information-store;
        application junos-dcerpc-msexchange-directory-nsp;
        application junos-sip;
    }
    # " 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;
    }
}
}
}
```

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 an Application Protocol" on page 60; for details about a specific protocol, see "ALG Descriptions" on page 70.

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;
            }
            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;
}
```

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.

The following example shows a complete configuration for SIP and its related statements:

```
[edit]
applications {
   application sip {
      application-protocol sip;
      protocol udp;
      destination-port 5060;
      inactivity-timeout 300;
      learn-sip-register;
   }
```

```
}
interfaces {
    sp-0/2/0 {
        services-options {
            syslog {
                host local {
                    services any;
                }
            }
        }
        unit 0 {
            family inet;
        }
    }
    ge-1/0/0 {
        description amazon_eth2;
        unit 0 {
            family inet {
                service {
                    input {
                        service-set test_sip;
                    }
                    output {
                        service-set test_sip;
                    }
                }
                address 10.200.1.1/30;
            }
        }
    }
    ge-1/1/0 {
        description maxtor_eth1;
        unit 0 {
            family inet {
                address 10.100.1.1/30;
            }
        }
    }
}
services {
    stateful-firewall {
        rule sip {
            match-direction input-output;
            term 0 {
                from {
                    applications sip;
                }
                then {
                    accept;
                }
            }
        }
    }
    ids {
        rule ids {
            match-direction input-output;
```

```
        term 0 {
            then {
                force-entry;
                logging {
                    threshold 1;
                    syslog;
                }
            }
        }
    }
  }
  service-set test_sip {
    syslog {
      host local {
        services any;
      }
    }
    stateful-firewall-rules sip;
    ids-rules ids;
    interface-service {
      service-interface sp-0/2/0;
    }
  }
}
```

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;
    learn-sip-register;
    protocol type;
    rpc-program-number number;
    sip-call-hold-timeout seconds;
    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 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 60.

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

application-protocol

Syntax application-protocol *protocol-name*;

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

Release Information Statement introduced before JUNOS Release 7.4.
login and sip options introduced in JUNOS Release 7.4.
ip option introduced in JUNOS 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
- dce-rpc
- dce-rpc-portmap
- dns
- exec
- ftp
- h323
- icmp
- iiop
- ip
- login
- netbios
- netshow
- realaudio
- rpc
- rpc-portmap
- rtsp
- shell
- sip
- snmp
- sqlnet
- tftp
- traceroute
- winframe

Usage Guidelines See "Configuring an Application Protocol" on page 60.

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 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 70.

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 Release 7.4.

Description Define the applications used in services.

Usage Guidelines See "Applications Configuration Guidelines" on page 59.

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

destination-port

Syntax destination-port *port-value*;

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

Release Information Statement introduced before JUNOS 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 65.

Usage Guidelines See "Configuring Source and Destination Ports" on page 65.

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 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 63.

Usage Guidelines See "Configuring the ICMP Code and Type" on page 63.

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

icmp-type

Syntax icmp-type *value*;

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

Release Information Statement introduced before JUNOS 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 63.

Usage Guidelines See "Configuring the ICMP Code and Type" on page 63.

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 Release 7.4.

Description Inactivity timeout period, in seconds.

Options *seconds*—Length of time the application is inactive before it times out.
Default: 60 seconds

Usage Guidelines See "Configuring the Inactivity Timeout Period" on page 68.

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

learn-sip-register

Syntax learn-sip-register;

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

Release Information Statement introduced in JUNOS Release 7.4.

Description Activate SIP register to accept potential incoming SIP calls.

Usage Guidelines See “Configuring SIP” on page 68.

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

protocol

Syntax protocol *type*;

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

Release Information Statement introduced before JUNOS 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 62.

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

rpc-program-number

Syntax rpc-program-number *number*;

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

Release Information Statement introduced before JUNOS 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 69.

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 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 68.

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

snmp-command

Syntax snmp-command *command*;

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

Release Information Statement introduced before JUNOS 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 69.

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 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 65.

Usage Guidelines See "Configuring Source and Destination Ports" on page 65.

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

ttl-threshold

Syntax ttl-threshold *number*;

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

Release Information Statement introduced before JUNOS 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 69.

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 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 70.

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

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 108
- Configuring Stateful Firewall Rule Sets on page 112
- Examples: Configuring Stateful Firewall Rules on page 112

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.
- 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 stateful firewall rules:

- Configuring Match Direction for Stateful Firewall Rules on page 108
- Configuring Match Conditions in Stateful Firewall Rules on page 109
- Configuring Actions in Stateful Firewall Rules on page 110

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);
```

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 440.

On the 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 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 *JUNOS Policy Framework 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 112.

If you omit the from term, the stateful firewall 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 protocol definitions you have configured at the [edit applications] hierarchy level; for more information, see "Applications Configuration Guidelines" on page 59.

- 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.

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 10 on page 111 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 10: 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	–

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;
        }
    }
  }
}
```

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;
      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 {
         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;
                }
                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 “Examples: Services Interfaces Configuration” on page 49.

Chapter 7

Summary of Stateful Firewall Configuration Statements

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

allow-ip-options

Syntax allow-ip-options [*value*s];

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 110.

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

application-sets

Syntax applications-sets *set-name*;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 109.

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

applications

Syntax applications [*application-name*s];

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 109.

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

destination-address

Syntax destination-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS Release 7.4.
any-unicast and except options introduced in JUNOS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS 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 109.

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 rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 7.6.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS 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 109.

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

destination-prefix-list

Syntax destination-prefix-list *list-name* <except>;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS 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 109.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

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 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 *JUNOS Policy Framework Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Stateful Firewall Rules" on page 108.

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 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 108.

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

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 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 108.

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

rule-set

Syntax
```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services stateful-firewall]

Release Information Statement introduced before JUNOS 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 112.

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

services

Syntax services stateful-firewall { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

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 Services Configuration Guidelines" on page 107.

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* | any-unicast) <except>;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS Release 7.4.
any-unicast and except options introduced in JUNOS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS 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 109.

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 Release 7.6.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS 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 109.

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

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 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 109.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

syslog

Syntax syslog;

Hierarchy Level [edit services stateful-firewall rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 110.

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

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 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 108.

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

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 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 110.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

Chapter 8

Network Address Translation Services Configuration Guidelines

To configure Network Address Translation (NAT) services, include the nat statement at the [edit services] hierarchy level:

```
[edit services]
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 [ transport-protocols ];
        }
        port (automatic | range low value high value) {
            random-allocation;
        }
    }
    rule rule-name {
        match-direction (input | output);
        term term-name {
            nat-type (full-cone | symmetric);
            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 {
                no-translation;
                translated {
                    destination-pool nat-pool-name;
                    destination-prefix destination-prefix;
```

```
                overload-pool overload-pool-name;
                overload-prefix overload-prefix;
                source-pool nat-pool-name;
                source-prefix source-prefix;
                translation-type (destination type | source type);
                translation-type {
                    source type;
                    destination type;
                }
            }
            syslog;
        }
    }
    rule-set rule-set-name {
        [ rule rule-names ];
    }
}
```

This chapter includes the following sections:

- Configuring Addresses and Ports for Use in NAT Rules on page 130
- Configuring NAT Rules on page 133
- Configuring NAT Rule Sets on page 139
- Examples: Configuring NAT Rules on page 139

Configuring Addresses and Ports for Use in NAT Rules

For information about configuring translated addresses, see the following sections:

- Configuring Pools of Addresses and Ports on page 130
- Specifying Destination and Source Prefixes when Pools Are Not Used on page 132
- Requirements for NAT Addresses on page 132
- Configuring IPv6 Multicast Filters on page 133

Configuring Pools of Addresses and Ports

You can use the **pool** statement to define the addresses (or prefixes), address ranges, and ports used for network address translation. You can also use the **pgcp** option with the **pool** statement to specify that NAT pool is used exclusively by the BGF. To configure the information, include the **pool** statement at the **[edit services nat]** hierarchy level:

```
[edit services nat]
pool nat-pool-name {
    address ip-prefix</prefix-length>;
    address-range low minimum-value high maximum-value;
    pgcp;
    port (automatic | range low minimum-value high maximum-value) {
        random-allocation;
    }
```

}

To configure pools for traditional NAT, specify either a destination pool or a source pool. To configure pools for twice NAT, specify both the destination pool and the source pool.

With static source NAT and dynamic source NAT, you can specify multiple IPv4 or IPv6 addresses (or prefixes) and IPv4 and IPv6 address ranges. Up to 10 prefixes or address ranges (or a combination) can be supported within a single pool.

With static destination NAT, you can also specify multiple address prefixes and address ranges in a single term. Multiple destination NAT terms can share a destination NAT pool. However, the netmask or range for the from address must be smaller or equal to the netmask or range for the destination pool address. If you define the pool to be larger than required, some addresses will not be used. For example, if you define the pool size as 100 addresses and the rule specifies only 80 addresses, the last 20 addresses in the pool are not used.

For constraints on specific translation types, see "Configuring Actions in NAT Rules" on page 137.

With source static NAT, the prefixes and address ranges cannot overlap between separate pools. However, source dynamic NAT (without NAPT) and destination static NAT allow more than one rule or service set to refer to the same pool, and allow multiple pools to have subnets that can overlap. A prefix pool can be used by multiple rules or terms.

NOTE: When you configure address pools for NAT and user access, these address pools can overlap with one another. To configure overlapping address pools, include the address or address-range statement at the [edit access address-pool *pool-name*] and [edit services nat pool *pool-name*] hierarchy level.

In an address range, the low value must be a lower number than the high value. When multiple address ranges and prefixes are configured, the prefixes are depleted first, followed by the address ranges.

When you specify a port for dynamic source NAT, address ranges are limited to a maximum of 32 addresses, for a total of approximately 2,000 flows. A dynamic NAT pool with no address port translation supports up to 65,535 addresses. There is no limit on the pool size for static source NAT.

The port statement specifies port assignment for the translated addresses. To configure automatic assignment of ports, include the port automatic statement at the [edit services nat pool *nat-pool-name*] hierarchy level. To configure a specific range of port numbers, include the port range low *minimum-value* high *maximum-value* statement at the [edit services nat pool *nat-pool-name*] hierarchy level. By default, the JUNOS Software allocates NAT ports sequentially. To configure random port allocation, include the random-allocation statement.

Specifying Destination and Source Prefixes when Pools Are Not Used

You can directly specify the destination or source prefix used in network address translation without configuring a pool. When you configure prefixes for twice NAT, you must specify both a destination prefix and a source prefix.

To configure the information, include the rule statement at the [edit services nat] hierarchy level:

```
[edit services nat]
rule rule-name {
    term term-name {
        then {
            translated {
                destination-prefix prefix;
                source-prefix prefix;
            }
        }
    }
}
```

Requirements for NAT Addresses

You must configure a specific address, a prefix, or the address-range boundaries:

- If you specify a specific address or prefix, it is assumed the translated address belongs to the inet.0 routing instance. The following addresses, while valid in inet.0, cannot be used for NAT translation:
  - 0.0.0.0/32
  - 127.0.0.0/8 (loopback)
  - 128.0.0.0/16 (martian)
  - 191.255.0.0/16 (martian)
  - 192.0.0.0/24 (martian)
  - 223.255.255.0/24 (martian)
  - 224.0.0.0/4 (multicast)
  - 240.0.0.0/4 (reserved)
  - 255.255.255.255 (broadcast)
- You can specify one or more IPv4 or IPv6 address prefixes in the pool statement and in the from clause of the NAT rule term. This enables you to configure source translation from a private subnet to a public subnet without defining a rule term for each address in the subnet. Destination translation cannot be configured by this method. For examples, see “Examples: Configuring NAT Rules” on page 139.
- When you configure static source NAT, the address prefix size you configure at the [edit services nat pool *pool-name*] hierarchy level must be larger than the

source-address prefix range configured at the [edit services nat rule *rule-name* term *term-name* from] hierarchy level. The source-address prefix range must also map to a single subnet or range of IPv4 or IPv6 addresses in the pool statement. Any pool addresses that are not used by the source-address prefix range are left unused; pools cannot be shared.

NOTE: When you include a NAT configuration that changes IP addresses, it might affect forwarding path features elsewhere in your router configuration, such as source class usage (SCU), destination class usage (DCU), filter-based forwarding, or other features that target specific IP addresses or prefixes.

NAT configuration might also affect routing protocols operation, because the protocol peering, neighbor, and interface addresses can be altered when routing protocols packets transit the Adaptive Services (AS) or Multiservices PIC.

Configuring IPv6 Multicast Filters

To enable multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery, you must include the ipv6-multicast-interfaces statement at the [edit services nat] hierarchy level.

```
[edit services nat]
ipv6-multicast-interfaces (all | interface-name) {
    disable;
}
```

By default, multicast filters are not enabled on media interfaces. To enable filters on all interfaces, include the ipv6-multicast-interfaces all statement. To enable filters on a specified interface only, include the ipv6-multicast-interfaces *interface-name* statement.

To disable filters that were previously enabled, include the disable statement.

Configuring NAT Rules

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

```
[edit services nat]
rule rule-name {
    match-direction (input | output);
    term term-name {
        nat-type (full-cone | symmetric);
        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 {
            no-translation;
            translated {
                destination-pool nat-pool-name;
                destination-prefix prefix;
                overload-pool overload-pool-name;
                overload-prefix overload-prefix;
                source-pool nat-pool-name;
                source-prefix prefix;
                translation-type (destination type | source type);
                translation-type {
                    source type;
                    destination type;
                }
            }
            syslog;
        }
    }
}
```

Each NAT 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.

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

- Configuring Match Direction for NAT Rules on page 134
- Configuring NAT Type for Terms in NAT Rules on page 135
- Configuring Match Conditions in NAT Rules on page 136
- Configuring Actions in NAT Rules on page 137

Configuring Match Direction for NAT Rules

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

```
[edit services nat 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 440.

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 NAT Type for Terms in NAT Rules

The NAT type specifies whether a particular term supports traditional NAT processing or full-cone NAT. A *full-cone* NAT is one in which all requests from the same internal IP address and port are mapped to the same external IP address and port. In addition, any external host can send a packet to the internal host by sending it to the mapped external address. Full-cone NAT is useful if you want to allow external hosts from the public network to connect to internal hosts using public IP addresses. However, we recommend that you use this feature along with strict firewall rules that allow only the intended traffic from the public network to reach the customer-edge router.

To configure the NAT type, include the nat-type statement at the [edit services nat rule *rule-name* term *term-name*] hierarchy level:

```
[edit services nat rule rule-name term term-name]
nat-type (full-cone | symmetric);
```

nat-type has two possible options:

- full-cone—Specifies that the term supports full-cone NAT.
- symmetric—Specifies that the term supports only traditional NAT; this is the default setting.

The following specifications and restrictions apply to full-cone NAT:

- As long as an internal host has a connection to an external host and uses source NAT, this feature allows any external host to connect back to the internal host over the public IP network.
- When the internal host terminates its connection to the external host, initiation of any new connections from external host to internal host over the public IP network is disallowed. Existing connections are not affected.
- Use of full-cone NAT enables the external-to-internal host connection to be independent from the internal-to-external host connection with regard to protocol and source and destination port.
- The aging mechanism for the external-to-internal host connection is similar to other host connections. Once the connection is established from the external host to the internal host, it is treated like any other network connection.

- Full-cone NAT is available with both source static and source dynamic NAT processing; for more information, see "Configuring Actions in NAT Rules" on page 137.
- It does not support Port Address Translation (PAT) or Network Address Port Translation (NAPT).
- It is not supported for use with twice NAT configurations.

For a configuration example, see "Example: Configuring Full-Cone NAT" on page 150.

Configuring Match Conditions in NAT Rules

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

```
[edit services nat 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>;
}
```

To configure traditional NAT and twice NAT, you can use the destination address, a range of destination addresses, the 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 *JUNOS Policy Framework 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 NAT rule. For an example, see "Examples: Configuring Stateful Firewall Rules" on page 112.

You can include application protocol definitions that you have configured at the [**edit applications**] hierarchy level; for more information, see "Applications Configuration Guidelines" on page 59:

- To apply one or more specific application protocol definitions, include the **applications** statement at the [**edit services nat 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 nat 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.

You can configure ALGs for ICMP and trace route under stateful firewall, NAT, or class of service (CoS) rules when twice NAT is configured in the same service set. Twice NAT does not support any other ALGs.

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.

Configuring Actions in NAT Rules

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

```
[edit services nat rule rule-name term term-name]
then {
    no-translation;
    syslog;
    translated {
        destination-pool nat-pool-name;
        destination-prefix destination-prefix;
        overload-pool overload-pool-name;
        overload-prefix overload-prefix;
        source-pool nat-pool-name;
        source-prefix source-prefix;
        translation-type (destination type | source type);
        translation-type {
            source type;
            destination type;
        }
    }
}
```

The no-translation statement allows you to specify addresses that you want to be excluded from NAT.

The destination-pool, destination-prefix, source-pool, and source-prefix statements specify addressing information that you define by including the pool statement at the [edit services nat] hierarchy level; for more information, see "Configuring Addresses and Ports for Use in NAT Rules" on page 130.

The overload-pool and overload-prefix statements specify a pool of addresses or an address prefix that can be used if the source pool becomes exhausted. If all the addresses in the source pool are in use, additional NAT sessions are supported using the overload pool. The overload pool must have NAPT configured.

For twice NAT, you can apply an overload pool for source addresses and combined source and destination addresses.

The syslog statement enables you to record an alert in the system logging facility.

The translation-type statement specifies what type of network address translation is used for source or destination traffic:

- destination static—Implement address translation for destination traffic without port mapping. This requires the size of the source address space to be the same or smaller than the size of the destination address space. You must specify a destination-pool name. The referenced pool can contain multiple addresses but no port configuration.
- source dynamic—Implement address translation for source traffic with NAPT. You must specify a source-pool name. The referenced pool must include a port or address configuration.

  If port automatic or port range is specified, port translation is used. If a port is not defined, the port value defaults to 1.

  The source dynamic option supports translating a large range of addresses to a smaller size pool. The requests from the source address range are assigned to the addresses in the pool until the pool is used up, and any additional requests are rejected. A NAT address assigned to a host is used for all concurrent sessions from that host. The address is released to the pool only after all the sessions for that host expire. This feature enables the router to share a few public IP addresses between several private hosts. Since all the private hosts might not simultaneously create sessions, they can share a few public IP addresses.
- source static—Implement address translation for source traffic without port mapping. The size of the pool address space must be greater than or equal to the source address space. You must specify a source-pool name. The referenced pool must contain exactly one address or prefix and no port configuration. You must include exactly one source-address value at the [edit services nat rule *rule-name* term *term-name* from] hierarchy level; if it is a prefix, the size must be less than or equal to the pool prefix size. Any addresses in the pool that are not matched in the source-address value remain unused, because a pool cannot be shared among multiple terms or rules.

For traditional NAT, you can configure either translation-type destination or translation-type source, but not both. To configure twice NAT, you specify both a translation-type destination and a translation-type source.

NOTE: You can statically assign NAT addresses from a dynamic NAT pool. This capability enables you to advertise one subnet that represents the NAT pool and use an address within that subnet for static rules. Statically assigned addresses are not reused for dynamic assignment. Statically assigned addresses from a dynamic pool can only be used for source static NAT and not for destination static NAT.

NOTE: When configuring NAT, if you specify the following addresses that do not match the NAT flow or NAT rule, the corresponding traffic is dropped:

- Addresses specified in the from destination-address statement, when you are using destination translation
- Addresses specified in the source NAT pool when you are using source translation

For more information on NAT methods, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

Configuring NAT Rule Sets

The rule-set statement defines a collection of NAT 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 nat] 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 NAT Rules

This section provides the following configuration examples:

- Example: Configuring Dynamic Source Translation on page 140
- Example: Configuring Static Source Translation on page 140
- Example: Configuring Dynamic and Static Source Translation on page 141
- Example: Configuring an Oversubscribed Pool with No Fallback on page 141
- Example: Configuring an Oversubscribed Pool with Fallback to NAPT on page 142
- Example: Configuring Static Source Translation with Multiple Prefixes and Address Ranges on page 142
- Example: Assigning Addresses from a Dynamic Pool for Static Use on page 143
- Example: Configuring NAT Rules Without Defining a Pool on page 144
- Example: Preventing Translation of Specific Addresses on page 144
- Example: Configuring NAT for Multicast Traffic on page 145
- Example: Configuring Twice NAT on page 149

- Example: Configuring Full-Cone NAT on page 150
- Example: Configuring NAT in Mixed IPv4 and IPv6 Networks on page 150

For additional examples that combine NAT configuration with other services and with virtual private network (VPN) routing and forwarding (VRF) tables, see "Examples: Services Interfaces Configuration" on page 49.

Example: Configuring Dynamic Source Translation

The following example configures dynamic source translation:

```
[edit services nat]
pool public {
    address-range low 192.16.2.1 high 192.16.2.32;
    port automatic;
}
rule Private-Public {
    match-direction input;
    term Translate {
        then {
            translated {
                source-pool public;
                translation-type source dynamic;
            }
        }
    }
}
```

Example: Configuring Static Source Translation

The following configuration sets up one-to-one mapping between a private subnet and a public subnet:

```
[edit services nat]
pool mypool {
    address 192.16.1.0/28; # public subnet
}
rule src-nat {
    match-direction input;
    term t1 {
        from {
            source-address 10.150.1.0/28; # private subnet
        }
        then {
            translated {
                source-pool mypool;
                translation-type source static;
            }
        }
    }
}
```

Example: Configuring Dynamic and Static Source Translation

In the following configuration, term1 configures source address translation for traffic from any private address to any public address. The translation is applied for all services. term2 performs destination address translation for Hypertext Transfer Protocol (HTTP) traffic from any public address to the server's virtual IP address. The virtual server IP address is translated to an internal IP address.

```
[edit services nat]
rule my-nat-rule {
   match-direction input;
   term my-term1 {
      from {
         source-address private;
         destination-address public;
      }
      then {
         translated {
            source-pool my-prefix-list; # pick address from a pool
            translation-type source dynamic; # dynamic NAT with port translation
         }
      }
   }
   term my-term2 {
      from {
         destination-address 192.168.137.3; # my server's virtual address
         application http;
      }
      then {
         translated {
            destination-pool nat-pool-name;
            translation-type destination static; # static destination NAT
         }
      }
   }
}
```

Example: Configuring an Oversubscribed Pool with No Fallback

The following configuration shows dynamic address translation from a large prefix to a small pool, translating a /24 subnet to a pool of 10 addresses. Sessions from the first 10 host sessions are assigned an address from the pool on a first-come, first-served basis, and any additional requests are rejected. Each host with an assigned NAT addresses can participate in multiple sessions.

```
[edit nat services]
pool my-pool {
   address-range low 10.10.10.1 high 10.10.10.10;
}
rule src-nat {
   match-direction input;
   term t1 {
      from {
```

```
            source-address 192.168.1.0/24;
        }
        then {
            translated {
                translation-type source dynamic;
                source-pool my-pool;
            }
        }
    }
}
```

Example: Configuring an Oversubscribed Pool with Fallback to NAPT

The following configuration shows dynamic address translation from a large prefix to a small pool, translating a /24 subnet to a pool of 10 addresses. When the addresses in the source pool (src-pool) are exhausted, network address translation is provided by the NAPT overload pool (pat-pool).

```
[edit services nat]
pool src-pool {
    address-range low 192.16.2.1 high 192.16.2.10;
}
pool pat-pool {
    address-range low 192.2.11 high 192.16.2.12;
    port automatic;
}
rule myrule {
    match-direction input;
    term myterm {
        from {
            source-address 10.150.1.0/24;
        }
        then {
            translated {
                source-pool src-pool;
                overload-pool pat-pool;
                translation-type source dynamic;
            }
        }
    }
}
```

Example: Configuring Static Source Translation with Multiple Prefixes and Address Ranges

The following configuration creates a static pool with an address prefix and an address range and uses static source NAT translation.

```
[edit services nat]
pool p1 {
    address 30.30.30.252/30;
    address-range low 20.20.20.1 high 20.20.20.2;
}
rule r1 {
    match-direction input;
```

```
    term {
        from {
            source-address {
                10.10.10.252/30;
            }
        }
        then {
            translated {
                source-pool p1;
                translation-type source static;
            }
        }
    }
}
```

Example: Assigning Addresses from a Dynamic Pool for Static Use

The following configuration statically assigns a subset of addresses that are configured as part of a dynamic pool (dynamic-pool) to two separate static pools (static-pool and static-pool2).

```
[edit services nat]
pool dynamic-pool {
    address 20.20.10.0/24;
}
pool static-pool {
    address-range low 20.20.10.10 high 10.20.10.12;
}
pool static-pool2 {
    address 20.20.10.15/32;
}
rule src-nat {
    match-direction input;
    term t1 {
        from {
            source-address 30.30.30.0/24;
        }
        then {
            translation-type source dynamic;
            source-pool dynamic-pool;
        }
    }
    term t2 {
        from {
            source-address 10.10.10.2;
        }
        then {
            translation-type source static;
            source-pool static-pool;
        }
    }
    term t3 {
        from {
            source-address 10.10.10.10;
        }
```

```
        then {
            translation-type source static;
            source-pool static-pool2;
        }
    }
}
```

Example: Configuring NAT Rules Without Defining a Pool

The following configuration performs network address translation using the source prefix 20.20.10.0/24 without defining a pool.

```
[edit services nat]
rule src-nat {
    match-direction input;
    term t1 {
        then {
            translation-type source dynamic;
            source-prefix 20.20.10.0/24;
        }
    }
}
```

The following configuration performs network address translation using the destination prefix 20.20.10.0/32 without defining a pool.

```
[edit services nat]
rule src-nat {
    match-direction input;
    term t1 {
        from {
            destination-address 10.10.10.10/32;
            then {
                translation-type destination static;
                destination-prefix 20.20.10.0/32;
            }
        }
    }
}
```

Example: Preventing Translation of Specific Addresses

The following configuration specifies that network address translation is not performed on incoming traffic from the source address 192.168.20.24/32. Dynamic NAT is performed on all other incoming traffic.

```
[edit services nat]
pool my-pool {
    address-range low 10.10.10.1 high 10.10.10.16;
    port-automatic;
}
rule src-nat {
    match-direction input;
    term t0 {
```

```
        from {
            source-address 192.168.20.24/32;
        }
        then {
            no-translation;
        }
    }
    term t1 {
        then {
            translated {
                translation-type source dynamic;
                source-pool my-pool;
            }
        }
    }
}
```

Example: Configuring NAT for Multicast Traffic

Figure 2 on page 145 illustrates the network setup for the following configuration, which allows IP multicast traffic to be sent to the Adaptive Services (AS) or Multiservices PIC.

Figure 2: Configuring NAT for Multicast Traffic

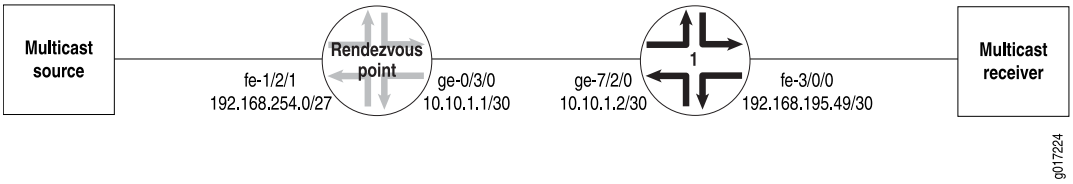

Rendezvous Point Configuration

On the rendezvous point (RP), all incoming traffic from the multicast source at 192.168.254.0/27 is sent to the static NAT pool mcast_pool, where its source is translated to 20.20.20.0/27. The service set nat_ss is a next-hop service set that allows IP multicast traffic to be sent to the AS or Multiservices PIC. The inside interface on the PIC is sp-1/1/0.1 and the outside interface is sp-1/1/0.2.

```
[edit services]
nat {
    pool mcast_pool {
        address 20.20.20.0/27;
    }
    rule nat_rule_1 {
        match-direction input;
        term 1 {
            from {
                source-address 192.168.254.0/27;
            }
        }
        then {
            translated {
```

```
            source-pool mcast_pool;
            translation-type source static;
          }
          syslog;
        }
      }
    }
    service-set nat_ss {
      allow-multicast;
      nat-rules nat_rule_1;
      next-hop-service {
        inside-service-interface sp-1/1/0.1;
        outside-service-interface sp-1/1/0.2;
      }
    }
```

The Gigabit Ethernet interface **ge-0/3/0** carries traffic out of the RP to Router 1. The adaptive services interface **sp-1/1/0** has two logical interfaces: **unit 1** is the inside interface for next-hop services and **unit 2** is the outside interface for next-hop services. Multicast source traffic comes in on the Fast Ethernet interface **fe-1/2/1**, which has the firewall filter **fbf** applied to incoming traffic.

```
[edit interfaces]
ge-0/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.1/30;
    }
  }
}
sp-1/1/0 {
  unit 0 {
    family inet;
  }
  unit 1 {
    family inet;
    service-domain inside;
  }
  unit 2 {
    family inet;
    service-domain outside;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      filter {
        input fbf;
      }
      address 192.168.254.27/27;
    }
  }
}
```

Multicast packets can only be directed to the AS or Multiservices PIC using a next-hop service set. In the case of NAT, you must also configure a VRF. Therefore, the routing instance **stage** is created as a "dummy" forwarding instance. To direct incoming packets to **stage**, you configure filter-based forwarding through a firewall filter called **fbf**, which is applied to the incoming interface **fe-1/2/1**. A lookup is performed in **stage.inet.0**, which has a multicast static route that is installed with the next hop pointing to the PIC's inside interface. All multicast traffic matching this route is sent to the PIC.

```
[edit firewall]
filter fbf {
   term 1 {
      then {
         routing-instance stage;
      }
   }
}
```

The routing instance **stage** forwards IP multicast traffic to the inside interface **sp-1/1/0.1** on the AS or Multiservices PIC:

```
[edit]
routing-instances stage {
   instance-type forwarding;
   routing-options {
      static {
         route 224.0.0.0/4 next-hop sp-1/1/0.1;
      }
   }
}
```

You enable OSPF and Protocol Independent Multicast (PIM) on the Fast Ethernet and Gigabit Ethernet logical interfaces over which IP multicast traffic enters and leaves the RP. You also enable PIM on the outside interface (**sp-1/1/0.2**) of the next-hop service set.

```
[edit protocols]
ospf {
   area 0.0.0.0 {
      interface fe-1/2/1.0 {
         passive;
      }
      interface lo0.0;
      interface ge-0/3/0.0;
   }
}
pim {
   rp {
      local {
         address 10.255.14.160;
      }
   }
   interface fe-1/2/1.0;
   interface lo0.0;
   interface ge-0/3/0.0;
```

```
    interface sp-1/1/0.2;
}
```

As with any filter-based forwarding configuration, in order for the static route in the forwarding instance stage to have a reachable next hop, you must configure routing table groups so that all interface routes are copied from inet.0 to the routing table in the forwarding instance. You configure routing tables inet.0 and stage.inet.0 as members of fbf_rib_group, so that all interface routes are imported into both tables.

```
[edit routing-options]
interface-routes {
    rib-group inet fbf_rib_group;
}
rib-groups fbf_rib_group {
    import-rib [ inet.0 stage.inet.0 ];
}
multicast {
    rpf-check-policy no_rpf;
}
```

Reverse path forwarding (RPF) checking must be disabled for the multicast group on which source NAT is applied. You can disable RPF checking for specific multicast groups by configuring a policy similar to the one in the example that follows. In this case, the no_rpf policy disables RPF check for multicast groups belonging to 224.0.0.0/4.

```
[edit policy-options]
policy-statement no_rpf {
    term 1 {
        from {
            route-filter 224.0.0.0/4 orlonger;
        }
        then reject;
    }
}
```

Router 1 Configuration

The Internet Group Management Protocol (IGMP), OSPF, and PIM configuration on Router 1 is as follows. Because of IGMP static group configuration, traffic is forwarded out fe-3/0/0.0 to the multicast receiver without receiving membership reports from host members.

```
[edit protocols]
igmp {
    interface fe-3/0/0.0 {
    }
}
ospf {
    area 0.0.0.0 {
        interface fe-3/0/0.0 {
            passive;
        }
        interface lo0.0;
        interface ge-7/2/0.0;
```

```
    }
    pim {
        rp {
            static {
                address 10.255.14.160;
            }
        }
        interface fe-3/0/0.0;
        interface lo0.0;
        interface ge-7/2/0.0;
    }
}
```

The routing option creates a static route to the NAT pool, mcast_pool, on the RP.

```
[edit routing-options]
static {
    route 20.20.20.0/27 next-hop 10.10.1.1;
}
```

Example: Configuring Twice NAT

In the following configuration, term1 configures source address translation and destination address translation for traffic from a specific destination address to a source address in a range of source addresses. Both destination and source pools are configured.

```
[edit services nat]
rule twice-nat {
    match-direction input;
    term my-term1 {
        from {
            destination-address {
                41.41.41.41/32;
            }
            source-address-range {
                low 10.58.254.34 high 10.58.254.35;
            }
        }
        then {
            translated {
                source-pool src-pool;
                destination-pool dst_pool;
                translation-type {
                    source static;
                    destination static;
                }
            }
        }
    }
}
```

Example: Configuring Full-Cone NAT

The following configuration example shows full-cone NAT with source static processing.

```
[edit services]
nat {
    pool static_nat_range {
        address-range low 10.200.253.1 high 10.200.253.5;
    }
    rule static_nat_rule {
        term sample-term {
            nat-type full-cone;
            from {
                source-address-range {
                    low 10.100.136.1 high 10.100.136.5;
                }
            }
            then {
                translated {
                    source-pool static_nat_range;
                    translation-type {
                        source static;
                    }
                }
            }
        }
    }
}
```

Example: Configuring NAT in Mixed IPv4 and IPv6 Networks

The following example shows NAT configuration in a network that combines IPv4 and IPv6 addressing.

```
interfaces {
    sp-2/0/0 {
        traceoptions {
            flag all;
        }
        services-options {
            syslog {
                host local {
                    services any;
                    log-prefix IPV6-SS;
                }
            }
            inactivity-timeout 200;
        }
        unit 0 {
            family inet;
            family inet6;
        }
        unit 1 {
```

```
        family inet;
        family inet6;
      }
      unit 1001 {
        family inet;
        family inet6;
      }
    }
    so-2/1/0 {
      description "services-art1 201/1";
      encapsulation ppp;
      sonet-options {
        fcs 32;
      }
      unit 0 {
        family inet {
          address 192.168.1.1/30;
        }
        family inet6 {
          service {
            input {
              service-set ss-ipv6;
            }
            output {
              service-set ss-ipv6;
            }
          }
          address 3ffe::1:1/64;
        }
      }
    }
    so-2/1/3 {
      description "services-art1 201/2";
      encapsulation ppp;
      sonet-options {
        fcs 32;
      }
      unit 0 {
        family inet {
          address 192.168.1.1/30;
        }
        family inet6 {
          address 4ffe::1:1/64;
        }
      }
    }
}
    routing-options {
      rib inet6.0 {
        static {
          route 5ffe::0:0/64 next-hop 3ffe::1:1;
          route 6ffe::0:0/64 next-hop 4ffe::1:1;
          route 192::168:1:0/112 next-hop 4ffe::1:2;
        }
      }
    }
```

```
services {
    service-set ss-ipv6 {
        stateful-firewall-rules test-ipv6;
        nat-rules src-nat;
        nat-rules src-nat-v6;
        interface-service {
            service-interface sp-2/0/0;
        }
    }
    stateful-firewall {
        rule test-ipv6 {
            match-direction input-output;
            term 1 {
                from {
                    source-address {
                        any-unicast;
                    }
                }
                then {
                    accept;
                }
            }
        }
    }
    nat {
        pool dst_pool {
            address 192.168.1.2/32;
        }
        pool src_pool {
            address 192.168.1.0/27;
            port automatic;
        }
        pool dst_pool_v6 {
            address 192::168:1:2/128;
        }
        pool src_pool_v6 {
            address 192::168:1:2/100;
        }
        rule src-nat {
            match-direction input;
            term t1 {
                from {
                    source-address {
                        3ffe::0:0/96;
                    }
                    destination-address {
                        4ffe::1:2/128;
                    }
                }
                then {
                    translated {
                        source-pool src_pool;
                        destination-pool dst_pool;
                        translation-type {
                            source dynamic;
                            destination static;
```

```
                }
            }
        }
    }
    rule src-nat-v6 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    5ffe::0:0/96;
                }
                destination-address {
                    6ffe::1:2/128;
                }
            }
            then {
                translated {
                    source-pool src_pool_v6;
                    destination-pool dst_pool_v6;
                    translation-type {
                        source static;
                        destination static;
                    }
                }
            }
        }
    }
}
```

Chapter 9

Summary of Network Address Translation Configuration Statements

The following sections explain each of the Network Address Translation (NAT) statements. The statements are organized alphabetically.

address

Syntax address *ip-prefix</prefix-length>*;

Hierarchy Level [edit services nat pool *nat-pool-name*]

Release Information Statement introduced before JUNOS Release 7.4.
prefix option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.

Description Specify the NAT pool prefix value.

Options *prefix*—Specify an IPv4 or IPv6 prefix value.

Usage Guidelines See "Configuring Addresses and Ports for Use in NAT Rules" on page 130.

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

address-range

Syntax address-range low *minimum-value* high *maximum-value*;

Hierarchy Level [edit services nat pool *nat-pool-name*]

Release Information Statement introduced before JUNOS Release 7.4.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.

Description Specify the NAT pool address range.

Options *minimum-value*—Lower boundary for the IPv4 or IPv6 address range.

maximum-value—Upper boundary for the IPv4 or IPv6 address range.

Usage Guidelines See "Configuring Addresses and Ports for Use in NAT Rules" on page 130.

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

application-sets

Syntax applications-sets *set-name*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 NAT Rules" on page 136.

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 nat rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define one or more application protocols to which the NAT services apply.

Options *application-name*—Name of the target application.

Usage Guidelines See "Configuring Match Conditions in NAT Rules" on page 136.

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

destination-address

Syntax destination-address (*address* | any-unicast) <except>;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS Release 7.4.
any-unicast and except options introduced in JUNOS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS 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) Prevent the specified address, prefix, or unicast packets from being translated.

Usage Guidelines See "Configuring Match Conditions in NAT Rules" on page 136.

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 nat rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 7.6.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS 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) Prevent the specified address range from being translated.

Usage Guidelines See "Configuring Match Conditions in NAT Rules" on page 136.

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

destination-pool

Syntax destination-pool *nat-pool-name*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the destination address pool for translated traffic.

Options *nat-pool-name*—Destination pool name.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

destination-prefix

Syntax destination-prefix *destination-prefix*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced in JUNOS Release 7.6.
destination-prefix option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.

Description Specify the destination prefix for translated traffic.

Options *destination-prefix*—IPv4 or IPv6 destination prefix value.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

destination-prefix-list

Syntax destination-prefix-list *list-name* <except>;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS 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 NAT Rules" on page 136.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

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>;
    source-address address (address | any-unicast) <except>;
    source-address-range low minimum-value high maximum-value <except>;
}
```

Hierarchy Level [edit services nat rule *rule-name* term *term-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify input conditions for the NAT term.

Options For information on match conditions, see the description of firewall filter match conditions in the *JUNOS Policy Framework Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Configuring NAT Rules" on page 133.

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

hint

Syntax hint [*hint-strings*];

Hierarchy Level [edit services nat pool *nat-pool-name* pgcp]

Release Information Statement introduced in JUNOS Release 9.0.

Description Configure a hint that enables the BGF to choose a NAT pool by direction rather than by virtual interface. The BGF matches the configured hint with a termination hint located in the Direction field of a nonstandard termination ID.

Default When no hint is configured, the BGF can choose any NAT pool associated with the virtual interface.

Options *hint-string*—Alphanumeric string of up to 3 characters that the BGF uses to match with a termination hint located in the Direction field of a nonstandard termination ID. You can also specify underscores (_) and hyphens (-) within the string. To specify a list of hints, use the format: [hint *xx* hint *yy*].

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

Related Topics *JUNOS Multiplay Solutions Guide*

ipv6-multicast-interfaces

Syntax
```
ipv6-multicast-interfaces (all | interface-name) {
    disable;
}
```

Hierarchy Level [edit services nat]

Release Information Statement introduced in JUNOS Release 9.1.

Description Enable multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery.

Options all—Enable filters on all interfaces.

disable—Disable filters on the specified interfaces.

interface-name—Enable filters on a specific interface only.

Usage Guidelines See “Configuring IPv6 Multicast Filters” on page 133.

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);

Hierarchy Level [edit services nat rule *rule-name*]

Release Information Statement introduced before JUNOS 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 NAT Rules" on page 133.

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

nat-type

Syntax nat-type (full-cone | symmetric);

Hierarchy Level [edit services nat rule *rule-name* term *term-name*]

Release Information Statement introduced in JUNOS Release 8.5.

Description Specify whether the term supports full-cone or traditional (symmetric) NAT.

Default symmetric

Options full-cone—Support full-cone NAT processing, in which all requests from the same internal IP address and port are mapped to the same external IP address and port.

symmetric—Support traditional NAT address matching only.

Usage Guidelines See "Configuring NAT Type for Terms in NAT Rules" on page 135

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

no-translation

Syntax no-translation;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then]

Release Information Statement introduced in JUNOS Release 7.6.

Description Specify that traffic is not to be translated.

Options none

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

overload-pool

Syntax overload-pool *overload-pool-name*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced in JUNOS Release 7.6.

Description Specify an address pool that can be used if the source pool becomes exhausted.

Options *overload-pool-name*—Name of the overload pool.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

overload-prefix

Syntax overload-prefix *overload-prefix*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced in JUNOS Release 7.6.

Description Specify the prefix that can be used if the source pool becomes exhausted.

Options *overload-prefix*—Prefix value.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

pgcp

Syntax
```
pgcp {
   hint [ hint-strings ];
   ports-per-session ports;
   remotely-controlled;
   transport [ transport-protocols ];
}
```

Hierarchy Level [edit services nat pool *nat-pool-name*]

Release Information Statement introduced in JUNOS Release 8.4.
remotely-controlled and ports-per-session statements added in JUNOS Release 8.5.
hint statement added in JUNOS Release 9.0.

Description Specify that the NAT pool is used exclusively by the BGF.

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

Related Topics *JUNOS Multiplay Solutions Guide*

pool

Syntax

```
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 [ transport-protocols ];
   }
   port (automatic | range low minimum-value high maximum-value);
}
```

Hierarchy Level [edit services nat]

Release Information Statement introduced before JUNOS Release 7.4.
pgcp statement added in JUNOS Release 8.4.
remotely-controlled and ports-per-session statements added in JUNOS Release 8.5.
hint statement added in JUNOS Release 9.0.

Description Specify the NAT name and properties.

Options *nat-pool-name*—Identifier for the NAT address pool.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Addresses and Ports for Use in NAT Rules" on page 130.

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

port

Syntax
```
port (automatic | range low minimum-value high maximum-value) {
    random-allocation;
}
```

Hierarchy Level [edit services nat pool *nat-pool-name*]

Release Information port statement introduced before JUNOS Release 7.4.
random-allocation statement introduced in JUNOS Release 9.3.

Description Specify the NAT pool port or range. You can configure an automatically assigned port or specify a range with minimum and maximum values.

Options automatic—Router-assigned port.

minimum-value—Lower boundary for the port range.

maximum-value—Upper boundary for the port range.

random-allocation—By default, the JUNOS Software allocates NAT ports sequentially. To configure random port allocation, include the random-allocation option.

Usage Guidelines See "Configuring Addresses and Ports for Use in NAT Rules" on page 130.

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

ports-per-session

Syntax ports-per-session *ports*;

Hierarchy Level [edit services nat pool *nat-pool-name* pgcp]

Release Information Statement introduced in JUNOS Release 8.4.

Description Configure the number of ports required to support Real-Time Transport Protocol (RTP), Real-Time Control Protocol (RTCP), Real-Time Streaming Protocol (RTSP), forward error correction (FEC) for voice and video flows on the Multiservices PIC.

Options *number-of-ports*—Number of ports to enable: 2 or 4 for combined voice and video services.
Default: 2

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

Related Topics *JUNOS Multiplay Solutions Guide*

remotely-controlled

Syntax remotely-controlled;

Hierarchy Level [edit services nat pool *nat-pool-name* pgcp]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure the addresses and ports in a NAT pool to be remotely controlled 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 Topics *JUNOS Multiplay Solutions Guide*

rule

Syntax

```
rule rule-name {
    match-direction (input | output);
    term term-name {
        nat-type (full-cone | symmetric);
        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 {
            no-translation;
            translated {
                destination-pool nat-pool-name;
                destination-prefix destination-prefix;
                overload-pool overload-pool;
                overload-prefix overload-prefix;
                source-pool nat-pool-name;
                source-prefix source-prefix;
                translation-type (destination type | source type);
                translation-type {
                    source type;
                    destination type;
                }
            }
            syslog;
        }
    }
}
```

Hierarchy Level [edit services nat],
[edit services nat rule-set *rule-set-name*]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See “Configuring NAT Rules” on page 133.

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

rule-set

Syntax
```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services nat]

Release Information Statement introduced before JUNOS 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 NAT Rule Sets" on page 139.

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

services

Syntax services nat { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the service rules to be applied to traffic.

Options nat—Identifies the NAT set of rules statements.

Usage Guidelines See "Network Address Translation Services Configuration Guidelines" on page 129.

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* | any-unicast) <except>;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS Release 7.4.
any-unicast and except options introduced in JUNOS Release 7.6.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS 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) Prevent the specified address or unicast packets from being translated.

Usage Guidelines See "Configuring Match Conditions in NAT Rules" on page 136.

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 nat rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 7.6.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS 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) Prevent the specified address range from being translated.

Usage Guidelines See "Configuring Match Conditions in NAT Rules" on page 136.

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

source-pool

Syntax source-pool *nat-pool-name*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the source address pool for translated traffic.

Options *nat-pool-name*—Source pool name.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

source-prefix

Syntax source-prefix *source-prefix*;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced in JUNOS Release 7.6.
source-prefix option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.

Description Specify the source prefix for translated traffic.

Options *source-prefix*—IPv4 or IPv6 source prefix value.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

source-prefix-list

Syntax source-prefix-list *list-name* <except>;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS 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 NAT Rules" on page 136.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

syslog

Syntax syslog;

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

term

Syntax

```
term term-name {
    nat-type (full-cone | symmetric);
    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 {
        no-translation;
        translated {
            destination-pool nat-pool-name;
            destination-prefix destination-prefix;
            overload-pool overload-pool;
            overload-prefix overload-prefix;
            source-pool nat-pool-name;
            source-prefix source-prefix;
            translation-type (destination type | source type);
            translation-type {
                source type;
                destination type;
            }
        }
        syslog;
    }
}
```

Hierarchy Level [edit services nat rule *rule-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the NAT term properties.

Options *term-name*—Identifier for the term.

Usage Guidelines See “Configuring NAT Rules” on page 133.

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

then

Syntax

```
then {
    no-translation;
    translated {
        destination-pool nat-pool-name;
        destination-prefix destination-prefix;
        overload-pool overload-pool-name;
        overload-prefix overload-prefix;
        source-pool nat-pool-name;
        source-prefix source-prefix;
        translation-type (destination type | source type);
        translation-type {
            source type;
            destination type;
        }
    }
    syslog;
}
```

Hierarchy Level [edit services nat rule *rule-name* term *term-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the NAT term actions.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring NAT Rules” on page 133.

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

translated

Syntax

```
translated {
    destination-pool nat-pool-name;
    source-pool nat-pool-name;
    translation-type (destination type | source type);
    translation-type {
        source type;
        destination type;
    }
}
```

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define properties for translated traffic.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

translation-type

See the following sections:

- translation-type (Traditional NAT) on page 176
- translation-type (Twice NAT) on page 176

translation-type (Traditional NAT)

Syntax translation-type (destination *type* | source *type*)

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the NAT types for traditional NAT.

Options *type*—You can specify source dynamic, source static, or destination static.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

translation-type (Twice NAT)

Syntax
```
translation-type {
    source type;
    destination type;
}
```

Hierarchy Level [edit services nat rule *rule-name* term *term-name* then translated]

Release Information Statement introduced in JUNOS Release 8.3.

Description Specify the NAT types for twice NAT.

Options *type*—You must specify destination static and either source dynamic or source static.

Usage Guidelines See "Configuring Actions in NAT Rules" on page 137.

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

transport

Syntax transport [*transport-protocols*];

Hierarchy Level [edit services nat pool *nat-pool-name* pgcp]

Release Information Statement introduced in JUNOS Release 9.2.

Description Configure the BGF to select a NAT pool based on transport protocol type.

Options [*transport-protocol*]—One or more transport protocols.
Values: rtp-avp, tcp, udp
Syntax: One or more protocols. If you specify more than one protocol, you must enclose all protocols in brackets.

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

Related Topics *JUNOS Multiplay Solutions Guide*

Chapter 10

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 {
                    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 Software 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 181
- Configuring IDS Rule Sets on page 187
- Examples: Configuring IDS Rules on page 188

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 "Stateful Firewall Services Configuration Guidelines" on page 107.

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;
                }
```

```
            }
            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 182
- Configuring Match Conditions in IDS Rules on page 183
- Configuring Actions in IDS Rules on page 184

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 440.

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 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 *JUNOS Policy Framework 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 112.

You can also include application protocol definitions that you have configured at the [edit applications] hierarchy level; for more information, see "Applications Configuration Guidelines" on page 59.

- 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 System Basics and Services Command Reference*.

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]
   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;
        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;
    }
  }
```

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;
      }
      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;
        }
        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;
                }
            }
        }
    }
}
```

Chapter 11

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 Release 7.4.

Description Specify the type of data to be aggregated.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Actions in IDS Rules" on page 184.

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

application-sets

Syntax application-sets *set-name*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 183.

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

applications

Syntax applications [*application-name*s];

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 183.

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

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 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 184.

Required Privilege Level

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

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 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 184.

Required Privilege Level

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

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 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 184.

Required Privilege Level

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

destination-address

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 Release 7.4.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS 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 183.

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 ids rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 7.6.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS 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 183.

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

destination-prefix

Syntax destination-prefix *prefix-value*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then aggregation]

Release Information Statement introduced before JUNOS 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 184.

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 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 184.

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

destination-prefix-list

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 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 183.

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

Related Topics *JUNOS Policy Framework 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 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 184.

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

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>;
    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 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 *JUNOS Policy Framework Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Conditions in IDS Rules" on page 183.

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

Syntax
```
logging {
    syslog;
    threshold rate;
}
```

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 184.

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 ids rule *rule-name*]

Release Information Statement introduced before JUNOS 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 183.

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 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 184.

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

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>;
            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*]

Release Information Statement introduced before JUNOS 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 181.

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

rule-set

Syntax
```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services ids]

Release Information Statement introduced before JUNOS 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 187.

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

services

Syntax
```
services ids { ... }
```

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS 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 "Intrusion Detection Service Configuration Guidelines" on page 179.

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

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 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 184.

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* | any-unicast) <except>;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS Release 7.4.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS 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 183.

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 ids rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 7.6.
minimum-value and *maximum-value* options enhanced to support IPv4 and IPv6 addresses in JUNOS 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 183.

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

source-prefix

Syntax source-prefix *prefix-value*;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then aggregation]

Release Information Statement introduced before JUNOS 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 184.

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 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 184.

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

source-prefix-list

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 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 183.

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

Related Topics *JUNOS Policy Framework 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 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 184.

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

syslog

Syntax syslog;

Hierarchy Level [edit services ids rule *rule-name* term *term-name* then logging]

Release Information Statement introduced before JUNOS 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 184.

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

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>;
        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 Release 7.4.

Description Define the IDS term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines See “Configuring IDS Rules” on page 181.

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

then

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 Release 7.4.

Description Define the IDS term actions.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring IDS Rules” on page 181.

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

threshold

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 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 184.

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

Chapter 12

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);
        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);
        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 {
            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;
}
```

This chapter includes the following sections:

- Minimum Security Association Configurations on page 215
- Configuring Security Associations on page 216
- Configuring IKE Proposals on page 222
- Configuring IKE Policies on page 225
- Configuring IPsec Proposals on page 231
- Configuring IPsec Policies on page 233
- Configuring IPsec Rules on page 235
- Configuring IPsec Rule Sets on page 242
- Configuring Dynamic Endpoints for IPsec Tunnels on page 242
- Tracing IPsec Operations on page 248
- Examples: Configuring IPsec Services on page 249

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 215
- Minimum Dynamic SA Configuration on page 215

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:

```
[edit services ipsec-vpn]
```

```
ike {
    proposal proposal-name {
        authentication-algorithm (md5 | sha1 | sha-256);
        authentication-method pre-shared-keys;
        dh-group (group1 | group2);
        encryption-algorithm algorithm;
    }
    policy policy-name {
        proposals [ ike-proposal-names ];
        pre-shared-key (ascii-text key | hexadecimal key);
    }
}
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);
    }
}
```

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 217.
- 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 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 221.

This section includes the following topics:

- Configuring Manual Security Associations on page 217
- Configuring Dynamic Security Associations on page 221
- Clearing Security Associations on page 222

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 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 217
- Configuring the Protocol for a Manual IPsec SA on page 218
- Configuring the Security Parameter Index on page 219
- Configuring the Auxiliary Security Parameter Index on page 219
- Configuring Authentication for a Manual IPsec SA on page 219
- Configuring Encryption for a Manual IPSec SA on page 220

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 23456789012345678901234;
   }
}
direction outbound {
   protocol esp;
   spi 24576;
   encryption {
      algorithm 3des-cbc;
      key ascii-text 123456789012345678901abcd;
   }
}
```

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;
   }
}
```

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:

```
[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.

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 Software 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 225 and “Configuring IKE Proposals” on page 222. For more information about IPsec policies and proposals, see “Configuring IPsec Policies” on page 233.

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 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.

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);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
}
```

This section includes the following topics:

- Configuring the Authentication Algorithm for an IKE Proposal on page 223
- Configuring the Authentication Method for an IKE Proposal on page 223
- Configuring the Diffie-Hellman Group for an IKE Proposal on page 224
- Configuring the Encryption Algorithm for an IKE Proposal on page 224
- Configuring the Lifetime for an IKE SA on page 225
- Example: Configuring an IKE Proposal on page 225

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:

- 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);
```

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.

group2 provides more security but requires more 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.
- 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 Software 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: For IKE proposals, there is only one SA lifetime value, specified by the JUNOS Software. IPsec proposals use a different mechanism; for more information, see "Configuring the Lifetime for an IPsec SA" on page 232.

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.

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);
    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 Mode for an IKE Policy on page 227
- Configuring the Proposals in an IKE Policy on page 227
- Configuring the Preshared Key for an IKE Policy on page 227
- Configuring the Local Certificate for an IKE Policy on page 228
- Configuring the Description for an IKE Policy on page 229
- Configuring Local and Remote IDs for IKE Phase 1 Negotiation on page 229
- Example: Configuring an IKE Policy on page 230

For an example of an IKE policy configuration, see "Example: Configuring an IKE Policy" on page 230.

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.

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 policyl, 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 223. 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 223. You must identify a local certificate that is sent to the peer during the IKE authentication phase.

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 System Basics Configuration Guide*. For complete examples of digital certificate configuration, see the *JUNOS Feature Guide*.

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 445.

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 System Basics Configuration Guide*.

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 242.

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.

```
[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;
      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 System Basics and Services Command Reference*.

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);
}
```

This section discusses the following topics:

- Configuring the Authentication Algorithm for an IPsec Proposal on page 231
- Configuring the Description for an IPsec Proposal on page 231
- Configuring the Encryption Algorithm for an IPsec Proposal on page 232
- Configuring the Lifetime for an IPsec SA on page 232
- Configuring the Protocol for a Dynamic SA on page 233

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.

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 Software 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.

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.

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);
    }
```

```
    proposals [ proposal-names ];
}
```

This section includes the following topics related to configuring an IPsec policy:

- Configuring the Description for an IPsec Policy on page 234
- Configuring Perfect Forward Secrecy on page 234
- Configuring the Proposals in an IPsec Policy on page 234
- Example: Configuring an IPsec Policy on page 235

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);
}
```

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.

group2 provides more security than **group1**, but requires 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 ];
```

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 System Basics and Services Command Reference*.

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.
- 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 236
- Configuring Match Conditions in IPsec Rules on page 237
- Configuring Actions in IPsec Rules on page 239

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 440.

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 *JUNOS Policy Framework 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** (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 440 and "Service Interface Configuration Guidelines" on page 479.

The JUNOS Software 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.

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 221.
- 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 217.

You can configure the following additional properties:

- Enabling IPsec Packet Fragmentation on page 240
- Configuring Destination Addresses for Dead Peer Detection on page 240
- Configuring or Disabling IPSec Anti-Replay on page 241

- Enabling System Log Messages on page 242
- Specifying the MTU for IPsec Tunnels on page 242

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:

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:

```
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, it would not fail the commit or produce any error message. However, the software internally sets 1024 as the antireplay window size if the configuration's window size exceeds 1024 for AS PICs.

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:

```
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.

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 227. 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 243
- Implicit Dynamic Rules on page 244
- Reverse Route Insertion on page 244
- Configuring an IKE Access Profile on page 245
- Referencing the IKE Access Profile in a Service Set on page 246
- Configuring the Interface Identifier on page 247
- Default IKE and IPsec Proposals on page 247

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 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 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 240.

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.

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 225.

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 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>;
        }
    }
}
```

NOTE: For dynamic peers, the JUNOS Software 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 225.

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.

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;
}
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. All interfaces referenced by the ipsec-inside-interface statement within a service set must belong to the same VRF.

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; if more than one value is shown, the first value is the default. For more information on IKE proposals, see “Configuring IKE Proposals” on page 222; for more information on IPsec proposals, see “Configuring IPsec Proposals” on page 231.

NOTE: RSA certificates are not supported with dynamic endpoint configuration.

Table 11: Default IKE and IPsec Proposals for Dynamic Negotiations

Statement Name	Values
Implicit IKE Proposal	
authentication-method	pre-shared keys
dh-group	group1, group2
authentication-algorithm	sha1, md5, sha-256
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256

Table 11: Default IKE and IPsec Proposals for Dynamic Negotiations *(continued)*

Statement Name	Values
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 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;
}
```

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.

```
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 445. 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.

Examples: Configuring IPsec Services

See the following sections:

- Example: Configuring Statically Assigned Tunnels on page 249
- Example: Configuring Dynamically Assigned Tunnels on page 252

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 {
    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 {
    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;
            }
         }
      }
   }
   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 3 on page 253):

- 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.

Figure 3: IPsec Dynamic Endpoint Tunneling Topology

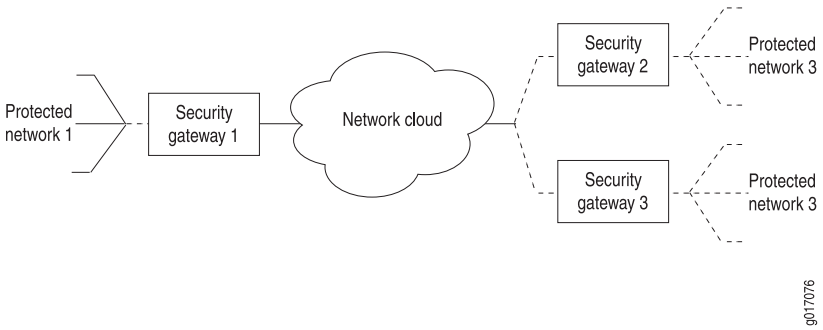

The examples in this section show the following configurations:

- Configuring a Next-Hop Style Service Set with Link-Type Tunnels on page 253
- Configuring a Next-Hop Style Service-Set with Policy-Based Tunnels on page 255

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.

```
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
    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;
            }
        }
    }
}
# 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
```

Chapter 13

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

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 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 241.

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

authentication

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 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 219.

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

authentication-algorithm

See the following sections:

- authentication-algorithm (IKE) on page 259
- authentication-algorithm (IPsec) on page 259

authentication-algorithm (IKE)

Syntax authentication-algorithm (md5 | sha1 | sha-256);

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.
sha-256 option added in JUNOS 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 223.

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

authentication-algorithm (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 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 231.

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

authentication-method

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 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 223.

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

auxiliary-spi

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 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 219. For information about SPI, see "Configuring the Security Parameter Index" on page 219 and spi.

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

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 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 240.

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

clear-dont-fragment-bit

Syntax clear-dont-fragment-bit;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 239.

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

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 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 222.

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 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 222.

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

description

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 policy *policy-name*],
[edit services ipsec-vpn ipsec proposal *proposal-name*]

Release Information Statement introduced before JUNOS 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 229, "Configuring the Description for an IPsec Proposal" on page 231, and "Configuring the Description for an IPsec Policy" on page 234.

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

destination-address

Syntax destination-address *address*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 237.

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

dh-group

Syntax dh-group (group1 | group2);

Hierarchy Level [edit services ipsec-vpn ike proposal *proposal-name*]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See "Configuring the Diffie-Hellman Group for an IKE Proposal" on page 224.

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

direction

Syntax

```
direction (inbound | outbound | bidirectional) {
    protocol (ah | bundle | esp);
    spi spi-value;
    auxiliary-spi spi-value;
    authentication {
        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 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 236.

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

dynamic

Syntax

```
dynamic {
    ike-policy policy-name;
    ipsec-policy policy-name;
}
```

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 221.

Required Privilege Level

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

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 Release 7.4.
aes-128-cbc, aes-192-cbc, and aes-256-cbc options added in JUNOS 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

- aes-256-cbc option, 64 hexadecimal characters

Usage Guidelines See "Configuring Encryption for a Manual IPSec SA" on page 220.

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

encryption-algorithm

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 Release 7.4.
aes-128-cbc, aes-192-cbc, and aes-256-cbc options added in JUNOS 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 224 and "Configuring the Encryption Algorithm for an IPsec Proposal" on page 232.

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

from

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 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 *JUNOS Policy Framework Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 236.

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

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);
        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);
        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 Release 7.4.

Description Configure IKE.

The statements are explained separately.

Usage Guidelines See “Configuring IKE Proposals” on page 222 and “Configuring IKE Policies” on page 225.

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

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 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 240.

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

Related Topics backup-remote-gateway

ipsec

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 Release 7.4.

Description Configure IPsec.

The statements are explained separately.

Usage Guidelines See "Configuring Security Associations" on page 216.

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

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 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 237 or "Configuring Dynamic Endpoints for IPsec Tunnels" on page 242.

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

lifetime-seconds

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 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 225 and "Configuring the Lifetime for an IPsec SA" on page 232.

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

local-certificate

Syntax local-certificate *identifier*;

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*]

Release Information Statement introduced in JUNOS 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 228.

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 Release 7.4.
ipv6_addr option added in JUNOS 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 229.

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

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 Release 7.4.

Description Define a manual IPsec SA.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Manual Security Associations” on page 217.

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

match-direction

Syntax match-direction (input | output);

Hierarchy Level [edit services ipsec-vpn rule *rule-name*]

Release Information Statement introduced before JUNOS 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 236.

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

mode

Syntax mode (aggressive | main);

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*]

Release Information Statement introduced before JUNOS 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 227.

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

no-anti-replay

Syntax no-anti-replay;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 241.

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

no-ipsec-tunnel-in-traceroute

Syntax no-ipsec-tunnel-in-traceroute;

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced in JUNOS 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 241.

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

perfect-forward-secrecy

Syntax
```
perfect-forward-secrecy {
   keys (group1 | group2);
}
```

Hierarchy Level [edit services ipsec-vpn ipsec policy *policy-name*]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See "Configuring Perfect Forward Secrecy" on page 234.

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

policy

See the following sections:

- policy (IKE) on page 277
- policy (IPsec) on page 278

policy (IKE)

Syntax

```
policy policy-name {
   description description;
   local-certificate identifier;
   local-id (ipv4_addr ipv4-address | ipv6-addr ipv6-address | key-id identifier);
   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 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 225.

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

policy (IPsec)

Syntax
```
policy policy-name {
   description description;
   perfect-forward-secrecy {
      keys (group1 | group2);
   }
   proposals [ proposal-names ];
}
```

Hierarchy Level [edit services ipsec-vpn ipsec]

Release Information Statement introduced before JUNOS 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 233.

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

pre-shared-key

Syntax pre-shared-key (ascii-text *key* | hexadecimal *key*);

Hierarchy Level [edit services ike policy *policy-name*]

Release Information Statement introduced before JUNOS 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 225.

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

proposal

See the following sections:

- proposal (IKE) on page 279
- proposal (IPsec) on page 280

proposal (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);
    encryption-algorithm algorithm;
    lifetime-seconds seconds;
}
```

Hierarchy Level [edit services ipsec-vpn ike]

Release Information Statement introduced before JUNOS 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 222.

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

proposal (IPsec)

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 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 231.

Required Privilege Level admin—To view this statement in the configuration.
admin-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 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 227 and "Configuring the Proposals in an IPsec Policy" on page 234.

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

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 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 218.

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 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 239.

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

remote-id

Syntax

```
remote-id {
    any-remote-id;
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
}
```

Hierarchy Level [edit services ipsec-vpn ike policy *policy-name*]

Release Information Statement introduced before JUNOS Release 7.4.
ipv6_addr option added in JUNOS Release 7.6.
any-remote-id option added in JUNOS 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 229.

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

rule

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 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.

Usage Guidelines See "Configuring Match Direction for IPsec Rules" on page 236.

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

rule-set

Syntax
```
rule-set rule-set-name {
    [ rule rule-names ];
}
```

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced before JUNOS 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 242.

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

services

Syntax services ipsec-vpn { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS 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 Services Configuration Guidelines" on page 213.

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 ipsec-vpn rule *rule-name* term *term-name* from]

Release Information Statement introduced before JUNOS 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 237.

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 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 219.

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

syslog

Syntax syslog;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced before JUNOS 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 239.

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

term

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 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 236.

Required Privilege Level interface—To view this statement in the configuration.

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

then

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 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 236.

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

traceoptions

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 Release 7.5.
level option added in JUNOS 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.

size *bytes*—Maximum trace file size.

Usage Guidelines See "Tracing IPsec Operations" on page 248.

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

tunnel-mtu

Syntax tunnel-mtu *bytes*;

Hierarchy Level [edit services ipsec-vpn rule *rule-name* term *term-name* then]

Release Information Statement introduced in JUNOS 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 242.

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

Related Topics mtu

Chapter 14

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 {
    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;
    }
    traceoptions {
      debug-level level;
      filter {
        protocol name;
      }
      flag flag;
      interfaces interface-name {
        debug-level level;
        flag flag;
      }
    }
  }
```

NOTE: L2TP configurations are supported only on M7i, M10i, and M120 routers.

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 System Basics Configuration Guide* or the *JUNOS Network Interfaces Configuration Guide*.

This chapter contains the following sections:

- L2TP Services Configuration Overview on page 293
- L2TP Minimum Configuration on page 294
- Configuring L2TP Tunnel Groups on page 296
- Configuring the Identifier for Logical Interfaces that Provide L2TP Services on page 300
- AS PIC Redundancy for L2TP Services on page 302
- Tracing L2TP Operations on page 303
- Examples: Configuring L2TP Services on page 304

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 System Basics Configuration Guide*.

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;
```

```
        }
    }
}
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;
  }
  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 297
- Configuring the Local Gateway Address and PIC on page 297

- Configuring Window Size for L2TP Tunnels on page 298
- Configuring Timers for L2TP Tunnels on page 298
- Hiding Attribute-Value Pairs for L2TP Tunnels on page 299
- Configuring System Logging of L2TP Tunnel Activity on page 299

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 System Basics Configuration Guide*. A profile example is included in "Examples: Configuring L2TP Services" on page 304.

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 Class of Service Configuration Guide*.

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;
```

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 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 included in inet.0. You can specify only one system logging hostname.

Table 12 on page 299 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 12: 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

Table 12: System Log Message Severity Levels *(continued)*

Severity Level	Description
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 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 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 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.

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 "$9$Vyb4ZHkPQ39mf9pORIexNdbgoZUjqP5"; # 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;
         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 487. For an example configuration, see “Examples: Configuring L2TP Services” on page 304. For information on operational mode commands, see the *JUNOS Interfaces Command Reference*.

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.

To trace L2TP operations, include the traceoptions statement at the [edit services l2tp] hierarchy level:

```
traceoptions {
    debug-level level;
    filter {
        protocol name;
    }
    flag flag;
    interfaces interface-name {
        debug-level level;
        flag flag;
    }
}
```

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
- 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;
        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/t0BEIeWFnRhrIXxbs2aJDHqf3nCP5";
            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;
    }
    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;
        }
    }
}
```

Chapter 15

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 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 299.

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

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the keepalive timer for L2TP tunnels.

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

Usage Guidelines See "Configuring Timers for L2TP Tunnels" on page 298.

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

hide-avps

Syntax hide-avps;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Hide L2TP attribute-value pairs if the secret shared between the two ends of the tunnel is known.

Default Attribute-value pairs that can be hidden are exposed, even if the secret information is known.

Usage Guidelines See "Hiding Attribute-Value Pairs for L2TP Tunnels" on page 299.

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

host

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 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 299.

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

l2tp-access-profile

Syntax l2tp-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See "Configuring Access Profiles for L2TP Tunnel Groups" on page 297.

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

local-gateway address

Syntax local-gateway address *address*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the local IP address for L2TP tunnel.

Options *address*—Local IP address.

Usage Guidelines See "Configuring the Local Gateway Address and PIC" on page 297.

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

log-prefix

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit services l2tp tunnel-group *group-name* syslog host *hostname*]

Release Information Statement introduced before JUNOS 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 299.

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

maximum-send-window

Syntax maximum-send-window *packets*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the size of the send window for L2TP tunnels, which limits the remote end's receive window size.

Options *packets*—Maximum number of packets the send window can hold at one time.
Default: 32

Usage Guidelines See "Configuring Window Size for L2TP Tunnels" on page 298.

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

ppp-access-profile

Syntax ppp-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS 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.

Options *profile-name*—Identifier for the PPP profile.

Usage Guidelines See "Configuring Access Profiles for L2TP Tunnel Groups" on page 297.

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

receive-window

Syntax receive-window *packets*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the size of the receive window for L2TP tunnels, which limits the number of packets the server processes concurrently.

Options *packets*—Maximum number of packets the receive window can hold at one time.
Default: 16

Usage Guidelines See "Configuring Window Size for L2TP Tunnels" on page 298.

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

retransmit-interval

Syntax retransmit-interval *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the maximum retransmit interval for L2TP tunnels.

Options *seconds*—Interval, in seconds, after which the server retransmits data if no acknowledgment is received.
Default: 30 seconds

Usage Guidelines See "Configuring Timers for L2TP Tunnels" on page 298.

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 sp-*fpc/pic/port*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the adaptive services interface responsible for handling L2TP processing.

Options sp-*fpc/pic/port*—AS or Multiservices PIC interface used for L2TP processing.

Usage Guidelines See "Configuring the Local Gateway Address and PIC" on page 297.

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

services

See the following sections:

- services (Hierarchy) on page 316
- services (L2TP System Logging) on page 317

services (Hierarchy)

Syntax services l2tp { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS 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 "Layer 2 Tunneling Protocol Services Configuration Guidelines" on page 291.

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

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 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 of L2TP Tunnel Activity" on page 299.

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

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 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.

Options The remaining statements are described separately.

Usage Guidelines See “Configuring System Logging of L2TP Tunnel Activity” on page 299.

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

traceoptions

Syntax

```
traceoptions {
    debug-level level;
    filter {
        protocol name;
    }
    flag flag;
    interfaces interface-name {
        debug-level level;
        flag flag;
    }
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure L2TP tracing operations. The messages are output to /var/log/l2tpd.

Options debug-level *level*—Trace level for PPP, L2TP, RADIUS, and User Datagram Protocol (UDP):

- detail—Detailed debug information.
- error—Errors only.
- packet-dump—Packet decoding information.

filter protocol *name*—Additional filter for the specified protocol:

- l2tp
- ppp
- radius
- udp

flag *flag*—Tracing operation to perform:

- all—Trace everything.
- configuration—Trace configuration events.
- protocol—Trace routing protocol events.
- routing-socket—Trace routing socket events.
- rpd—Trace routing protocol process events.

interfaces *interface-name*—Apply L2TP traceoptions to a specific adaptive services interface.

- debug-level *level*—Trace level for the interface:

- detail—Detailed debug information.
- error—Errors only.
- extensive—All PIC debug information.

- flag *flag*—Tracing operations for the interface:
  - 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.

Usage Guidelines See “Tracing L2TP Operations” on page 303.

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

tunnel-group

Syntax

```
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;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the L2TP tunnel properties.

Options *group-name*—Identifier for the tunnel group.

The remaining statements are explained separately.

Usage Guidelines See “Configuring L2TP Tunnel Groups” on page 296.

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

tunnel-timeout

Syntax tunnel-timeout *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before JUNOS 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

Usage Guidelines See "Configuring Timers for L2TP Tunnels" on page 298.

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

Chapter 16

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 977. 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 Feature Guide*.

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 324
- Configuring LSQ Interface Redundancy Across Multiple Routers Using SONET APS on page 326
- Configuring LSQ Interface Redundancy in a Single Router Using SONET APS on page 328
- Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces on page 328
- Configuring CoS Scheduling Queues on Logical LSQ Interfaces on page 337
- Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces on page 341
- Reserving Bundle Bandwidth for Link-Layer Overhead on LSQ Interfaces on page 343
- Configuring Multiclass MLPPP on LSQ Interfaces on page 343
- Oversubscribing Interface Bandwidth on LSQ Interfaces on page 344
- Configuring Guaranteed Minimum Rate on LSQ Interfaces on page 348

- Configuring Link Services and CoS on Services PICs on page 351
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on page 354
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 359
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 365
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 369
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 on page 377
- Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 378
- Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 379
- Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 381

Layer 2 Service Package Capabilities and Interfaces

As described in "Enabling Service Packages" on page 35, 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—The section "Configuring CoS Scheduling Queues on Logical LSQ Interfaces" on page 337 describes how the JUNOS CoS components work on link services IQ (lsq) interfaces. For detailed information about JUNOS CoS components, see the *JUNOS 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 398 and "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 341.

- 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.
- 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 Interfaces Configuration Guidelines" on page 1087.

NOTE: Interface type sp is created because it is needed by the JUNOS Software. 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 337.

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 Network Interfaces Configuration Guide*.

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 Network Interfaces Configuration Guide*.

The following sections describe how to configure failover properties:

- Configuring the Association between LSQ and SONET Interfaces on page 326
- Configuring SONET APS Interoperability with Cisco Systems FRF.16 on page 327
- Restrictions on APS Redundancy for LSQ Interfaces on page 328

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.

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.

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 326. 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 System Basics Configuration Guide*.

Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces

You can configure failure recovery on M Series and T Series routers that have multiple AS or Multiservices PICs with lsq- interfaces by specifying a virtual 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 329
- Restrictions on Redundant LSQ Interfaces on page 330
- Configuring Link State Replication for Redundant Link PICs on page 331
- Examples: Configuring Redundant LSQ Interfaces for Failure Recovery on page 333

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 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.

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.

NOTE: The hot-standby option is supported only with MLPPP and CRTP configurations. The warm-standby option is also supported with LSQ FRF.15 and FRF.16 bundle configurations and with L2TP and flow monitoring services.

Certain combinations of hot-standby and warm-standby configuration are not permitted and result in a configuration error. The following examples are permitted:

- 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.
- 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.

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 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. If you configure the warm-standby option, you can include multiple instances of the primary interface.

- 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 326. 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 337.

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 Network Interfaces Configuration Guide*.

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 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 coc3-2/0/0.

```
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.

```
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 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;
               }
            }
         }
      }
   }
   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;
    }
    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;
    }
}
```

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;
    }
    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;
    }
    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 Class of Service Configuration Guide*.

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 338.
- priority—The transmit priority (low, high, strict-high); for more information, see "Configuring Scheduler Priority" on page 339.
- shaping-rate—The subscribed transmit rate; for more information, see "Configuring Scheduler Shaping Rate" on page 339.
- drop-profile-map—The random early detection (RED) drop profile; for more information, see "Configuring Drop Profiles" on page 339.

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 363. 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. 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.

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 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 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.
         ...
      }
   }
   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 982.

To configure fragmentation properties on a queue, include the fragmentation-maps statement at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            (fragment-threshold bytes | 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 343.

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 354
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16 on page 359
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using MLPPP and LFI on page 365
- Configuring LSQ Interfaces for Single Fractional T1 or E1 Interfaces Using FRF.12 on page 369
- Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.15 on page 377
- Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP on page 378
- Configuring LSQ Interfaces as T3 or OC3 Bundles Using FRF.12 on page 379
- Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP on page 381

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 985.

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 985.

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 Software 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 396.

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.

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.

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);
   ```

   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 348.

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*);

   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.

   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 Class of Service Configuration Guide*. For an example showing how the delay-buffer rates are applied, see "Example: Oversubscribing an LSQ Interface" on page 347.

   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 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 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;
   ```

   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 Class of Service Configuration Guide*.

Example: Oversubscribing an LSQ 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;
        }
    }
    interfaces {
        lsq-1/3/0 {
            unit 0 {
                output-traffic-control-profile tc_0;
            }
            unit 1 {
                output-traffic-control-profile tc_1;
            }
        }
    }
}
```

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 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.

   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 Class of Service Configuration Guide*. For an example showing how the delay-buffer rates are applied, see "Example: Configuring Guaranteed Minimum Rate" on page 351.

   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 Class of Service Configuration Guide*.

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 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 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
        }
        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 35 and "Layer 2 Service Package Capabilities and Interfaces" on page 324.

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 {
   address address;
}
```

For more information about these statements, see the "Link and Multilink Services Configuration Guidelines" on page 973.

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 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 Services Configuration Guidelines” on page 973.

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;
    }
}
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 357.

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 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;
      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 343. For more information about fragmentation maps, see “Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces” on page 341.

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 982.

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 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;
    }
}
[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;
    }
  }
  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;
    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 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 "Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using FRF.16" on page 359.

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.

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.

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 982.

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;
   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;
    }
    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 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 368.

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 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 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 341.

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 982.

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;
}
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 {
      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.

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 372.

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 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 341.

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 982.

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 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 {
    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 {
        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;
    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;
            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 369. 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 Software 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 eample 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;
   }
}
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 Software does not currently support CRTP over Frame Relay. For more information, see "Voice Services Configuration Guidelines" on page 395.

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;
    }
}
```

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 Network Interfaces Configuration Guide*.

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.

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 343.

```
[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.
        }
    }
}
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 369.

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 Network Interfaces Configuration Guide*.

NOTE: Do not configure RED profiles on ATM logical interfaces that are bundled. Drops do not occur at the ATM interface.

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;
        }
      }
    }
  }
  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;
    }
  }
```

Chapter 17

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 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 327.

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

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 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 341.

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

fragment-threshold

Syntax fragment-threshold *bytes*;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before JUNOS 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

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 341.

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

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 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 341.

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

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 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 341.

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]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring LSQ Interface Redundancy in a Single Router Using Virtual Interfaces" on page 328.

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

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 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 337.

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 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 326.

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

multilink-class

Syntax multilink-class *number*;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before JUNOS 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

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 341 and "Configuring Multiclass MLPPP on LSQ Interfaces" on page 343.

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

Related Topics 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 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 343.

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

no-fragmentation

Syntax no-fragmentation;

Hierarchy Level [edit class-of-service fragmentation-maps forwarding-class *class-name*]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 341.

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

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 Release 7.4.
Support at the [edit interfaces *interface-name* ppp-options] hierarchy level added in JUNOS 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 326.

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

per-unit-scheduler

Syntax per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For channelized OC12 IQ, channelized T3 IQ, channelized E1 IQ, E3 IQ, and Gigabit Ethernet IQ interfaces only, enable association of scheduler map names with logical interfaces.

Usage Guidelines See "Configuring Link Services and CoS on Services PICs" on page 351.

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

preserve-interface

Syntax preserve-interface;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced in JUNOS 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 331.

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

Related Topics *JUNOS Network Interfaces Configuration Guide*

primary

Syntax primary *interface-name*;

Hierarchy Level [edit interfaces rlsq*number* redundancy-options]

Release Information Statement introduced in JUNOS 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 328.

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 {
   (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 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 328.

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 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 328.

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

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 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 326.

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 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 328.

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

Chapter 18

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 Software Interfaces and Routing Configuration Guide*.

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 351.

Link services IQ interfaces use a bundle configuration. For more information, see "Link Services IQ Interfaces Configuration Guidelines" on page 323 and "Multilink and Link Services Logical Interface Configuration Overview" on page 977.

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 398 and "Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces" on page 341.

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 396
- Configuring Encapsulation for Voice Services on page 399

- Configuring Network Interfaces for Voice Services on page 399
- Examples: Configuring Voice Services on page 400

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 396
- Configuring Compression of Voice Traffic on page 397
- Configuring Delay-Sensitive Packet Interleaving on page 398
- Example: Configuring Compression of Voice Traffic on page 398

Configuring the Logical Interface Address for the MLPPP Bundle

To configure the logical address for the MLPPP bundle, include the address statement:

```
address address {
   ...
}
```

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 Network Interfaces Configuration Guide*.

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 Software 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 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;
        }
        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 Network Interfaces Configuration Guide*. 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 400.

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 400
- Configuring the Compression Interface with PPP Encapsulation on page 400

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.

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 976.

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;
    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;
        }
    }
}
```

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.
          # 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;
      }
   }
}
```

Chapter 19

Summary of Voice Services Configuration Statements

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

address

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 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 396; for a general discussion of address statement options, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics
JUNOS Network Interfaces Configuration Guide for other statements that do not affect services interfaces.

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 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 400.

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 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 397.

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

compression-device

Syntax compression-device *interface-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 in JUNOS Release 7.5.

Description Specify the compression interface for voice services traffic.

Usage Guidelines See "Configuring the Compression Interface with PPP Encapsulation" on page 400.

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 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 399; for information about encapsulation statement options used with other interface types, see the *JUNOS Network Interfaces Configuration Guide*.

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

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 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 397.

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

family

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 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 399; for a general discussion of family statement options, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

fragment-threshold

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 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 398.

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 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 Network Interfaces Configuration Guide*.

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

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 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 releases that base the RTP context value on link speed.

Usage Guidelines See "Configuring Compression of Voice Traffic" on page 397.

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 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 397.

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

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 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 397.

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 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 397.

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

unit

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 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 396; for a general discussion of logical interface properties, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

Chapter 20

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 {
        sip-text {
            dscp (alias | bits);
            forwarding-class class-name;
        }
        sip-video {
            dscp (alias | bits);
            forwarding-class class-name;
        }
        sip-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 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;
                }
            }
        }
    }
```

```
    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 416
- Configuring CoS Rules on page 416
- Configuring CoS Rule Sets on page 421
- Examples: Configuring CoS on Services Interfaces on page 422

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.

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 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 417
- Configuring Match Conditions In CoS Rules on page 418
- Configuring Actions in CoS Rules on page 419
- Example: Configuring CoS Rules on page 421

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:

```
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 440.

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 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 *JUNOS Policy Framework 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 112.

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.

You can also include application protocol definitions you have configured at the [edit applications] hierarchy level; for more information, see "Applications Configuration Guidelines" on page 59.

- 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 422 or the *JUNOS 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.

For information about some additional CoS actions, see the following sections:

- Configuring Application Profiles for Use as CoS Rule Actions on page 420
- Configuring Reflexive and Reverse CoS Rule Actions on page 420

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 {
   sip-text {
      dscp (alias | bits);
      forwarding-class class-name;
   }
   sip-video {
      dscp (alias | bits);
      forwarding-class class-name;
   }
   sip-voice {
      dscp (alias | bits);
      forwarding-class class-name;
   }
}
```

The application-profile statement includes three fixed components, sip-text, sip-video, and sip-voice. You can set the appropriate dscp and forwarding-class values for each component within the application profile.

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.

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;
            }
        }
    }
}
```

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 leve 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 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;
    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;
  }
}
```

Chapter 21

Summary of Class-of-Service Configuration Statements

The following sections explain each of the class-of-service (CoS) statements. The statements are organized alphabetically.

application-profile

Syntax

```
application-profile profile-name {
    sip-text {
        dscp (alias | bits);
        forwarding-class class-name;
    }
    sip-video {
        dscp (alias | bits);
        forwarding-class class-name;
    }
    sip-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 Release 8.1.

Description Define a CoS application profile.

Options *profile-name*—Identifier for the application profile.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Application Profiles for Use as CoS Rule Actions" on page 420.

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

application-sets

Syntax applications-sets *set-name*;

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 8.1.

Description Define one or more target application sets.

Options *set-name*—Name of the target application set.

Usage Guidelines See "Configuring Match Conditions In CoS Rules" on page 418.

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

applications

Syntax applications [*application-name*];

Hierarchy Level [edit services cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 8.1.

Description Define one or more applications to which the CoS services apply.

Options *application-name*—Name of the target application.

Usage Guidelines See "Configuring Match Conditions In CoS Rules" on page 418.

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

destination-address

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 Release 8.1.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.

Description Specify the destination address for rule matching.

Options *address*—Destination IPv4 or IPv6 address or prefix value.

Usage Guidelines See "Configuring Match Conditions In CoS Rules" on page 418.

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

destination-prefix-list

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 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 418.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

dscp

Syntax dscp (*alias* | *bits*);

Hierarchy Level [edit services cos application-profile *profile-name* (sip-text | sip-video | sip-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 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 419.

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

forwarding-class

Syntax forwarding-class *class-name*;

Hierarchy Level [edit services cos application-profile *profile-name* (sip-text | sip-video | sip-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 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 419.

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

from

Syntax
```
from {
    application-sets set-name;
    applications [ application-names ];
    destination-address 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 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 *JUNOS Policy Framework Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Configuring CoS Rules" on page 416.

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 cos rule *rule-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring CoS Rules" on page 416.

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

(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 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.

Usage Guidelines See “Configuring Reflexive and Reverse CoS Rule Actions” on page 420.

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

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;
            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 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 416.

Required Privilege Level

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

rule-set

Syntax
```
rule-set rule-set-name {
    [ rule rule-name ];
}
```

Hierarchy Level [edit services cos]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring CoS Rule Sets" on page 421.

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

services

Syntax services cos { ... }

Hierarchy Level [edit]

Release Information Statement introduced in JUNOS 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 Configuration Guidelines" on page 415.

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

sip-text

Syntax
```
sip-text {
    dscp (alias | bits);
    forwarding-class class-name;
}
```

Hierarchy Level [edit services cos application-profile *profile-name*]

Release Information Statement introduced in JUNOS Release 8.1.

Description Enable a predefined application profile for handling text data packets.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Application Profiles for Use as CoS Rule Actions" on page 420.

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

sip-video

Syntax
```
sip-video {
    dscp (alias | bits);
    forwarding-class class-name;
}
```

Hierarchy Level [edit services cos application-profile *profile-name*]

Release Information Statement introduced in JUNOS Release 8.1.

Description Enable a predefined application profile for handling video data packets.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Application Profiles for Use as CoS Rule Actions" on page 420.

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

sip-voice

Syntax
```
sip-voice {
    dscp (alias | bits);
    forwarding-class class-name;
}
```

Hierarchy Level [edit services cos application-profile *profile-name*]

Release Information Statement introduced in JUNOS Release 8.1.

Description Enable a predefined application profile for handling voice data packets.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Application Profiles for Use as CoS Rule Actions" on page 420.

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 cos rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 8.1.
address option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.

Description Source address for rule matching.

Options *address*—Source IPv4 or IPv6 address or prefix value.

Usage Guidelines See "Configuring Match Conditions In CoS Rules" on page 418.

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

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 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 CoS Rules" on page 418.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

syslog

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 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.

Usage Guidelines See "Configuring Actions in CoS Rules" on page 419.

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

term

Syntax

```
term term-name {
   from {
      application-sets set-name;
      applications [ application-names ];
      destination-address 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 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 416.

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

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 Release 8.1.

Description Define the CoS term actions.

The remaining statements are explained separately.

Usage Guidelines See "Configuring Actions in CoS Rules" on page 419.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

Chapter 22

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);
    (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;
        }
    }
}
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 440
- Configuring Service Rules on page 444
- Configuring IPsec Service Sets on page 445
- Configuring Service Set Limitations on page 450
- Configuring System Logging for Service Sets on page 450
- Enabling Services PICs to Accept Multicast Traffic on page 452
- Tracing Services PIC Operations on page 452
- Example: Configuring Service Sets on page 455

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 440
- Configuring Next-Hop Service Sets on page 441
- Determining Traffic Direction on page 442

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 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.

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 455.

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).

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:

```
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.

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 "Intrusion Detection Service Configuration Guidelines" on page 179.
- You configure IP Security (IPsec) rules at the [edit services ipsec-vpn] hierarchy level; for more information, see "IPsec Services Configuration Guidelines" on page 213.
- You configure Network Address Translation (NAT) rules at the [edit services nat] hierarchy level; for more information, see "Network Address Translation Services Configuration Guidelines" on page 129.
- You configure Packet Gateway Control Protocol (PGCP) rules at the [edit services pgcp] hierarchy level; for more information, see "PGCP Configuration Guidelines for the BGF Feature" on page 511.
- You configure stateful firewall rules at the [edit services stateful-firewall] hierarchy level; for more information, see "Stateful Firewall Services Configuration Guidelines" on page 107.

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);
([ 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 Software 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 "Dynamic Application Awareness for JUNOS Software" on page 685.

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 445
- Configuring IKE Access Profiles for IPsec Service Sets on page 446
- Configuring Certification Authorities for IPsec Service Sets on page 447
- Configuring or Disabling Antireplay Service on page 447
- Clearing the Don't-Fragment Bit on page 448
- Configuring Passive-Mode Tunneling on page 449
- Configuring the Tunnel MTU Value on page 449

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 235.

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:

```
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 System Basics Configuration Guide*. For more information about IPsec digital certificate configuration, see "Configuring IPsec Rules" on page 235.

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 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.

```
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 249.

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 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 "Intrusion Detection Service Configuration Guidelines" on page 179.

- 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 System Basics Configuration Guide*.

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 {
        services severity-level;
        facility-override facility-name;
```

```
        log-prefix prefix-value;
    }
}
```

Configure the host statement with a hostname that specifies the system log target server. The hostname local directs system log messages to the Routing Engine. The hostname must be included in inet.0. You can specify only one system log host.

Table 13 on page 451 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 13: 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 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 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 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" on page 145.

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.

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 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 453
- Configuring the Number and Size of Adaptive Services Log Files on page 453
- Configuring Access to the Log File on page 453
- Configuring a Regular Expression for Lines to Be Logged on page 454
- Configuring the Trace Operations on page 454

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;
```

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:

```
flag {
    all;
    configuration;
    routing-protocol;
    routing-socket;
    snmp;
}
```

Table 14 on page 454 describes the meaning of the adaptive services tracing flags.

Table 14: 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

Table 14: Adaptive Services Tracing Flags *(continued)*

Flag	Description	Default Setting
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 {
                service-set my-output-service-set;
            }
        }
    }
}
```

Chapter 23

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 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 452.

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

allow-multicast

Syntax allow-multicast;

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced in JUNOS 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 452.

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

anti-replay-window-size

Syntax anti-replay-window-size *bits*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in JUNOS 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 445 or "Configuring or Disabling IPSec Anti-Replay" on page 241.

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

clear-dont-fragment-bit

Syntax clear-dont-fragment-bit;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in JUNOS 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 445 or "Configuring Actions in IPsec Rules" on page 239.

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

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 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 450.

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

host

Syntax
```
host hostname {
    facility-override facility-name;
    interface-service prefix-value;
    services severity-level;
}
```

Hierarchy Level [edit services service-set *service-set-name* syslog]

Release Information Statement introduced before JUNOS 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 450.

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 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 444.

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

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 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 242 or "Configuring IPsec Service Sets" on page 445.

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 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 440.

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

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 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 444.

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 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 444.

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

local-gateway

Syntax local-gateway *address*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced before JUNOS 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 444.

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

log-prefix

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit services service-set *service-set-name* syslog host *hostname*]

Release Information Statement introduced before JUNOS 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 450.

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

logging

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 Release 8.0.

Description Define global services properties.

Options The remaining statement is explained separately.

Usage Guidelines See "Tracing Services PIC Operations" on page 452.

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 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 450.

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 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 444.

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

next-hop-service

Syntax

```
next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
    service-interface-pool name;
}
```

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced before JUNOS Release 7.4.
service-interface-pool option added in JUNOS 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.

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.

Usage Guidelines See "Configuring Service Sets to be Applied to Services Interfaces" on page 440.

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

no-anti-replay

Syntax no-anti-replay;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in JUNOS 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 445 or "Configuring or Disabling IPSec Anti-Replay" on page 241.

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 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 445.

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

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 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 440.

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 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 440.

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

service-set

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;
    }
    (ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
    max-flows number;
    (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;
        service-interface-pool name;
    }
    (pgcp-rules rule-name | pgcp-rule-sets rule-set-name);
    (stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);
    syslog {
        host hostname {
            facility-override facility-name;
            log-prefix prefix-value;
            services severity-level;
        }
    }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before JUNOS Release 7.4. The pgcp-rules and pgcp-rule-sets options were added in Release 8.4.

Description Define the service set.

Options *service-set-name*—Identifies the service set.

The remaining statements are explained separately.

Usage Guidelines See “Service Set Configuration Guidelines” on page 439.

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

services

See the following sections:

- services (Hierarchy) on page 472
- services (System Logging) on page 473

services (Hierarchy)

Syntax services { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the service rules to be applied to traffic.

Usage Guidelines See “Service Set Configuration Guidelines” on page 439.

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

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 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 450.

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

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 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 444.

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

syslog

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 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 483.

Options The remaining statements are described separately.

Usage Guidelines See "Configuring System Logging for Service Sets" on page 450.

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

tcp-mss

Syntax tcp-mss *number*;

Hierarchy Level [edit services service-set *service-set-name*]

Release Information Statement introduced in JUNOS 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 450.

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

traceoptions

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 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).
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 452.

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 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 445.

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

tunnel-mtu

Syntax tunnel-mtu *bytes*;

Hierarchy Level [edit services service-set *service-set-name* ipsec-vpn-options]

Release Information Statement introduced in JUNOS 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 445 or "Specifying the MTU for IPsec Tunnels" on page 242.

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

Related Topics mtu

Chapter 24

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 Software detects its presence and applies the appropriate configuration to it. For more information on the general configuration of interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

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:

```
(sp-fpc/pic/port | rspnumber) {
    services-options {
        inactivity-timeout seconds;
        open-timeout seconds;
        syslog {
            host hostname {
                services severity-level;
                facility-override facility-name;
                log-prefix prefix-value;
            }
        }
    }
}
```

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 inet {
        address address {
            ...
        }
        mtu bytes;
        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);
    }
}
```

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;
    }
}
```

This chapter contains the following sections:

- Services Interface Naming Overview on page 480
- Configuring the Address and Domain for Services Interfaces on page 481
- Configuring Default Timeout Settings for Services Interfaces on page 482
- Configuring System Logging for Services Interfaces on page 483
- Enabling Fragmentation on GRE Tunnels on page 484
- Applying Filters and Services to Interfaces on page 485
- Configuring AS or Multiservices PIC Redundancy on page 487
- Examples: Configuring Services Interfaces on page 490

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.
- rlsq—Redundancy LSQ interface.
- rsp—Redundancy adaptive services interface.
- 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 Network Interfaces Configuration Guide*.

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 440.

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.
- 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 60.

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 86,400 seconds. Any value you configure in the intrusion detection service (IDS) definition overrides the value specified here; for more information, see "Intrusion Detection Service Configuration Guidelines" on page 179.

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 450.

To configure interface-wide default system logging values, include the syslog statement at the [edit interfaces *interface-name* services-options] hierarchy level:

```
syslog {
    host hostname {
        services severity-level;
        facility-override facility-name;
        log-prefix prefix-value;
    }
}
```

Configure the host statement with a hostname 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 included in inet.0. You can specify only one system logging hostname.

Table 15 on page 483 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 15: 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

Table 15: System Log Message Severity Levels *(continued)*

Severity Level	Description
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 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:

```
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:

```
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;
            ...
        }
    }
}
```

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.

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

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 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 440.

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:

- 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 *JUNOS Policy Framework 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 490.

For more information on applying filters to interfaces, see the *JUNOS Network Interfaces Configuration Guide*. For general information on filters, see the *JUNOS Policy Framework 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.

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 System Basics and Services Command Reference*.
- The driver watchdog timer expires.
- The request interface switchover command is issued. For more information, see the *JUNOS Interfaces Command Reference*.

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 rsp number]
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 328.

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 837.

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 222.

- 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 490.

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;
            }
        }
    }
}
```

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;
}
```

Chapter 25

Summary of Service Interface Configuration Statements

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

address

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 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 481; for a general discussion of address statement options, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics
JUNOS Network Interfaces Configuration Guide for other statements that do not affect services interfaces.

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 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.

Usage Guidelines See "Enabling Fragmentation on GRE Tunnels" on page 484.

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

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 logical-systems *logical-system-name* interfaces sp-*fpc/pic/port* unit *logical-unit-number*]

Release Information

Statement introduced before JUNOS Release 7.4.

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

ipsec-interface-id *name*—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 | dedicated)—Specifies whether a logical interface can host one (dedicated) or multiple (shared) sessions at one time.

Usage Guidelines

See "Configuring the Identifier for Logical Interfaces that Provide L2TP Services" on page 300 (for L2TP configurations) or "Configuring Dynamic Endpoints for IPsec Tunnels" on page 242 (for IPsec configurations).

Required Privilege Level

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

facility-override

Syntax facility-override *facility-name*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before JUNOS 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 483.

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

family

Syntax

```
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> ];
        }
    }
}
```

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 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 481 or; for a general discussion of family statement options, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

host

Syntax
```
host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
}
```

Hierarchy Level [edit interfaces *interface-name* services-options syslog]

Release Information Statement introduced before JUNOS 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 485.

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 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 482.

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

input

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 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 485.

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 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 Network Interfaces Configuration Guide*.

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

log-prefix

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before JUNOS 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 483.

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

open-timeout

Syntax open-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before JUNOS 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 86,400 seconds

Usage Guidelines See "Configuring Default Timeout Settings for Services Interfaces" on page 482.

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 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 485.

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

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 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 485.

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

primary

Syntax primary *interface-name*;

Hierarchy Level [edit interfaces (rsp0 | rsp1) redundancy-options]

Release Information Statement introduced before JUNOS 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 487.

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 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 487.

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 (rsp0 | rsp1) redundancy-options]

Release Information Statement introduced before JUNOS 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 487.

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 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 485.

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

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 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 481.

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

service-filter

Syntax service-filter *filter-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 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, the router software assumes that 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.

Usage Guidelines See "Applying Filters and Services to Interfaces" on page 485.

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

service-set

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 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 485.

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

services

Syntax services *severity-level*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before JUNOS 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 483.

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

services-options

Syntax

```
services-options {
   inactivity-timeout seconds;
   open-timeout seconds;
   syslog {
      host hostname {
         facility-override facility-name;
         log-prefix prefix-value;
         services severity-level;
      }
   }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the service options to be applied on an interface.

Options The remaining statements are explained separately.

Usage Guidelines See “Service Interface Configuration Guidelines” on page 479.

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

syslog

Syntax

```
syslog {
    host hostname {
        services severity-level;
        facility-override facility-name;
        log-prefix prefix-value;
    }
}
```

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before JUNOS 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 483.

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

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 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 "Service Interface Configuration Guidelines" on page 479; for a general discussion of logical interface properties, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

Chapter 26

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;
                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;
            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;
      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 (All Streams) {
            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 {
    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 source-port;
        }
        disable-session-mirroring;
    }
}
media-service media-service-name {
    nat-pool nat-pool-name;
}
rule rule-name {
    gateway gateway-name;
    media-service media-service-name;
}
rule-set rule-set-name {
    rule rule-name1;
    rule rule-name2;
    rule rule-name3;
}
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 source-port;
    }
    disable-session-mirroring;
}
traceoptions {
    file <filename filename> <files number> <match regexg> <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 number {
        media-service media-service-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 *JUNOS Multiplay Solutions Guide*.

Chapter 27

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.

administrative

See the following sections:

- administrative (Control Association) on page 518
- administrative (Virtual Interface) on page 519

administrative (Control Association)

Syntax administrative (forced-905 | 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 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 Topics *JUNOS Multiplay Solutions Guide*

administrative (Virtual Interface)

Syntax administrative (forced-905 | forced-906 | 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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

cancel-graceful

See the following sections:

- cancel-graceful (Control Association) on page 523
- cancel-graceful (Virtual Interface) on page 524

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

cleanup-timeout

Syntax cleanup-timeout *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

data-inactivity-detection

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 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 Topics *JUNOS Multiplay Solutions Guide*

default

Syntax default *trace-level*;

Hierarchy Level [edit services pgcp gateway *gateway-name* traceoptions flag]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 source-port;
}
```

Hierarchy Level
[edit services pgcp session-mirroring]
[edit services pgcp gateway *gateway-name* session-mirroring]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

destination-address

Syntax
```
destination-address destination-address;
```

Hierarchy Level
[edit services pgcp session-mirroring delivery-function *delivery-function-name*]
[edit services pgcp gateway *gateway-name* session-mirroring delivery-function *delivery-function-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

destination-port

Syntax destination-port *destination-port*;

Hierarchy Level [edit services pgcp session-mirroring delivery-function *delivery-function-name*]
[edit services pgcp gateway *gateway-name* session-mirroring delivery-function *delivery-function-name*]

Release Information Statement introduced in JUNOS 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 65535

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 Topics *JUNOS Multiplay Solutions Guide*

detect

Syntax detect;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties inactivity-timer inactivity-timeout]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

disable-session-mirroring

Syntax
```
disable-session-mirroring;
```

Hierarchy Level [edit services pgcp session-mirroring]
[edit services pgcp gateway *gateway-name* session-mirroring]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

dscp

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

file

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 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 Topics *JUNOS Multiplay Solutions Guide*

flag

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 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 Topics *JUNOS Multiplay Solutions Guide*

gateway

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;
        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 {
            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 {
    audit-observed-events-returns;
    encoding {
        no-dscp-bit-mirroring;
        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 (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 source-port;
            }
            disable-session-mirroring;
        }
    }
}
```

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in JUNOS Release 8.4.
graceful-restart option introduced in JUNOS Release 8.5.
h248-options option introduced in JUNOS Release 8.5.
h248-properties option introduced in JUNOS Release 8.5.
monitor option introduced in JUNOS Release 9.0.
session-mirroring option introduced in JUNOS Release 9.2.
data-inactivity-detection option introduced in JUNOS Release 9.3.
overload-control option introduced in JUNOS Release 9.3.
platform option introduced in JUNOS Release 9.6.
h248–profile option introduced in JUNOS Release 10.0.
ipsec-transport-security-association option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

gateway-address

Syntax gateway-address *gateway-address*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.4.
local-controller option introduced in JUNOS Release 9.4.
remote-controller option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

gateway-port

Syntax gateway-port *gateway-port*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

graceful

See the following sections:

- graceful (Control Association) on page 548
- graceful (Virtual Interface) on page 549

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

graceful-restart

Syntax
```
graceful-restart {
    maximum-synchronization-mismatches seconds;
    seconds;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

h248-options

Syntax

```
h248-options {
    audit-observed-events-returns;
    encoding {
        no-dscp-bit-mirroring;
        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;
    }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS Release 8.5.
audit-observed-events-returns option introduced in JUNOS Release 9.3.
encoding option introduced in JUNOS Release 9.3.
service-change options introduced in JUNOS Release 9.3.

use-lower-case option introduced in JUNOS Release 9.5.
h248-profile option introduced in JUNOS Release 10.0.

Description Configure options that affect virtual BGF H.248 behavior.

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 Topics *JUNOS Multiplay Solutions Guide*

h248-profile

Syntax
```
h248-profile {
    profile-name profile-name;
    profile-version version-number;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options]

Release Information Statement introduced in JUNOS 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 Topics Configuring the H.248 Profile

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;
        }
        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;
            }
        }
    }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS Release 8.5.
diffserv option introduced in JUNOS Release 9.0.
inactivity-timer option introduced in JUNOS Release 9.2.
traffic-management option introduced in JUNOS Release 9.2.
application-data-inactivity-detection option introduced in JUNOS Release 9.3.
event-timestamp-notification option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 H.248 control association.the.

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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

inactivity-delay

Syntax inactivity-delay *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

inactivity-duration

Syntax inactivity-duration *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* data-inactivity-detection]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

ipsec-transport-security-association

Syntax ipsec-transport-security-association *security-association-name*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics
- Protecting H.248 Messages and Mirrored Sessions with IPsec Overview
- Configuring IPsec for H.248 Messages in Transport Mode

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 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 Topics *JUNOS Multiplay Solutions Guide*

max-burst-size

See the following sections:

- max-burst-size (All Streams) on page 562
- max-burst-size (RTCP Streams) on page 563

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 Release 9.2.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

max-burst-size (RTCP Streams)

Syntax

```
max-burst-size {
    rtcp {
        (fixed-value bytes | percentage percentage);
    }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties traffic-management]

Release Information Statement introduced in JUNOS Release 9.2.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

max-concurrent-calls

Syntax max-concurrent-calls *number-of-calls*;

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 9.2.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

maximum-waiting-delay

Syntax maximum-waiting-delay *milliseconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-timers]

Release Information Statement introduced in JUNOS 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.

Related Topics *JUNOS Multiplay Solutions Guide*

media

Syntax
```
media {
    rtcp;
    rtp;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name* monitor]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

media-service

Syntax
```
media-service media-service-name {
    nat-pool nat-pool-name;
}
```

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in JUNOS Release 8.4.

Description Configure media services for the BGF configuration. Media services are applied to PGCP packets.

Options *media-service-name*—Identifier for the media service name.

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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 30,000

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 30,000

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 30,000

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

mgc-provisional-response-timer-value

Syntax
```
mgc-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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 3,000

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 3,000

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 3,000

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 30,000

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 30,000

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 30,000

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

Related Topics *JUNOS Multiplay Solutions Guide*

monitor

Syntax
```
monitor {
    media {
        rtcp;
        rtp;
    }
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

nat-pool

Syntax nat-pool *nat-pool-name*;

Hierarchy Level [edit services pgcp media-service *media-service-name*]

Release Information Statement introduced in JUNOS Release 8.4.

Description Specify a Network Address Translation (NAT) pool for the media service. You can specify only one pool for each media service.

Options *nat-pool-name*—Identifier for the NAT address pool.

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 *gateway-name* session-mirroring delivery-function *deliver-function-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 29,000

maximum *milliseconds*—Maximum interval within which the gateway controller waits for a response to transactions from the virtual BGF.
Range: 500 through 29,000

minimum *milliseconds*—Minimum interval within which the gateway controller waits for a response to transactions from the virtual BGF.
Range: 500 through 29,000

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.5.
maximum and minimum options introduced in JUNOS 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 29,000

maximum *milliseconds*—Maximum interval within which the virtual BGF waits for a response to transactions from the gateway controller.
Range: 500 through 29,000

minimum *milliseconds*—Minimum interval within which the virtual BGF waits for a response to transactions from the gateway controller.
Range: 500 through 29,000

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

notification-rate-limit

Syntax notification-rate-limit *rate*;

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 9.3.
reject-all-commands-threshold and reject-new-calls-threshold options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

peak-data-rate

See the following sections:

- peak-data-rate (All Streams) on page 584
- peak-data-rate (RTCP) on page 585

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 Release 9.2.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

platform

Syntax
```
platform {
    device interface-name;
    routing-engine;
}
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics Configuring the H.248 Profile

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 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 Topics Configuring the H.248 Profile

queue-limit-percentage

Syntax queue-limit-percentage *percentage*;

Hierarchy Level [edit services pgcp gateway *gateway-name* overload-control]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

reject-all-commands-threshold

Syntax reject-all-commands-threshold *percentage*;

Hierarchy Level [edit services pgcp gateway *gateway-name* overload-control]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 8.4.
service-interface option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

rtcp

Syntax rtcp;

Hierarchy Level [edit services pgcp gateway *gateway-name* monitor media]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

rtp

Syntax rtp;

Hierarchy Level [edit services pgcp gateway *gateway-name* monitor media]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

rule

Syntax
```
rule rule-name {
    gateway gateway-name;
    media-service media-service-name;
}
```

Hierarchy Level
[edit services pgcp],
[edit services service-set *service-set-name*]

Release Information Statement introduced in JUNOS Release 8.4.

Description Specify the rule that the router uses when it applies the media service.

Options *rule-name*—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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

sbc-utils

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 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.

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics inactivity-delay

latch-deadlock-delay

JUNOS Multiplay Solutions Guide

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

service-state

See the following sections:

- service-state (Virtual BGF) on page 598
- service-state (Virtual Interface) on page 599

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

services

Syntax services pgcp { ... }

Hierarchy Level [edit]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 source-port;
    }
    disable-session-mirroring;
}
```

Hierarchy Level [edit services pgcp]
[edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

source-address

Syntax source-address *source-address*;

Hierarchy Level [edit services pgcp session-mirroring delivery-function *deliver-function-name*]
[edit services pgcp gateway *gateway-name* session-mirroring delivery-function *deliver-function-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

source-port

Syntax source-port *source-port*;

Hierarchy Level [edit services pgcp session-mirroring delivery-function *deliver-function-name*]
[edit services pgcp gateway *gateway-name* session-mirroring delivery-function *deliver-function-name*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

sustained-data-rate

See the following sections:

- sustained-data-rate (All Streams) on page 603
- sustained-data-rate (RTCP Streams) on page 604

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 Release 9.2.
maximum and minimum options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

timerx

Syntax timerx *seconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties hanging-termination-detection]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

tmax-retransmission-delay

Syntax tmax-retransmission-delay *milliseconds*;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-timers]

Release Information Statement introduced in JUNOS 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.

Related Topics *JUNOS Multiplay Solutions Guide*

traceoptions

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 Release 8.4.
Statement extensively revised in JUNOS 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.

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

use-lower-case

Syntax use-lower-case;

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-options]

Release Information Statement introduced in JUNOS 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

virtual-interface

Syntax
```
virtual-interface number {
   media-service media-service-name;
   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 Release 8.4.
service-state option introduced in JUNOS Release 9.0.
service-interface option introduced in JUNOS Release 9.3.

Description Configure a virtual interface for the BGF.

Options *number*—Identifier for the interface.
Range: 0 through 1023

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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

Chapter 28

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 613

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;
}
```

Chapter 29

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 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 Topics *JUNOS Multiplay Solutions Guide*

pool

Syntax
```
pool pool-name {
    interface interface-name.unit-number;
}
```

Hierarchy Level [edit services service-interface-pools]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

Chapter 30

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;
            }
            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-policies [ policy-and-policy-set-names ];
        new-transaction-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 {
                        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 [ contact-fields ];
                method {
                    method-invite;
                    method-message;
                    method-options;
                    method-publish;
                    method-refer;
                    method-register;
                    method-subscribe;
                }
                request-uri [ uri-fields ];
                source-address [ ip-addresses ];
            }
            then {
                (accept | reject);
                admission-control admission-control-profile;
                message-manipulation {
                    forward-manipulation {
                        manipulation-rule-name;
                    }
                    reverse-manipulation {
                        manipulation-rule-name;
                    }
                }
                route {
                    egress-service-point service-point-name;
                    next-hop (request-uri | address ipv4-address <port port-number>
                        <transport-protocol (udp|tcp)>);
                }
                trace;
            }
        }
    }
    new-transaction-policy-set policy-set-name {
        policy-name [ policy-name ];
    }
    timers {
        inactive-call seconds;
        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;
                    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 *JUNOS Multiplay Solutions Guide*.

Chapter 31

Summary of Border Signaling Gateway Configuration Statements

The following sections explain each of the border signaling gateway statements. The statements are organized alphabetically.

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics Configuring DNS Resolution for Locating SIP Servers

admission-control

See the following sections:

- admission-control (Border Signaling Gateway) on page 623
- admission-control (New Transaction Policy) on page 624

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics Configuring DNS Resolution for Locating SIP Servers

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

data-inactivity-detection

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 Release 9.6.

Description Configure a default value for media-realm for each new call. The BGF uses media-realm to locate a virtual interface with the same value to determine the media-service 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

dscp

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 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 Topics *JUNOS Multiplay Solutions Guide*

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]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

file

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 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.

world-readable | no-world-readable—Allow all users to use the log file or disallow all users from using the log file.
Default: 1048576

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

Related Topics *JUNOS Multiplay Solutions Guide*

flag

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 Release 9.4.

Description Configure trace options for components of the BSG.

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 Topics *JUNOS Multiplay Solutions Guide*

forward-manipulation

Syntax
```
forward-manipulation {
    manipulation-rule-name;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policyterm then message-manipulation]

Release Information Statement introduced in JUNOS 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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Required Privilege Level interface—To view this statement in the configuration.

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

Related Topics *JUNOS Multiplay Solutions Guide*

from

See the following sections:

- from (New Call Usage Policy) on page 638
- from (New Transaction Policy) on page 639
- from (Service Class) on page 640

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 Release 9.4.
regular-expression options introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

from (New Transaction Policy)

Syntax

```
from {
    contact {
        regular-expression [ regular-expression ];
    }
    method {
        method-invite;
        method-message;
        method-options;
        method-publish;
        method-refer;
        method-register;
        method-subscribe;
    }
    request-uri {
        regular-expression [ regular-expression ];
    }
    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 Release 9.4.
regular-expression options introduced in JUNOS 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.

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—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.

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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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-policies [ policy-and-policy-set-names ];
      new-transaction-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 {
                    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 [ contact-fields ];
            method {
                method-invite;
                method-message;
                method-options;
                method-publish;
                method-refer;
                method-register;
```

```
                method-subscribe;
            }
            request-uri [ uri-fields ];
            source-address [ ip-addresses ];
          }
          then {
            (accept | reject);
            admission-control admission-control-profile;
            message-manipulation {
              forward-manipulation {
                manipulation-rule-name;
              }
              reverse-manipulation {
                manipulation-rule-name;
              }
            }
            route {
              egress-service-point service-point-name;
              next-hop (request-uri | (address ipv4-address <port port-number>
                <transport-protocol (udp | tcp)>);
            }
            trace;
          }
        }
      }
      new-transaction-policy-set policy-set-name {
        policy-name [ policy-names ];
      }
      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;
            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 Release 9.4.
data-inactivity-detection option introduced in JUNOS Release 9.6.
message-manipulation option introduced in JUNOS Release 9.6.
message-manipulation-rules option introduced in JUNOS Release 9.6.
name-resoltuion-cache option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

inactivity-duration

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

media-policy

Syntax
```
media-policy {
    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 Release 9.4.
no-anchoring statement introduced in JUNOS Release 9.5.
service-class statement introduced in JUNOS Release 9.5.

Description Configure the service class to be applied to traffic that matches the new call usage policy.

Options 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 gateway *gateway-name* embedded-spdf] hierarchy level.

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 term then]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics 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 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 Topics Configuring DNS Resolution for Locating SIP Servers

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 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 Topics *JUNOS Multiplay Solutions Guide*

minimum

Syntax minimum *trace-level*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* traceoptions flag *flag*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics Configuring DNS Resolution for Locating SIP Servers

new-call-usage-policies

Syntax new-call-usage-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 Release 9.4.

Description Assign new call usage policies or policy sets to the service point. All the packets arriving at the service point are matched to 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 Topics *JUNOS Multiplay Solutions Guide*

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 {
                    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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

new-transaction-policies

Syntax
```
new-transaction-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 Release 9.4.

Description Assign new transaction policies or policy sets to the service point. All packets arriving at the service point are matched to these policies.

Options [*policy-and-policy-set-names*]—Names of new call usage 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 Topics *JUNOS Multiplay Solutions Guide*

new-transaction-policy

Syntax

```
new-transaction-policy policy-name {
    term term-name {
        from {
            contact [ contact-fields ];
            method {
                method-invite;
                method-message;
                method-options;
                method-publish;
                method-refer;
                method-register;
                method-subscribe;
            }
            request-uri [ uri-fields ];
            source-address [ ip-addresses ];
        }
        then {
            (accept | reject);
            admission-control admission-control-profile;
            message-manipulation {
                forward-manipulation {
                    manipulation-rule-name;
                }
                reverse-manipulation {
                    manipulation-rule-name;
                }
            }
            route {
                egress-service-point service-point-name;
                next-hop (request-uri | address ipv4-address <port port-number>
                    <transport-protocol (udp | tcp)>);
            }
            trace;
        }
    }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip]

Release Information Statement introduced in JUNOS Release 9.4.
message-manipulation option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

next-hop

Syntax next-hop (request-uri | 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]

Release Information Statement introduced in JUNOS Release 9.4.

Description Specify the SIP entity towards which SIP requests are sent.

Options request-uri—All requests and responses on the dialog are routed according to SIP. 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

reverse-manipulation

Syntax
```
reverse-manipulation {
    manipulation-rule-name;
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy term then message-manipulation]

Release Information Statement introduced in JUNOS 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 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 Topics *JUNOS Multiplay Solutions Guide*

route

Syntax

```
route {
    egress-service-point service-point-name;
    next-hop {
        (request-uri | address ipv4-address <port port-number> <transport-protocol (tcp |
            udp)>);
    }
}
```

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 Release 9.4.

Description Configure the next-hop destination and egress service point for a new transaction policy.

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 Topics *JUNOS Multiplay Solutions Guide*

sbc-utils

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 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.

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

service-interface

See the following sections:

- service-interface (Gateway) on page 663
- service-interface (Service Point) on page 663

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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-policies [ policy-and-policy-set-names ];
        new-transaction-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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

service-policies

Syntax
```
service-policies {
    new-call-usage-policies [ policy-and-policy-set-names ];
    new-transaction-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 Release 9.4.

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 Topics *JUNOS Multiplay Solutions Guide*

services

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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.
- 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 Topics *JUNOS Multiplay Solutions Guide*

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{
                    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 [ contact-fields ];
```

```
            method {
               method-invite;
               method-message;
               method-options;
               method-publish;
               method-refer;
               method-register;
               method-subscribe;
            }
            request-uri [ uri-fields ];
            source-address [ ip-addresses ];
         }
         then {
            (accept | reject);
            admission-control admission-control-profile;
            message-manipulation {
               forward-manipulation {
                  manipulation-rule-name;
               }
               reverse-manipulation {
                  manipulation-rule-name;
               }
            }
            route {
               egress-service-point service-point-name;
               next-hop (request-uri | address ipv4-address | <port port-number>
                  <transport-protocol (udp | tcp)>);
            }
            trace;
         }
      }
   }
   new-transaction-policy-set policy-set-name {
      policy-name [ policy-names ];
   }
   timers {
      inactive-callseconds;
      timer-c seconds;
   }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name*]

Release Information Statement introduced in JUNOS Release 9.4.
data-inactivity-detection option introduced in JUNOS 9.6.
message-manipulation-rules option introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

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 Topics *JUNOS Multiplay Solutions Guide*

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 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.

Related Topics *JUNOS Multiplay Solutions Guide*

term

See the following sections:

- term (New Call Usage Policy) on page 675
- term (New Transaction Policy) on page 676
- term (Service Class) on page 677

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;
         }
         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 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 Topics *JUNOS Multiplay Solutions Guide*

term (New Transaction Policy)

Syntax

```
term term-name {
    from {
        contact [ contact-fields ];
        method {
            method-invite;
            method-message;
            method-options;
            method-publish;
            method-refer;
            method-register;
            method-subscribe;
        }
        request-uri [ uri-fields ];
        source-address [ ip-addresses ];
    }
    then {
        (accept | reject);
        admission-control admission-control-profile;
        message-manipulation {
            forward-manipulation {
                manipulation-rule-name;
            }
            reverse-manipulation {
                manipulation-rule-name;
            }
        }
        route {
            next-hop (request-uri | address ipv4-address | <port port-number> |
                <transport-protocol (udp|tcp)>);
            egress-service-point service-point-name;
        }
        trace;
    }
}
```

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip new-transaction-policy *policy-name*

Release Information Statement introduced in JUNOS Release 9.4.
message-manipulation introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

then

See the following sections:

- then (New Call Usage Policy) on page 678
- then (New Transaction Policy) on page 679
- then (Service Class) on page 680

then (New Call Usage Policy)

Syntax

```
then {
    media-policy {
        data-inactivity-detection {
            inactivity-duration seconds;
        }
        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 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 Topics *JUNOS Multiplay Solutions Guide*

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;
        }
    }
    route {
        next-hop (request-uri | address ipv4-address | <port port-number> | <transport-protocol
            (udp|tcp)>);
        egress-service-point service-point-name;
    }
    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 Release 9.4.
message-manipulation option introduced in JUNOS Release 9.6.

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*.

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.

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

Related Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

timer-c

Syntax timer-c *seconds*;

Hierarchy Level [edit services border-signaling-gateway gateway *gateway-name* sip timers]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

traceoptions

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*]

Release Information Statement introduced in JUNOS 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

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 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 Topics *JUNOS Multiplay Solutions Guide*

Part 3

Dynamic Application Awareness for JUNOS Software

- Dynamic Application Awareness for JUNOS Software Overview on page 687
- Application Identification Configuration Guidelines on page 691
- Summary of Application Identification Configuration Statements on page 701
- Application-Aware Access List Configuration Guidelines on page 725
- Summary of AACL Configuration Statements on page 731
- Local Policy Decision Function Configuration Guidelines on page 743
- Summary of L-PDF Configuration Statements on page 749

Chapter 32

Dynamic Application Awareness for JUNOS Software 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).

This chapter includes the following:

- IDP Overview on page 687
- APPID Overview on page 688
- AACL Overview on page 688
- L-PDF Overview on page 689

IDP Overview

The Dynamic Application Awareness for JUNOS Software set of services adds support for Intrusion Detection and Prevention (IDP) functionality using Deep Packet Inspection (DPI) technology to Juniper Networks MX Series Ethernet Services Routers equipped with Multiservices DPCs and M120 or M320 Multiservice Edge Routers equipped with Multiservices 400 PICs. The IDP functionality is already supported on Juniper Networks J Series Services Routers and SRX Series Services Gateways running the JUNOS Software and is described in the *JUNOS Software Security Configuration Guide*. The same CLI statements and commands are used on all JUNOS platforms, but IDP is incorporated as a component of service sets only on the specified Juniper Networks M Series and MX Series routers.

NOTE: IDP depends on APPID for definition and detection of some Layer 7 applications. Therefore you must configure APPID along with IDP on the service set. Before configuring IDP policy, you must download the APPID application package. Only one service set can be applied to a single interface when APPID functionality is used.

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 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. Operational commands for monitoring and regulating IDP activity use the clear/request/show security idp command syntax.

NOTE: On 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 source port, destination port, source address, and destination address. However, for ICMP flows, the destination port is 0, so that any ICMP flow matching source port, source address, and destination address would be blocked. For more information on the ip-action statement, see the *JUNOS Software CLI Reference*.

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 M120 or M320 routers equipped with Multiservices 400 PICs. 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 Configuration Guidelines" on page 691 and "Summary of Application Identification Configuration Statements" on page 701. For more information on the operational commands, see the *JUNOS System Basics and Services Command Reference*.

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 currently supported only on MX Series routers equipped with Multiservices DPCs.

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 Configuration Guidelines" on page 725 and "Summary of AACL Configuration Statements" on page 731. For more information on the operational command, see the *JUNOS System Basics and Services Command Reference*.

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); it is currently supported only on Juniper Networks MX Series Ethernet Services routers equipped with Multiservices DPCs. 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 Configuration Guidelines" on page 439.

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 Configuration Guidelines" on page 743 and "Summary of L-PDF Configuration Statements" on page 749. For more information on the operational commands, see the *JUNOS System Basics and Services Command Reference*.

Chapter 33

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;
    max-checked-bytes bytes;
    min-checked-bytes bytes;
    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 {
```

```
        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;
    }
    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;
    }
}
```

This chapter contains the following sections:

- Defining an Application Identification on page 692
- Configuring APPID Rules on page 694
- Configuring Application Profiles on page 695
- Configuring Application Groups on page 695
- Configuring Global APPID Properties on page 696
- Configuring Automatic Download of Application Package Updates on page 697
- Tracing APPID Operations on page 697
- Examples: Configuring Application Identification Properties on page 699

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:

- 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.

For a configuration example, see “Examples: Configuring Application Identification Properties” on page 699.

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.

- 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 699.

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 694. 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 Configuration Guidelines" on page 439.

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;
    }
    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 699.

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:

```
[edit services]
application-identification {
  application-system-cache-timeout seconds;
  max-checked-bytes bytes;
  min-checked-bytes bytes;
  no-application-identification;
  no-application-system-cache;
  no-clear-application-system-cache;
  no-signature-based;
}
```

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.
- no-application-identification—Disable all application identification methods.

- 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-signature-based—Disable signature-based application identification method.

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.

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 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 698
- Configuring the Number and Size of APPID Log Files on page 698
- Configuring Access to the Log File on page 699
- Configuring a Regular Expression for Lines to Be Logged on page 699
- Configuring the Tracing Flags on page 699

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;
```

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;
}
```

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;
        junos:mgcp;
        junos:sip;
    }
}
```

Chapter 34

Summary of Application Identification Configuration Statements

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

address

Syntax

```
address address-name {
    destination {
        ip address/netmask;
        port-range {
            tcp (port | range);
            udp (port | range);
        }
    }
    source {
        ip address/netmask;
        port-range {
            tcp (port | range);
            udp (port | range);
        }
    }
    order number;
}
```

Hierarchy Level [edit services application-identification application-rule *rule-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring APPID Rules" on page 694.

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

application

See the following sections:

- application (Defining) on page 702
- application (Including in Rule) on page 703

application (Defining)

Syntax

```
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;
    }
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Defining an Application Identification" on page 692.

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

application (Including in Rule)

Syntax application *application-name*;

Hierarchy Level [edit services application-identification rule *rule-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Identify the application for inclusion in a rule.

Options *application-name*—Identifier for the application.

Usage Guidelines See "Configuring APPID Rules" on page 694.

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

application-group

Syntax
```
application-group group-name {
   application-groups {
      application-group-name;
   }
   applications {
      application-name;
   }
   index number;
   disable;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring Application Groups" on page 695.

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

application-groups

Syntax application-groups {
application-group-name;
}

Hierarchy Level [edit services application-identification application-group *group-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Identify the list of application groups for inclusion in a larger application group. The name statement is mandatory.

Options *application-group-name*—Identifier for the application group. Maximum length is 32 characters.

Usage Guidelines See “Configuring Application Groups” on page 695.

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

application-system-cache-timeout

Syntax application-system-cache-timeout *time*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Configure the lifetime for entries in the application system cache.

Options *time*— Lifetime for system cache entries, in seconds.

Usage Guidelines See “Configuring Global APPID Properties” on page 696.

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

applications

Syntax
```
applications {
application-name;
}
```

Hierarchy Level [edit services application-identification application-group *group-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring Application Groups" on page 695.

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

automatic

Syntax
```
automatic {
  interval hour;
  start-time time;
}
```

Hierarchy Level [edit services application-identification download]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring Automatic Download of Application Package Updates" on page 697.

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

destination

Syntax

```
destination {
    ip address/netmask;
    port-range {
        tcp (port | range);
        udp (port | range);
    }
}
```

Hierarchy Level [edit services application-identification rule *rule-name* address *address-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Define destination properties for application-identification rule processing.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring APPID Rules” on page 694.

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

disable

See the following sections:

- disable (APPID Application) on page 707
- disable (APPID Application Group) on page 707
- disable (APPID Port Mapping) on page 708

disable (APPID Application)

Syntax disable;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable this application definition.

Usage Guidelines See "Defining an Application Identification" on page 692.

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

disable (APPID Application Group)

Syntax disable;

Hierarchy Level [edit services application-identification application-group *group-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable application group properties.

Usage Guidelines See "Configuring Application Groups" on page 695.

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

disable (APPID Port Mapping)

Syntax disable;

Hierarchy Level [edit services application-identification application *application-name* port-mapping]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable port-mapping properties for application identification.

Usage Guidelines See “Defining an Application Identification” on page 692.

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

disable-global-timeout-override

Syntax disable-global-timeout-override;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in JUNOS Release 10.0.

Description Disallow overriding a global inactivity or session timeout.

Usage Guidelines See “Defining an Application Identification” on page 692.

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

download

Syntax
```
download {
    automatic {
        interval hour;
        start-time time;
    }
    url url;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Define application download properties.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring Automatic Download of Application Package Updates" on page 697.

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

idle-timeout

Syntax idle-timeout *seconds*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in JUNOS 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: 60
Range: 1 through 604,800

Usage Guidelines See "Defining an Application Identification" on page 692.

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

inactivity-non-tcp-timeout

Syntax inactivity-non-tcp-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in JUNOS 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

Usage Guidelines See "Defining an Application Identification" on page 692.

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

inactivity-tcp-timeout

Syntax inactivity-tcp-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in JUNOS Release 10.0.

Description Define the inactivity timeout period for TCP established sessions in seconds.

Options *seconds*—Timeout period.
Range: 4 through 86,400

Usage Guidelines See "Defining an Application Identification" on page 692.

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

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 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

Usage Guidelines See "Defining an Application Identification" on page 692 and "Configuring Application Groups" on page 695.

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

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 Release 9.5.

Description Define an IP address and netmask for identifying the traffic destination or source.

Options *address/netmask*—IP address and netmask.

Usage Guidelines See "Configuring APPID Rules" on page 694.

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

max-checked-bytes

Syntax max-checked-bytes *bytes*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Specify the maximum number of bytes to be inspected.

Options *bytes*—Maximum number of bytes.
Range: 0 through 100,000

Usage Guidelines See "Configuring Global APPID Properties" on page 696.

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

min-checked-bytes

Syntax min-checked-bytes *bytes*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Specify the minimum number of bytes to be inspected.

Options *bytes*—Minimum number of bytes.
Range: 0 through 2000

Usage Guidelines See "Configuring Global APPID Properties" on page 696.

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

no-application-identification

Syntax no-application-identification;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable all application identification methods.

Usage Guidelines See “Configuring Global APPID Properties” on page 696.

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

no-application-system-cache

Syntax no-application-system-cache;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable storing application identification results in the application system cache.

Usage Guidelines See “Configuring Global APPID Properties” on page 696.

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

no-clear-application-system-cache

Syntax no-clear-application-system-cache;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable clearing the application system cache.

Usage Guidelines See “Configuring Global APPID Properties” on page 696.

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

no-signature-based

Syntax no-signature-based;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS Release 9.5.

Description Disable the signature-based application identification method.

Usage Guidelines See “Configuring Global APPID Properties” on page 696.

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

order

Syntax order *number*;

Hierarchy Level [edit services application-identification rule *rule-name* address]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See “Configuring APPID Rules” on page 694.

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

port-mapping

Syntax

```
port-mapping {
   port-range {
      tcp (port | range);
      udp (port | range);
   }
   disable;
}
```

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Define port-mapping properties for application identification.

Options The remaining statements are explained separately.

Usage Guidelines See "Defining an Application Identification" on page 692.

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

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 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 *port*—Port number.

range—Numeric port range. Enter this specification as [*minimum-value* – *maximum-value*].

Usage Guidelines See "Defining an Application Identification" on page 692 and "Configuring APPID Rules" on page 694.

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

profile

Syntax

```
profile profile-name {
   [ rule-set rule-set-name ];
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See “Configuring Application Profiles” on page 695.

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

rule

See the following sections:

- rule (Configuring) on page 717
- rule (Including in Rule Set) on page 718

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 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.

Usage Guidelines See “Configuring APPID Rules” on page 694.

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

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 Release 9.5.

Description Identify rules for inclusion in application rule set.

Options *rule-name*—Unique identifier for the rule.

Usage Guidelines See “Configuring APPID Rules” on page 694.

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

rule-set

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

Hierarchy Level [edit services application-identification],
[edit services application-identification profile]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See “Configuring APPID Rules” on page 694.

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

services

Syntax services application-identification { ... }

Hierarchy Level [edit]

Release Information services statement introduced before JUNOS Release 7.4.
application-identification statement introduced in JUNOS 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.

Usage Guidelines See "Application Identification Configuration Guidelines" on page 691.

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

session-timeout

See the following sections:

- session-timeout (interfaces) on page 720
- session-timeout (application-identification) on page 721
- [xref target has no title]

session-timeout (interfaces)

Syntax session-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced in JUNOS Release 10.0.

Description Define session lifetime globally for the Multiservices interface in seconds.

Options *seconds*—Duration of session.
Range: 4 through 86,400

Usage Guidelines See "Defining an Application Identification" on page 692.

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

session-timeout (application-identification)

Syntax session-timeout *seconds*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in JUNOS 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

Usage Guidelines See "Defining an Application Identification" on page 692.

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

source

Syntax
```
source {
   ip address/netmask;
   port-range {
      tcp (port | range);
      udp (port | range);
   }
}
```

Hierarchy Level [edit services application-identification rule *rule-name* address *address-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Define source properties for application-identification rule processing.

Options The remaining statements are explained separately.

Usage Guidelines See "Configuring APPID Rules" on page 694.

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

traceoptions

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 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.

Usage Guidelines See "Tracing APPID Operations" on page 697.

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

type

Syntax type *type*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Defining an Application Identification" on page 692.

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

type-of-service

Syntax type-of-service *service-type*;

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Defining an Application Identification" on page 692.

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

url

Syntax url *url*;

Hierarchy Level [edit services application-identification download]

Release Information Statement introduced in JUNOS Release 9.5.

Description Define the URL for application package downloads.

Options *url*—Download URL.

Usage Guidelines See “Configuring Automatic Download of Application Package Updates” on page 697.

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

Chapter 35

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 ];
                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 726
- Configuring AACL Rule Sets on page 729
- Example: Configuring AACL Rules on page 729

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;
            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;
        }
    }
}
```

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 726
- Configuring Match Conditions in AACL Rules on page 727
- Configuring Actions in AACL Rules on page 728

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 440.

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;
    source-address address <any-unicast>;
    source-address-range low minimum-value high maximum-value;
    source-prefix-list list-name;
}
```

Only IPv4 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 *JUNOS Policy Framework 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 the **source-prefix-list** statement in the AACL rule. For an example, see "Example: Configuring AACL Rules" on page 729.

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 "Application Identification Configuration Guidelines" on page 691.

- 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.

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 | none) | forwarding-class
        class-name);
}
```

You must include one of the following actions:

- 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 | 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.
  - none—Same as not specifying count as an action.
- 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 *JUNOS Policy Framework Configuration Guide*.

Configuring AACL Rule Sets

The rule-set statement defines a collection of AACLl 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.

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:

```
[edit serivces 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;
        }
    }
}
```

Chapter 36

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

Syntax applications [*application-names*];

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring Match Conditions in AACL Rules" on page 727.

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

application-groups

Syntax application-groups [*application-group-names*];

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See “Configuring Match Conditions in AACL Rules” on page 727.

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

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 Release 9.5.

Description Indicates that any application group defined in the database is considered a match.

Usage Guidelines See “Configuring Match Conditions in AACL Rules” on page 727.

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

destination-address

Syntax destination-address *address*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 9.5.

Description Specify the destination address for rule matching.

Options *address*—Destination IPv4 address or prefix value.

Usage Guidelines See "Configuring Match Conditions in AACL Rules" on page 727.

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*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 9.5.

Description Specify the destination address range for rule matching.

Options *minimum-value*—Lower boundary for the IPv4 address range.

maximum-value—Upper boundary for the IPv4 address range.

Usage Guidelines See "Configuring Match Conditions in AACL Rules" on page 727.

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

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 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.

Usage Guidelines See "Configuring Match Conditions in AACL Rules" on page 727.

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

from

Syntax
```
from {
    application-group-any;
    application-groups [ application-groups ];
    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;
}
```

Hierarchy Level [edit services aacl rule *rule-name* term *term-name*]

Release Information Statement introduced before JUNOS 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 *JUNOS Policy Framework Configuration Guide*.

The remaining statements are explained separately.

Usage Guidelines See "Configuring AACL Rules" on page 726.

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 aacl rule *rule-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring Match Direction for AACL Rules" on page 726.

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

rule

Syntax

```
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;
            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;
        }
    }
}
```

Hierarchy Level [edit services aacl],
[edit services aacl rule-set *rule-set-name*]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See “Configuring AACL Rules” on page 726.

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

rule-set

Syntax
```
rule-set rule-set-name {
   [rule rule-names ];
}
```

Hierarchy Level [edit services aacl]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring AACL Rule Sets" on page 729.

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

services

Syntax services aacl { ... }

Hierarchy Level [edit]

Release Information aacl statement introduced in JUNOS 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.

Usage Guidelines See "Application-Aware Access List Configuration Guidelines" on page 725.

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 aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 9.5.

Description Specify the source address for rule matching.

Options *address*—Source IPv4 address or prefix value.

Usage Guidelines See “Configuring Match Conditions in AACL Rules” on page 727.

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*;

Hierarchy Level [edit services aacl rule *rule-name* term *term-name* from]

Release Information Statement introduced in JUNOS Release 9.5.

Description Specify the source address range for rule matching.

Options *minimum-value*—Lower boundary for the IPv4 address range.

maximum-value—Upper boundary for the IPv4 address range.

Usage Guidelines See “Configuring Match Conditions in AACL Rules” on page 727.

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

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 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.

Usage Guidelines See "Configuring Match Conditions in AACL Rules" on page 727.

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

term

Syntax

```
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;
      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;
   }
}
```

Hierarchy Level [edit services aacl rule *rule-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Define the AACL term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

Usage Guidelines See “Configuring AACL Rules” on page 726.

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

then

Syntax

```
then {
    (accept | discard);
    count (application | application-group | application-group-any | none);
    forwarding-class class-name;
    policer policer-name;
}
```

Hierarchy Level [edit services aacl rule *rule-name* term *term-name*]

Release Information Statement introduced in JUNOS Release 9.5.
policer statement added in JUNOS Release 9.6.

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 | 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.
  - 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 *JUNOS Policy Framework Configuration Guide*.

Usage Guidelines See "Configuring AACL Rules" on page 726.

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

Related Topics *JUNOS Policy Framework Configuration Guide*

Chapter 37

Local Policy Decision Function Configuration Guidelines

This chapter includes the following sections:

- Configuring Statistics Profiles on page 743
- Applying L-PDF Profiles to Service Sets on page 746
- Tracing L-PDF Operations on page 747

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 743
- Configuring an AACL Statistics Profile on page 744

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.

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. 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 744. You are encouraged 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 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 *JUNOS 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-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:

- 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
- **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

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 15 through 1440 minutes.

For more information on configuring profiles, see the *JUNOS 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 ];
```

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 sets 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 {
            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 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 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
- rtsock—Routing socket traces
- statistics—Statistics traces
- subscriber—Subscriber traces

Chapter 38

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.

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 Release 10.0.

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 743.

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

aacl-statistics-profile

Syntax

```
aacl-statistics-profile profile-name {
    aacl-fields {
        field-name;
    }
    file filename;
    report-interval minutes;
}
```

Hierarchy Level [edit system services local-policy-decision-function statistics],
[edit services service-set *service-set-name*]

Release Information Statement introduced in JUNOS Release 10.0.

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.

report-interval *minutes*—Frequency at which statistics are recorded, in minutes.
Default: 15 minutes
Range: 15 through 1440 minutes

The remaining statements are explained separately.

Usage Guidelines See "Configuring Statistics Profiles" on page 743.

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

Related Topics For more information on profiles, see the *JUNOS Network Management Configuration Guide*.

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 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 743.

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

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 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 743.

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

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 Release 10.0.

Description Specify L-PDF properties.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring Statistics Profiles” on page 743.

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

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 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 743.

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

Related Topics For more information on profiles, see the *JUNOS Network Management Configuration Guide*.

statistics

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 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 743.

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

traceoptions

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 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
- 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 *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 747.

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.

Part 4
Encryption Services

- Encryption Overview on page 761
- Encryption Interfaces Configuration Guidelines on page 763
- Summary of Encryption Configuration Statements on page 773

Chapter 39

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 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)*

Chapter 40

Encryption Interfaces Configuration Guidelines

To enable encryption interfaces, you can configure the following properties:

- Configuring Encryption Interfaces on page 763
- Configuring Filters for Traffic Transiting the ES PIC on page 765
- Configuring an ES Tunnel Interface for a Layer 3 VPN on page 770
- Configuring ES PIC Redundancy on page 770
- Configuring IPsec Tunnel Redundancy on page 771

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.

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 765.

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 Network Interfaces Configuration Guide*.

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
      }
   }
}
```

Configuring Filters for Traffic Transiting the ES PIC

This section contains the following topics:

- Traffic Overview on page 765
- Configuring the Security Association on page 766
- Configuring an Outbound Traffic Filter on page 767
- Applying the Outbound Traffic Filter on page 768
- Configuring an Inbound Traffic Filter on page 768
- Applying the Inbound Traffic Filter to the Encryption Interface on page 769

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 4 on page 765, 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 *JUNOS Policy Framework Configuration Guide*.

Figure 4: Example: IPsec Tunnel Connecting Security Gateways

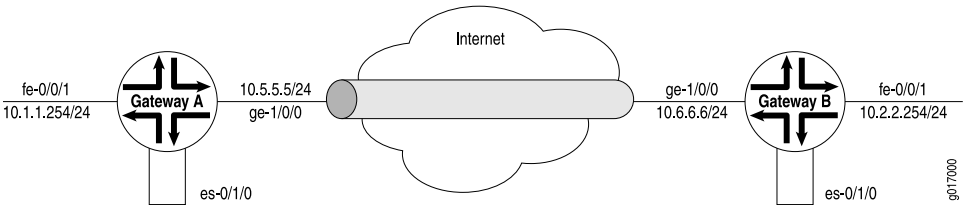

The SA and ES interface for security Gateway A are configured as follows:

```
[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 System Basics Configuration Guide*. For information about applying the SA to an interface, see "Specifying the Security Association Name for Encryption Interfaces" on page 764.

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 *JUNOS Policy Framework 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 4 on page 765). 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.

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 *JUNOS Policy Framework Configuration Guide*.

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 769. For more information about firewall filters, see the *JUNOS Policy Framework 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 769.

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 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;
```

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 System Basics Configuration Guide*, the *JUNOS Policy Framework Configuration Guide*, and "Example: Configuring an Inbound Traffic Filter" on page 769.

```
[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 5 on page 771 shows IPsec primary and backup tunnels.

Figure 5: IPsec Tunnel Redundancy

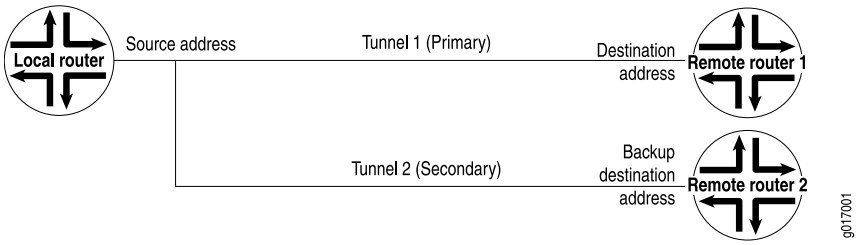

To configure IPsec tunnel redundancy, include the backup-destination statement at the [edit interfaces unit *logical-unit-number* tunnel] hierarchy level:

```
backup-destinationaddress;
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 Interfaces Configuration Guidelines" on page 1087.

Chapter 41

Summary of Encryption Configuration Statements

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

address

Syntax
```
address address {
    destination address;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced before JUNOS 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 763.

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

backup-destination

Syntax backup-destination *destination-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before JUNOS 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 771.

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

Related Topics destination

backup-interface

Syntax backup-interface *interface-name*;

Hierarchy Level [edit interfaces *interface-name* es-options]

Release Information Statement introduced before JUNOS 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 770.

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

destination

Syntax destination *destination-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 Release 7.4.

Description For tunnel and encryption interfaces, specify the remote address.

Options *destination-address*—Address of the remote side of the connection.

Usage Guidelines See "Configuring Encryption Interfaces" on page 763, "Configuring Traffic Sampling" on page 791, and "Configuring Flow Monitoring" on page 799.

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

es-options

Syntax
```
es-options {
    backup-interface interface-name;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS 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 770.

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

family

Syntax

```
family inet {
    ipsec-sa sa-name;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS 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 "Encryption Interfaces Configuration Guidelines" on page 763; for a general discussion of family statement options, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

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 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 765.

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 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 Network Interfaces Configuration Guide*.

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

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 Release 7.4.

Description Specify the IP Security (IPsec) SA name associated with the interface.

Options *sa-name*—IPsec SA name.

Usage Guidelines See "Encryption Interfaces Configuration Guidelines" on page 763.

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

Related Topics *JUNOS 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 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 763, "Configuring Traffic Sampling" on page 791, and "Configuring Flow Monitoring" on page 799.

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

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 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 "Encryption Interfaces Configuration Guidelines" on page 763 and "Tunnel Interfaces Configuration Guidelines" on page 1087.

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

Related Topics *JUNOS VPNs Configuration Guide*

unit

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 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 "Encryption Interfaces Configuration Guidelines" on page 763; for a general discussion of logical interface properties, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

Part 5

Flow Monitoring and Discard Accounting Services

- Flow Monitoring and Discard Accounting Overview on page 783
- Flow Monitoring and Discard Accounting Configuration Guidelines on page 787
- Summary of Flow-Monitoring Configuration Statements on page 841
- Flow Collection Configuration Guidelines on page 899
- Summary of Flow Collection Configuration Statements on page 911
- Dynamic Flow Capture Configuration Guidelines on page 929
- Flow-Tap Configuration Guidelines on page 941
- Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements on page 949

Chapter 42

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 Ethernet Services 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 on page 783
- Active Flow Monitoring on page 784

Passive Flow Monitoring

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 analyzers. Figure 6 on page 784 shows a typical topology for the passive flow-monitoring application.

Figure 6: Passive Monitoring Application Topology

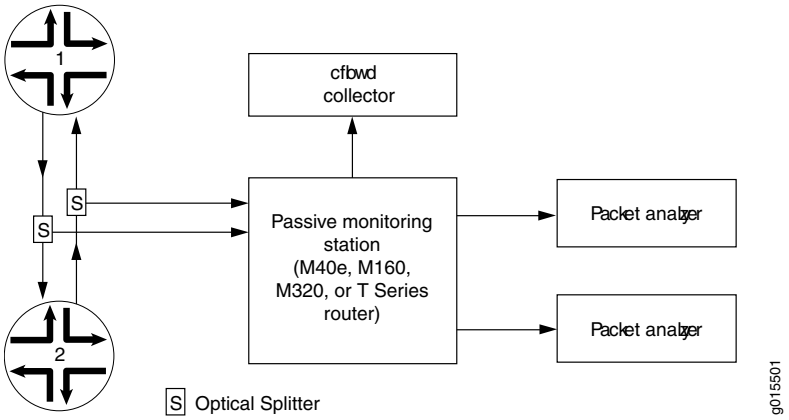

Traffic travels normally between Router 1 and Router 2. To redirect IPv4 traffic, you insert an optical splitter on the interface between these two routers. The optical splitter copies and redirects the traffic to the monitoring station, which is an M40e, M160, M320, or T Series router. The optical cable 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

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.

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
- Port mirroring and forwarding

- Port mirroring and discard accounting
- Sampling and port mirroring on different sets of traffic

Figure 7 on page 786 shows a sample topology.

Figure 7: Active Monitoring Configuration Topology

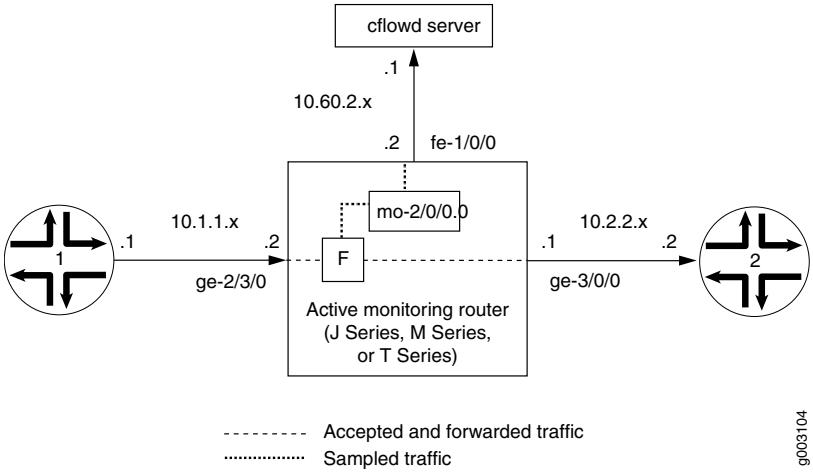

In Figure 7 on page 786, traffic from Router 1 arrives on the monitoring router's Gigabit Ethernet ge-2/3/0 interface. The exit interface on the monitoring router leading to destination Router 2 is ge-3/0/0, but this could be any interface type (such as SONET, Gigabit Ethernet, and so on). The export 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.

Chapter 43

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);
}
(at-fpc/pic/port | fe-fpc/pic/port | ge-fpc/pic/port) {
    passive-monitor-mode;
}
so-fpc/pic/port {
    unit logical-unit-number {
        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 (Forwarding Options) group-names {
   interface interface-name {
      next-hop address;
   }
}
port-mirroring {
   input {
      rate rate;
      run-length number;
   }
```

```
    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;
    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;
                }
                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 {
        file filename {
            files number;
            size bytes;
            (world-readable | no-world-readable);
        }
    }
}
```

NOTE: For the complete [edit forwarding-options] hierarchy, see the *JUNOS Policy Framework 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 ];
            }
            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 791
- Configuring Flow Monitoring on page 799
- Enabling Flow Aggregation on page 804
- Configuring Flow Aggregation to Use Version 5 or Version 8 cflowd on page 804
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 807

- Configuring Sampling Instances on page 814
- Directing Replicated Flows to Multiple Flow Servers on page 815
- Logging cflowd Flows Before Export on page 817
- Configuring Port Mirroring on page 818
- Load Balancing Among Multiple Monitoring Interfaces on page 828
- Configuring Discard Accounting on page 831
- Enabling Passive Flow Monitoring on page 833
- Configuring Services Interface Redundancy with Flow Monitoring on page 837

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 791
- Configuring Traffic Sampling on page 792
- Disabling Traffic Sampling on page 793
- Sampling Once on page 794
- Configuring Traffic Sampling Output on page 794
- Tracing Traffic Sampling Operations on page 796
- Traffic Sampling Examples on page 796

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* {

```
    term term-name {
        then {
            sample;
            accept;
        }
    }
}
```

For more information about firewall filter actions and action modifiers, see the *JUNOS Policy Framework 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 {
        family inet {
            max-packets-per-second number;
            rate number;
        }
    }
}
```

Configuring Traffic Sampling

To configure traffic sampling on any logical interface, include the input statement at the [edit forwarding-options sampling] hierarchy level:

```
input {
    family inet {
        max-packets-per-second number;
        rate number;
        run-length number;
    }
}
```

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.

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 8 on page 793).

Figure 8: Configure Sampling Rate

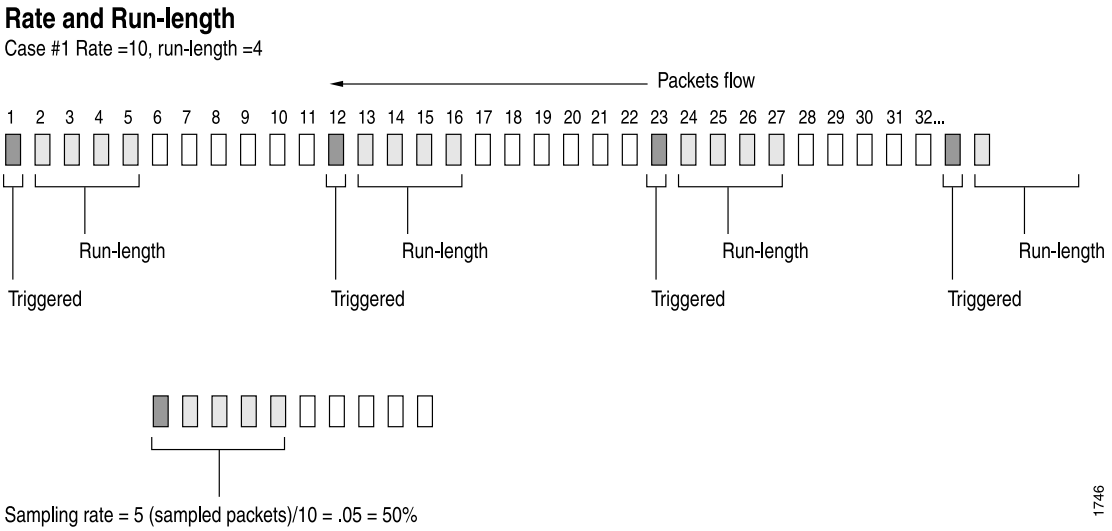

The rate statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, x number of packets out of every 10 is sampled, where x = run-length + 1. By default, the rate is 0, which means that no traffic is sampled.

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.

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. For more information about the output file formats, see "Configuring Traffic Sampling Output" on page 794.

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 output] hierarchy level:

```
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);
    label-position {
        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;
}
```

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 804.

The aggregate-export-interval statement is described in "Configuring Discard Accounting" on page 831, and the flow-active-timeout and flow-inactive-timeout statements are described in "Configuring Flow Monitoring" on page 799.

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 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
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 {
file filename <files number <size bytes> <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 796
- Example: Sampling All Traffic from a Single IP Address on page 797
- Example: Sampling All FTP Traffic on page 798

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;
      }
   }
}
```

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;
    }
  }
  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;
      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;
        }
    }
    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;
        }
    }
}
```

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;
    }
  }
  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 799
- Configuring Flow-Monitoring Properties on page 800
- Example: Configuring Flow Monitoring on page 803

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:

```
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);
      }
    }
```

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.

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 System Basics Configuration Guide*.

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 801
- Exporting Flows on page 801
- Configuring Time Periods when Flow Monitoring is Active and Inactive on page 802

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 Software 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 Configuration Guidelines” on page 899.

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 804.

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 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 Feature Guide*. For information on cflowd, see "Enabling Flow Aggregation" on page 804.

```
[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;
            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;
         }
      }
   }
}
```

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 as well as IPv4 and IPv6 traffic, and to combine configuration statements between the two formats.

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 804
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 807
- Directing Replicated Flows to Multiple Flow Servers on page 815
- Logging cflowd Flows Before Export on page 817

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 cflowd statement:

```
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;
    version format;
}
```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling output]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

You must configure the **family inet** statement on the monitoring interface and supply a destination address, as in the following example:

```
[edit interfaces]
sp-3/0/0 {
    unit 0 {
        family inet {
            address 10.2.3.4/32 {
                destination 192.68.7.8;
            }
        }
    }
}
```

Specification of **unit 0** is required, but the address values for the **address** and **destination** statements are arbitrary (that is, any address is accepted).

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 includeing 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 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 only one version 5 or one version 8 flow format at the [edit forwarding-options sampling output] hierarchy level for Routing Engine-based sampling. 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 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 818.

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 output cflowd *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 Software complies with Version 2.1b1 of cflowd. If you do not include the caida-compliant statement in the configuration, the JUNOS Software 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.

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, or a combination of IPv4 and MPLS 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 Table 3 on page 36 or the appropriate hardware documentation.

The following sections contain additional information:

- Configuring the Traffic to be Sampled on page 808
- Configuring the Version 9 Template Properties on page 808
- Restrictions on page 809
- Fields Included in Each Template Type on page 810
- MPLS Sampling Behavior on page 811
- Verification on page 812
- Examples: Configuring Version 9 Flow Templates on page 812

Configuring the Traffic to be Sampled

To specify sampling of IPv4, IPv6, or MPLS 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) {
max-packets-per-second number;
rate number;
run-length number;
```

You can include either family inet or family mpls.

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) {
        label-position [ positions ];
    }
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the template *name* statement.
- You then specify each template for the appropriate type of traffic by including the ipv4-template, ipv6–template, mpls-ipv4-template, or mpls-template statement.
- 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 output cflowd] 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;
            }
        }
    }
}
```

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
- 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
- Source MAC Address
- Destination MAC Address
- 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

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

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 MPLS Applications Configuration Guide*.

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).

   Resulting behavior is that MPLS packets are sent to 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.

   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 System Basics and Services Command Reference*.

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];
            }
         }
      }
   }
}
```

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;
      }
    }
    output {
      flow-active-timeout 60;
      flow-inactive-timeout 30;
      cflowd 1.2.3.4 {
        port 2055;
        version9 {
          template mpls-ipv4-template-1;
        }
      }
      interface sp-7/0/0 {
        source-address 1.1.1.1;
      }
    }
  }
}
```

Configuring Sampling Instances

You can configure active sampling by defining a sampling instance that specifies a name for the sampling parameters and binding the instance name to a particular 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 each sampling destination. This configuration is supported on MX Series, M120, M320, and T Series routers (but not on TX Matrix routers) and on the cflowd version5/8 and flow aggregation version 9 templates.

To implement this feature, you include the **instance** statement at the [**edit forwarding-options sampling**] hierarchy level:

```
instance instance-name { # named instances of sampling parameters
   input { # input parameters common to all protocol families for this instance
      rate number;
      run-length number;
   }
   family (inet | inet6 | mpls) {
      input { # input parameters specific to the designated protocol family
         rate number;
         run-length number;
      }
      output {
         flow-server server-name {
            version number {
               template template-name;
            }
         }
         interface {
            source-address address ;
         }
      }
   }
}
```

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.

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;
    }
}
```

For more information about chassis configuration, see the *JUNOS System Configuration Guide*

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 815
- Directing Replicated Version 9 Flow Aggregates to Multiple Servers on page 816

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.

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 {
        input {
            family inet {
                rate 1;
            }
        }
        output {
            cflowd 10.10.3.2 {
                port 2055;
                version 5;
                source-address 192.168.164.119;
            }
```

```
            cflowd 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 {
    sampling {
        input {
            family inet {
                rate 1;
            }
        }
        output {
            cflowd 10.10.3.2 {
                port 2055;
                version9 {
                    template {
                        ipv4;
                    }
                }
            }
            cflowd 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 cflowd *hostname*] hierarchy level:

```
[edit forwarding-options sampling output cflowd 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 794.

NOTE: Because the local-dump statement adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

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.

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]
input {
    rate rate;
    run-length number;
}
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>;
}
```

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.

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 796.

For more information about port mirroring, see the following sections:

- Configuring Tunnels on page 819
- Port Mirroring with Next-Hop Groups on page 820
- Filter-Based Forwarding with Multiple Monitoring Interfaces on page 821
- Restrictions on page 821
- Configuring Port Mirroring on Services Interfaces on page 822
- Examples: Configuring Port Mirroring on page 823

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 Interfaces Configuration Guidelines” on page 1087.

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 (Forwarding Options) 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.
- 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 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 also 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 Ethernet Services Routers Solutions Guide*.

On MX 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;
            }
        }
    }
}
```

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 823. 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.

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 Routing Protocols Configuration Guide*. For an example of FBF applied to an output interface, see "Examples: Configuring Port Mirroring" on page 823.

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, M320, and T Series routers, simultaneous IPv4 and IPv6 port mirroring is supported. Multiple next-hop mirroring is not supported.
- 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.
- 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.
- 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.

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 {
   input {
      family inet {
         rate 1;
      }
   }
   output {
      cflowd 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;
      }
}
```

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.
  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;
}
[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;
        }
        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;
                    }
                }
            }
        }
    }
}
```

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.

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 *JUNOS Policy Framework 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 {
    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 Network Interfaces Configuration Guide*.

- 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 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;
}
```

For more information on routing table groups, see the *JUNOS 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 *JUNOS Policy Framework 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;
            }
        }
    }
}
hash-key {
    family inet {
        layer-3;
    }
}
```

For more information on hash keys, see the *JUNOS Policy Framework 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;
        }
    }
}
```

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 804. The **flow-active-timeout** and **flow-inactive-timeout** statements are described in "Configuring Flow Monitoring" on page 799.

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 M40e, M160, M320, 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 interfaces.

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;
```

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.

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;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see “Configuring Flow-Monitoring Interfaces” on page 799.

For conformity with 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 834
- Example: Enabling Passive Flow Monitoring on page 835

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 MPLS Applications Configuration Guide*.

Removing MPLS Labels from Incoming Packets

The JUNOS Software can forward only IPv4 packets to a Monitoring Services, Adaptive Services, or Multiservices PIC. IPv4 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 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 ];
}
```

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.
- 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 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

Example: Enabling 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;
            }
        }
    }
}
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;
    }
}
```

Configuring Services Interface Redundancy with Flow Monitoring

Active monitoring services configurations on AS or Multiservices PICs 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 487. For information on operational mode commands, see the *JUNOS Interfaces Command Reference*.

A sample configuration follows.

```
interface {
    rsp0 {
        redundancy-options {
            primary sp-0/0/0;
            secondary sp-1/3/0;
        }
        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 {
        input {
            family inet {
                rate 1;
                run-length 0;
                max-packets-per-second 65535;
            }
        }
```

```
      output {
         cflowd 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;
         }
      }
   }
}
```

Chapter 44

Summary of Flow-Monitoring Configuration Statements

The following sections explain each of the flow-monitoring configuration statements. The statements are organized alphabetically.

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 number;
            engine-type number;
            source-address address;
        }
    }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS 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 831.

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

address

Syntax
```
address address {
    destination address;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced before JUNOS 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 799 or "Configuring Traffic Sampling" on page 791.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* 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 output]

Release Information Statement introduced before JUNOS 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 831.

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

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 output cflowd *hostname*]

Release Information Statement introduced before JUNOS 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 Software 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 804.

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

autonomous-system-type

Syntax autonomous-system-type (origin | peer);

Hierarchy Level [edit forwarding-options sampling output cflowd *hostname*]

Release Information Statement introduced before JUNOS 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 804.

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

cflowd

See the following sections:

- cflowd (Discard Accounting and Sampling) on page 846
- cflowd (Flow Monitoring) on page 847

cflowd (Discard Accounting and Sampling)

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 address;
   version format;
}
```

Hierarchy Level [edit forwarding-options accounting *name* output],
[edit forwarding-options sampling output]

Release Information Statement introduced before JUNOS 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.

You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options sampling 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 804.

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

cflowd (Flow Monitoring)

Syntax
```
cflowd hostname {
    port port-number;
}
```

Hierarchy Level [edit forwarding-options monitoring *name* family inet output]

Release Information Statement introduced before JUNOS 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 804.

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

core-dump

Syntax
```
(core-dump | no-core-dump);
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before JUNOS 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 Software saves the current core file (0) and the four previous core files, which are numbered from 1 through 4 (from newest to oldest):

- core-dump—Enable the core dumping operation.
- no-core-dump—Disable the core dumping operation.

Usage Guidelines See "Configuring Flow Monitoring" on page 799.

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

destination

Syntax destination *destination-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before JUNOS Release 7.4.

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 1087, "Configuring Traffic Sampling" on page 791, and "Configuring Flow Monitoring" on page 799.

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

disable

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 output file]

Release Information Statement introduced before JUNOS Release 7.4.
Statement added to port-mirror hierarchy in JUNOS Release 9.6.

Description Disable traffic accounting, port mirroring, or sampling.

Usage Guidelines See "Configuring Traffic Sampling" on page 791 or "Configuring Port Mirroring" on page 818.

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 Release 9.6.

Description Disable all port mirroring instances globally.

Usage Guidelines See "Configuring Port Mirroring" on page 818.

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

engine-id

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 output interface *interface-name*]

Release Information Statement introduced before JUNOS 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 791, "Configuring Flow Monitoring" on page 799, or "Configuring Discard Accounting" on page 831.

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

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 output interface *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the engine type number for flow monitoring and accounting services.

Options *number*—Platform-specific accounting interface type.

Usage Guidelines See "Configuring Traffic Sampling" on page 791, "Configuring Flow Monitoring" on page 799, or "Configuring Discard Accounting" on page 831.

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

export-format

Syntax export-format *format*;

Hierarchy Level [edit forwarding-options monitoring *name* output]

Release Information Statement introduced before JUNOS 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 801.

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

Related Topics version

family

See the following sections:

- family (Interfaces) on page 851
- family (Monitoring) on page 852
- family (Port Mirroring) on page 853
- family (Sampling) on page 853

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 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 799.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other options not used with services interfaces.

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 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 799.

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

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 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 818.

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

family (Sampling)

Syntax

```
family (inet | inet6 | mpls) {
    max-packets-per-second number;
    rate number;
    run-length number;
}
```

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced before JUNOS Release 7.4.
mpls option introduced in Release 8.3.
inet6 option introduced in Release 9.4.

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.

Usage Guidelines See "Configuring Traffic Sampling" on page 791.

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

file

See the following sections:

- file (Sampling) on page 854
- file (Trace Options) on page 854

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 output]

Release Information Statement introduced before JUNOS Release 7.4.

Description Collect the traffic samples in a file.

The statements are explained separately.

Usage Guidelines See "Configuring Traffic Sampling" on page 791.

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 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 796.

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

filename

Syntax filename *filename*;

Hierarchy Level [edit forwarding-options sampling output file]

Release Information Statement introduced before JUNOS 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 791.

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 output file],
[edit forwarding-options sampling traceoptions file]

Release Information Statement introduced before JUNOS 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 818 or "Configuring Traffic Sampling" on page 791.

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

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 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 799.

Required Privilege Level

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

Related Topics *JUNOS Policy Framework Configuration Guide* or the *JUNOS System Basics Configuration Guide*

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 output],
[edit services flow-monitoring version9]

Release Information Statement introduced before JUNOS 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 802 or "Configuring the Version 9 Template Properties" on page 808.

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

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 Release 7.4.

Description Configure flow collection.

Options cflowd-collector—cflowd collector.

collector-pic—Collector PIC.

Usage Guidelines See "Exporting Flows" on page 801.

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

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 output],
[edit services flow-monitoring version9]

Release Information Statement introduced before JUNOS 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 802 or "Configuring the Version 9 Template Properties" on page 808.

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

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 ];
            }
            option-refresh-rate packets packets seconds seconds;
            template-refresh-rate packets packets seconds seconds;
        }
    }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in JUNOS 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 807.

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

forwarding-options

Syntax forwarding-options { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure traffic forwarding.

The statements that apply to services interfaces are explained separately. For other statements, see the *JUNOS Policy Framework Configuration Guide*.

Usage Guidelines See "Flow Monitoring and Discard Accounting Configuration Guidelines" on page 787.

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

input

See the following sections:

- input (Port Mirroring) on page 861
- input (Sampling) on page 861

input (Port Mirroring)

Syntax

```
input {
    rate number;
    run-length number;
}
```

Hierarchy Level [edit forwarding-options port-mirroring],
[edit forwarding-options port-mirroring instance *instance-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure port mirroring on a logical interface.

The statements are explained separately.

Usage Guidelines See "Configuring Port Mirroring" on page 818.

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

input (Sampling)

Syntax

```
input {
    family (inet | inet6 | mpls) {
        max-packets-per-second number;
        rate number;
        run-length number;
    }
}
```

Hierarchy Level [edit forwarding-options sampling],
[edit forwarding-options sampling instance *instance-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure traffic sampling on a logical interface.

The statements are explained separately.

Usage Guidelines See "Configuring Traffic Sampling" on page 791.

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

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 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 799.

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

instance

See the following sections:

- instance (Port Mirroring) on page 863
- instance (Sampling) on page 864

instance (Port Mirroring)

Syntax

```
instance instance-name {
    disable;
    input {
        rate number;
        maximum-packet-length number;
    }
    family (inet | inet6 | vpls) {
        output {
            next-hop-group group-name;
        }
    }
}
```

Hierarchy Level [edit forwarding-options port-mirror]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure a sampling instance.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Sampling Instances” on page 814.

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

instance (Sampling)

Syntax

```
instance instance-name {
    disable;
    input {
        rate number;
        run-length number;
    }
    family (inet | inet6 | vpls) {
        output {
            interfaceaddress;
            flow-server server-name {
                version number {
                    template template-name;
                }
            }
        }
    }
}
```

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced in JUNOS Release 9.6.

Description Configure a sampling instance.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Sampling Instances” on page 814.

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

interface

See the following sections:

- interface (Accounting or Sampling) on page 865
- interface (Monitoring) on page 866
- interface (Port Mirroring) on page 866

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 output]

Release Information Statement introduced before JUNOS 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 831 or "Configuring Traffic Sampling" on page 791.

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

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 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 799.

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 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 818.

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 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 Network Interfaces Configuration Guide* 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 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 807.

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 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 807.

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

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 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 807.

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 output cflowd *hostname*]

Release Information Statement introduced before JUNOS 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 804.

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

max-packets-per-second

Syntax max-packets-per-second *number*;

Hierarchy Level [edit forwarding-options sampling input family (inet | mpls)]

Release Information Statement introduced before JUNOS 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 791.

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

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 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 799.

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

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 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 807.

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 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 807.

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

multiservice-options

Syntax
```
multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
}
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port*]

Release Information Statement introduced before JUNOS 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 799.

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

next-hop

Syntax next-hop *address*;

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6) output interface *interface-name*]

Release Information Statement introduced before JUNOS 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 818.

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

next-hop-group

See the following sections:

- next-hop-group (Forwarding Options) on page 873
- next-hop-group (Port Mirroring) on page 874

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 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. 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 818.

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

next-hop-group (Port Mirroring)

Syntax next-hop-group *group-name*;

Hierarchy Level [edit forwarding-options port-mirror family (inet | vpls) output],
[edit forwarding-options port-mirror instance *instance-name* family (inet | vpls) output],

Release Information Statement introduced in JUNOS 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.

Options *group-name*—Name of next-hop group.

Usage Guidelines See "Port Mirroring with Next-Hop Groups" on page 820.

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 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 818.

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

no-stamp

See stamp

no-syslog

See syslog

no-world-readable

See world-readable

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 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 807.

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

output

See the following sections:

- output (Accounting) on page 876
- output (Monitoring) on page 877
- output (Port Mirroring) on page 877
- output (Sampling) on page 878

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 Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Usage Guidelines See “Configuring Discard Accounting” on page 831.

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

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 Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The statements are explained separately.

Usage Guidelines See "Configuring Flow Monitoring" on page 799.

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 Release 7.4.

Description Configure output interfaces and flow properties.

The statements are explained separately.

Usage Guidelines See "Configuring Port Mirroring" on page 818.

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

output (Sampling)

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);
        label-position {
            template template-name;
        }
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
    }
    file {
        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;
    }
}
```

Hierarchy Level [edit forwarding-options sampling],
[edit forwarding-options sampling instance *instance-name* family *family*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure cflowd, output files and interfaces, and flow properties.

The statements are explained separately.

Usage Guidelines See “Configuring Traffic Sampling” on page 791.

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 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 799.

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

passive-monitor-mode

Syntax passive-monitor-mode;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS 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 833.

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

Related Topics multiservice-options

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 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.

This statement has no effect on IP packets with more than two MPLS labels. 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 834.

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

Related Topics *JUNOS Network Interfaces Configuration Guide*

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 output cflowd *hostname*]

Release Information Statement introduced before JUNOS 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 804.

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

port-mirroring

Syntax

```
port-mirror {
    input {
        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 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 818.

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

rate

Syntax rate *number*;

Hierarchy Level [edit forwarding-options port-mirroring input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling family (inet | inet6 | mpls)],
[edit forwarding-options sampling instance *instance-name* input]

Release Information Statement introduced before JUNOS 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 818 or "Configuring Traffic Sampling" on page 791.

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 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 833.

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

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 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 833.

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 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 834.

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

Related Topics *JUNOS Network Interfaces Configuration Guide*

run-length

Syntax run-length *number*;

Hierarchy Level [edit forwarding-options port-mirroring input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling family (inet | inet6 | mpls)],
[edit forwarding-options sampling instance *instance-name* input]

Release Information Statement introduced before JUNOS Release 7.4.

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

Usage Guidelines See "Configuring Port Mirroring" on page 818 or "Configuring Traffic Sampling" on page 791.

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

sample-once

Syntax sample-once;

Hierarchy Level [edit forwarding-options sampling]

Release Information Statement introduced in JUNOS Release 9.6.

Description Sample active monitoring traffic only once.

Usage Guidelines See "Configuring Traffic Sampling" on page 791.

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

sampling

See the following sections:

- sampling (Forwarding Options) on page 886
- sampling (Interfaces) on page 887

sampling (Forwarding Options)

Syntax

```
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;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            label-position {
                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 {
        file filename <files number <size bytes> <world-readable | no-world-readable>;
    }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure traffic sampling.

The statements are explained separately.

Usage Guidelines See "Configuring Traffic Sampling" on page 791.

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

sampling (Interfaces)

Syntax sampling *direction*;

Hierarchy Level [edit interfaces mo-*fpc/pic/port* unit *logical-unit-number* family inet]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the direction of traffic to be sampled.

Options 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.

Usage Guidelines See "Configuring Flow Monitoring" on page 799.

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

services

Syntax services { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS 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 807.

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 output file],
[edit forwarding-options sampling traceoptions file]

Release Information Statement introduced before JUNOS 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 818 or "Configuring Traffic Sampling" on page 791.

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 forwarding-options accounting *name* output interface *interface-name*],
[edit forwarding-options monitoring *name* family inet output interface *interface-name*],
[edit forwarding-options sampling output cflowd *hostname*],
[edit forwarding-options sampling output interface *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the source address for monitored packets.

Options *address*—Interface source address.

Usage Guidelines See "Configuring Discard Accounting" on page 831, "Configuring Flow Monitoring" on page 799, or "Configuring Traffic Sampling" on page 791.

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 output file]

Release Information Statement introduced before JUNOS 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 791.

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

syslog

Syntax (syslog | no-syslog);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before JUNOS 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 799.

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

template

See the following sections:

- template (Forwarding Options) on page 891
- template (Services) on page 892

template (Forwarding Options)

Syntax template *template-name*;

Hierarchy Level [edit forwarding-options sampling output cflowd version9]

Release Information Statement introduced in JUNOS 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 807.

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

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 ];
    }
    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 Release 8.3.

Description Specify the flow aggregation version 9 template properties. The remaining statements are explained separately.

Options *template-name*—Name of version 9 template.

Usage Guidelines See “Configuring Flow Aggregation to Use Version 9 Flow Templates” on page 807.

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

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 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 807.

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

traceoptions

Syntax
```
traceoptions {
   file filename <files number> <size bytes> <world-readable | no-world-readable>;
}
```

Hierarchy Level [edit forwarding-options port-mirroring],
[edit forwarding-options sampling]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure traffic sampling tracing operations.

The statements are explained separately.

Usage Guidelines See "Tracing Traffic Sampling Operations" on page 796.

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

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 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 Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do no affect services interfaces.

version

Syntax version *format*;

Hierarchy Level [edit forwarding-options accounting *name* output cflowd *hostname*],
[edit forwarding-options sampling output cflowd *hostname*]

Release Information Statement introduced before JUNOS 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 804.

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

Related Topics export-format

version9

See the following sections:

- version9 (Forwarding Options) on page 896
- version9 (Services) on page 897

version9 (Forwarding Options)

Syntax
```
version9 {
   template template-name;
}
```

Hierarchy Level [edit forwarding-options sampling output cflowd]

Release Information Statement introduced in JUNOS 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 807.

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

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 ];
      }
      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 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 807.

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 output file],
[edit forwarding-options sampling traceoptions file]

Release Information Statement introduced before JUNOS 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 818 or "Configuring Traffic Sampling" on page 791.

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

Chapter 45

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

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 901
- Sending cflowd Records to Flow Collector Interfaces on page 904
- Configuring Flow Collection Mode and Interfaces on Services PICs on page 904
- Example: Configuring Flow Collection on page 905

Configuring Flow Collection

This section describes the following tasks for configuring flow collection:

- Configuring Destination FTP Servers for Flow Records on page 901
- Configuring a Packet Analyzer on page 902
- Configuring File Formats on page 902
- Configuring Interface Mappings on page 903
- Configuring Transfer Logs on page 903
- Configuring Retry Attempts on page 904

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.

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 {
    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:

- 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 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 9 on page 905 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. 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 9: Flow Collector Interface Topology Diagram

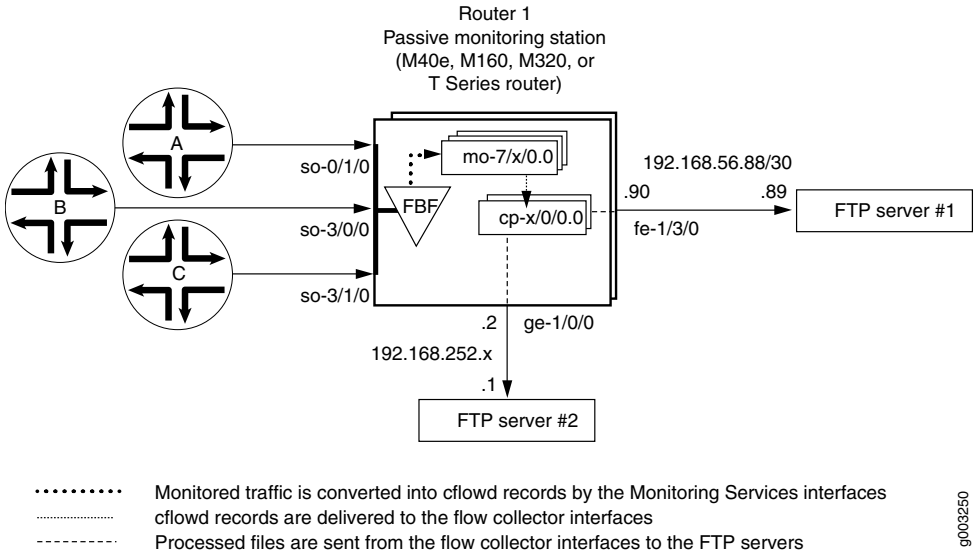


```
[edit]
chassis {
   fpc 6 {
      pic 0 {
         monitoring-services {
            application flow-collector; # This converts a Monitoring Services II or
                                        # Multiservices 400 PIC into a flow collector interface.
         }
      }
   }
   fpc 7 {
      pic 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;
      }
    }
}
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;
        }
    }
}
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 "$9$IXJK8xN-w2oZdbZDHmF30O1"; # SECRET-DATA
            }
            "ftp://user@192.168.252.1//tmp/collect2/" { # The secondary FTP server.
                password "$9$eIbvL7-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";
                }
            }
        ]
    }
}
```

Chapter 46

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 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 902.

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

analyzer-id

Syntax analyzer-id *name*;

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before JUNOS 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 902.

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 Release 7.4.

Description Specify the destination for transfer logs.

Options The statements are explained separately.

Usage Guidelines See “Configuring Transfer Logs” on page 903.

Required Privilege Level interface—To view this statement in the configuration.

collector

Syntax collector *interface-name*;

Hierarchy Level [edit services flow-collector interface-map]

Release Information Statement introduced before JUNOS 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 903.

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 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 902.

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

destinations

Syntax
```
destinations {
   ftp:url {
      password "password";
   }
}
```

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before JUNOS 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 901.

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 Release 7.4.

Description Configure the filename prefix for log files.

Options *prefix*—Filename identifier.

Usage Guidelines See “Configuring Transfer Logs” on page 903.

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

file-specification

See the following sections:

- file-specification (File Format) on page 915
- file-specification (Interface Mapping) on page 915

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 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 902.

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 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 903.

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

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 Release 7.4.

Description Define the flow collection.

Options The statements are explained separately.

Usage Guidelines See “Service Set Configuration Guidelines” on page 439.

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

ftp

See the following sections:

- ftp (Flow Collector Files) on page 919
- ftp (Transfer Log Files) on page 920

ftp (Flow Collector Files)

Syntax ftp:*url*;

Hierarchy Level [edit services flow-collector destination]

Release Information Statement introduced before JUNOS 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

Usage Guidelines See "Configuring Destination FTP Servers for Flow Records" on page 901.

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 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 903.

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 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 903.

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

maximum-age

Syntax maximum-age *minutes*;

Hierarchy Level [edit services flow-collector transfer-log-archive]

Release Information Statement introduced before JUNOS 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 903.

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

name-format

Syntax name-format *"format"*;

Hierarchy Level [edit services flow-collector file-specification variant *variant-number*]

Release Information Statement introduced before JUNOS 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).
- {year}—In the format YYYY; for example, 1970

- {year_abbr}—From 00 through 99

Usage Guidelines See "Configuring File Formats" on page 902.

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

password

See the following sections:

- password (Flow Collector File Servers) on page 924
- password (Transfer Log File Servers) on page 924

password (Flow Collector File Servers)

Syntax password "*password*";

Hierarchy Level [edit services flow-collector destination ftp:*url*]

Release Information Statement introduced before JUNOS 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 901.

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 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 903.

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

retry

Syntax retry *number*;

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before JUNOS 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 904.

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 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 904.

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

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 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 902.

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 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 903.

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

username

Syntax username *user-name*;

Hierarchy Level [edit services flow-collector transfer-log-archive archive-sites]

Release Information Statement introduced before JUNOS 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 903.

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 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 902.

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

Chapter 47

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 929
- Configuring Dynamic Flow Capture on page 930
- Example: Configuring Dynamic Flow Capture on page 937

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-01.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 Services Configuration Guidelines" on page 213.

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 10 on page 930 shows a sample topology. The number of control sources and content destinations is arbitrary.

Figure 10: Dynamic Flow Capture Topology

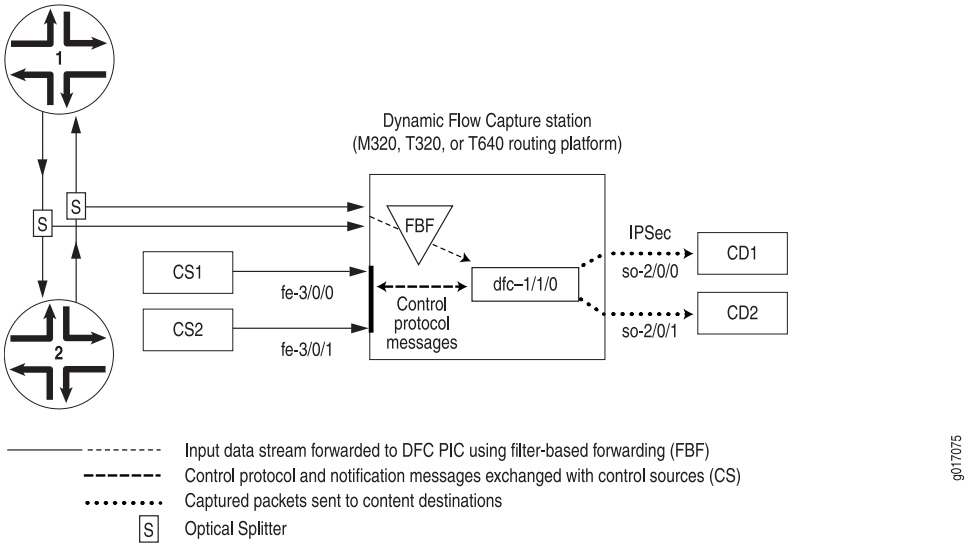

Liberal Sequence Windowing

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.

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;
}
```

This section describes the following tasks for configuring dynamic flow capture:

- Configuring the Capture Group on page 931
- Configuring the Content Destination on page 932
- Configuring the Control Source on page 933
- Configuring the DFC PIC Interface on page 934
- Configuring System Logging on page 935
- Configuring Thresholds on page 935
- Limiting the Number of Duplicates of a Packet on page 936

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;
    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:

```
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.
- 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:

```
[edit interfaces dfc-0/0/0]
unit 0 {
   family inet {
      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;
}
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 System Basics Configuration Guide*.

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 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
- 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;
```

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:

```
interfaces dfc-0/0/0 {
    unit 0 {
        family inet {
            address 2.1.0.0/32 { # DFC PIC address
                destination 10.36.100.1; # DFC PIC address used by
                # the control sources to correspond with
                # the monitoring platform
            }
        }
    }
}
unit 1 { # receive data packets on this logical interface
    family inet;
}
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;
        }
    }
}
```

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 833.

```
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;
            }
        }
    }
}
```

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;
    }
    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;
      }
    }
}
```

Chapter 48

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 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 M120, 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. 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 945.

NOTE: For information about dynamic flow capture, see "Dynamic Flow Capture Configuration Guidelines" on page 929. 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:

```
flow-tap {
    interface interface-name;
}
```

Other statements are configured at the [edit interfaces] and [edit system] hierarchy levels.

This chapter contains the following sections:

- Flow-Tap Architecture on page 942
- Configuring the Flow-Tap Service on page 943
- Configuring FlowTapLite on page 945
- Examples: Configuring Flow-Tap Services on page 946

Flow-Tap Architecture

The 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.
- 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 on IPsec tunnels, see "IPsec Services Configuration Guidelines" on page 213.
- Dynamic filters—The Packet Forwarding Engine automatically generates a firewall filter that is applied to all IPv4 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 AS or Multiservices PIC that is configured for flow-tap service. The AS 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;
        }
    }
    term LEA2_filter {
        from {
            source-address 10.1.1.1;
            source-port 23;
        }
        then {
            flow-tap;
        }
    }
```

}

Figure 11 on page 943 shows a sample topology that uses two mediation devices and two content destinations.

Figure 11: Flow-Tap Topology

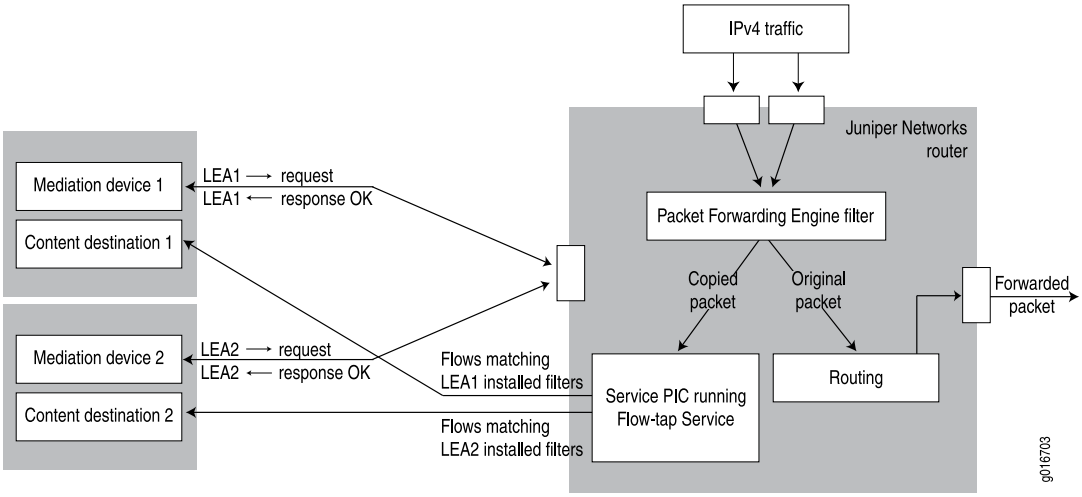

Configuring the Flow-Tap Service

This section describes the following tasks for configuring flow-tap service:

- Configuring the Flow-Tap Interface on page 943
- Strengthening Flow-Tap Security on page 944
- Restrictions on Flow-Tap Services on page 944

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*;

You can assign any AS or Multiservices PIC in the active monitoring router for flow-tap service, and use any logical unit on the PIC.

NOTE: You cannot configure dynamic flow capture (DFC) 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;
```

```
    }
}
```

Strengthening Flow-Tap Security

You can add an extra level of security to 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.
- **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 System Basics Configuration Guide*.

Restrictions on Flow-Tap Services

The following restrictions apply to flow-tap services:

- You cannot configure dynamic flow capture (DFC) and flow-tap features on the same router simultaneously.
- When the DFC process or the AS 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.

- If flow-tap is configured, you cannot configure the filter action then syslog for any firewall filter running on the same platform.
- Running the flow-tap application over an IPsec tunnel on the same router can cause packet loops and is not supported.

Configuring FlowTapLite

A lighter version of the flow-tap application is available on MX Series routers and also on M120 and M320 routers with Enhanced III FPCs. All of the functionality resides in the Packet Forwarding Engine rather than a service PIC or DPC.

NOTE: On M320 routers only, if the replacement of FPCs results in a mode change, you must restart the DFC process manually by disabling and then re-enabling the CLI configuration.

FlowTapLite uses the same DTCP-SSH architecture to install the 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-lite statement at the [edit services] hierarchy level:

```
flow-tap-lite {
   tunnel-interface interface-name;
}
```

For the Packet Forwarding Engine to encapsulate of the intercepted packet, it must send the packet to a tunnel logical (vt-) interface. You will need to allocate a tunnel interface and assign it to the dynamic flow control process (dfcd) 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 on this configuration, see the *JUNOS System Basics Configuration Guide*.

To configure the logical interfaces and assign them to dfcd, include the following configuration:

```
interfaces {
   vt-fpc/pic/port {
      unit 0 {
         family inet;
      }
   }
}
```

NOTE: If a service PIC or DPC is available, you can use its tunnel interface for the same purpose.

Examples: Configuring Flow-Tap Services

The following example shows all parts of a complete flow-tap configuration.

```
services {
   flow-tap {
      interface sp-1/2/0.100;
   }
}
interfaces {
   sp-1/2/0 {
      unit 100 {
         family inet;
      }
   }
}
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”;
         }
      }
   }
```

```
}
```

The following example shows a FlowTapLite configuration:

```
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;
            }
        }
    }
    services {
        flow-tap-lite {
            tunnel-interface vt-0/0/0.0;
        }
    }
```

Chapter 49

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

Syntax address *address*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in JUNOS 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 932.

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

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 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 933.

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

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 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 931.

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

content-destination

Syntax

```
content-destination identifier {
    address 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 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 932.

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

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 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 933.

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 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 936.

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

Related Topics g-duplicates-dropped-periodicity, max-duplicates

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 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 Configuration Guidelines” on page 929.

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

flow-tap

Syntax
```
flow-tap {
    interface interface-name;
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in JUNOS Release 8.1.

Description Enable the flow-tap application on an interface.

Options The remaining statements are explained separately.

Usage Guidelines See "Flow-Tap Configuration Guidelines" on page 941.

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

flow-tap-lite

Syntax
```
flow-tap-lite {
    tunnel-interface interface-name;
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in JUNOS Release 9.5.

Description Enable the FlowTapLite application on an interface. This 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 tunnel-interface *interface-name*—Specify the tunnel interface name.

Usage Guidelines See "Configuring FlowTapLite" on page 945.

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

g-duplicates-dropped-periodicity

Syntax g-duplicates-dropped-periodicity *seconds*;

Hierarchy Level [edit services dynamic-flow-capture]

Release Information Statement introduced in JUNOS 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 936.

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

Related Topics duplicates-dropped-periodicity

g-max-duplicates

Syntax g-max-duplicates *number*;

Hierarchy Level [edit services dynamic-flow-capture]

Release Information Statement introduced in JUNOS 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 936.

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

Related Topics max-duplicates

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 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 932.

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

Related Topics hard-limit-target

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 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 932.

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

Related Topics hard-limit

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 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 935.

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

interface

Syntax interface sp-*fpc*/*pic*/*port*.*logical-unit*;

Hierarchy Level [edit services flow-tap]

Release Information Statement introduced in JUNOS 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.

Usage Guidelines See "Configuring the Flow-Tap Interface" on page 943.

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

interfaces

Syntax interfaces *interface-name*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in JUNOS 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 934.

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

max-duplicates

Syntax max-duplicates *number*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name*]

Release Information Statement introduced in JUNOS 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 936.

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

Related Topics g-max-duplicates

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 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 933.

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

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 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 935.

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 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 933.

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

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 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 935.

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 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 933.

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

services

Syntax services dynamic-flow-capture { ... },
services flow-tap {...}

Hierarchy Level [edit]

Release Information dynamic-flow-capture statement introduced in JUNOS Release 7.4.
flow-tap statement introduced in JUNOS 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 "Dynamic Flow Capture Configuration Guidelines" on page 929 or "Flow-Tap Configuration Guidelines" on page 941.

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 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 933.

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

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 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 932.

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 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 932.

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

Related Topics soft-limit

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 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 933.

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

ttl

Syntax ttl *hops*;

Hierarchy Level [edit services dynamic-flow-capture capture-group *client-name* content-destination *identifier*]

Release Information Statement introduced in JUNOS 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 932.

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

Part 6
Link and Multilink Services

- Link and Multilink Services Overview on page 969
- Link and Multilink Services Configuration Guidelines on page 973
- Summary of Multilink and Link Services Configuration Statements on page 1011

Chapter 50

Link and Multilink Services Overview

The Multilink Protocol enables 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 Juniper Networks JUNOS Software supports several MP-based services PICs: 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 "Link Services IQ Interfaces Configuration Guidelines" on page 323. For more information about voice services, see "Voice Services Configuration Guidelines" on page 395.

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.

For link services IQ (lsq) interfaces, JUNOS Software CoS components are fully supported and are handled normally on M Series and T Series routers, as described in the *JUNOS Class of Service Configuration Guide*. There are some restrictions on J Series Services Routers; for more information on link services IQ configuration, see "Link Services IQ Interfaces Configuration Guidelines" on page 323.

The Link Services and Multilink Services PICs support the following MP encapsulation types:

- Multilink Point-to-Point Protocol (MLPPP)
- Multilink Frame Relay (MLFR)

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.

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*

NOTE: Endpoint Discriminator Class compatibility checking is enabled on MLPPP interfaces. Prior to JUNOS 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.

Chapter 51

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;
```

```
    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 974
- Configuring the Number of Bundles on Link Services PICs on page 975
- Configuring the Links in a Multilink or Link Services Bundle on page 976
- Multilink and Link Services Logical Interface Configuration Overview on page 977
- Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 979
- Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 979
- Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 981
- Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 982
- Configuring MRRU on Multilink and Link Services Logical Interfaces on page 982
- Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 983
- Configuring DLCIs on Link Services Logical Interfaces on page 984
- Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 985
- Configuring Link Services Physical Interfaces on page 988
- Configuring CoS on Link Services Interfaces on page 992
- Examples: Configuring Multilink Interfaces on page 997
- Examples: Configuring Link Interfaces on page 1000

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 16 on page 975. 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 16: Multilink 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.

For configuration information, see the following sections:

- Configuring the Number of Bundles on Link Services PICs on page 975
- Configuring the Links in a Multilink or Link Services Bundle on page 976

Configuring the Number of Bundles on Link Services PICs

The default number of bundles per Link Services PIC is 16, ranging from ls-*fpc*/*pic*/*port*:0 to ls-*fpc*/*pic*/*port*:15.

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 1001.

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*;

Each Link Services PIC can accommodate a maximum of 256 MLFR UNI NNI bundles. For more information, see the *JUNOS 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 "Examples: Configuring Link Interfaces" on page 1000.

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.

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 12 on page 976). The physical interface is usually connected to networks capable of supporting MLPPP or MLFR (FRF.15 or FRF.16).

Figure 12: Multilink Interface Configuration

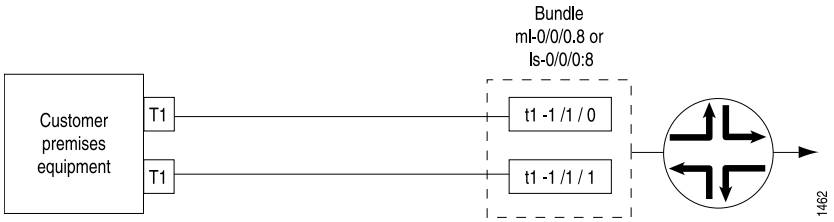

The following sample configuration refers to the topology in Figure 12 on page 976 and configures a multilink or link services bundle over a T1 connection (for which the T1 physical interface is already configured).

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 the section "Default Settings for Multilink and Link Services Logical Interfaces" on page 978.

For general information about logical unit properties or family inet properties, see the *JUNOS Network Interfaces Configuration Guide*. 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 976.

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 iterface properties, see the *JUNOS Network Interfaces Configuration Guide*.

Default Settings for Multilink and Link Services Logical Interfaces

Table 17 on page 978 lists the default settings for multilink and link services statements, together with the other permitted values or value ranges.

Table 17: 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 Table 18 on page 988 for statements that apply to link services physical interfaces only.

Related Topics

- Configuring Encapsulation for Multilink and Link Services Logical Interfaces on page 979
- Configuring the Drop Timeout Period on Multilink and Link Services Logical Interfaces on page 979
- Limiting Packet Payload Size on Multilink and Link Services Logical Interfaces on page 981
- Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces on page 982
- Configuring MRRU on Multilink and Link Services Logical Interfaces on page 982
- Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces on page 983

- Configuring DLCIs on Link Services Logical Interfaces on page 984
- Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces on page 985

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 Network Interfaces Configuration Guide*.

You can also configure physical interface encapsulation on link services interfaces. For more information, see "Configuring Encapsulation for Link Services Physical Interfaces" on page 989.

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.

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 990.

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.

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 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:

```
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 Software 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 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:

```
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*]

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 991.

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 341.

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 992.

Per-bundle CoS queuing is supported on link services IQ interfaces (lsq). For more information about link services IQ interfaces, see "Link Services IQ Interfaces Configuration Guidelines" on page 323.

The JUNOS Software 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 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 {
        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;
      }
      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 988.

The following sections explain how to configure link services physical interfaces:

- Default Settings for Link Services Interfaces on page 988
- Configuring Encapsulation for Link Services Physical Interfaces on page 989
- Configuring Acknowledgment Timers on Link Services Physical Interfaces on page 989
- Configuring Differential Delay Alarms on Link Services Physical Interfaces with MLFR FRF.16 on page 990
- Configuring Keepalives on Link Services Physical Interfaces on page 991

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 977.

Default Settings for Link Services Interfaces

Table 18 on page 988 lists the default settings for link services statements, together with the other permitted values or value ranges.

Table 18: 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
MRRU	1504 bytes	1500 through 4500 bytes

Table 18: Link Services Physical Interface Statements for MLFR FRF.16 *(continued)*

Option	Default Value	Possible Values
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 Network Interfaces Configuration Guide*.

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 979.

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 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.

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 Software uses ITU Q.933 Annex A LMIs for FRF.16. To use ITU Annex A LMIs, 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;
```

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.

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 Class of Service Configuration Guide*. For more information and detailed configuration examples, see "Link Services IQ Interfaces Configuration Guidelines" on page 323.

For link services (ls-) interfaces, CoS functions differently depending on whether the interface is configured on a Link Services PIM in a J Series router or on a Link Services PIC in an M Series or T Series router. For more information, see the following section or the *JUNOS Software Interfaces and Routing Configuration Guide*:

- CoS for Link Services Interfaces on M Series and T Series Routers on page 992
- Example: Configuring CoS on Link Services Interfaces on page 993

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 19 on page 992 summarizes how CoS queues work on link services (ls) interfaces.

Table 19: 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.
- 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 985 and the *JUNOS Policy Framework 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 a Link Services PIC installed on a J Series router, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Packets that do not match the firewall filters are sent to a queue that performs load balancing by sending fragments to all constituent links.

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;
        }
    }
}
[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 {
            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 {
                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 Network Interfaces Configuration Guide*.

The examples in this section show the following configurations:

- Example: Configuring a Multilink Interface with MLPPP on page 997
- Example: Configuring a Multilink Interface with MLPPP over ATM 2 Interfaces on page 998
- Example: Configuring a Multilink Interface with MLFR FRF.15 on page 999

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 {
```

```
    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;
      }
```

```
    }
  }
  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 {
            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 Network Interfaces Configuration Guide*.

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 20 on page 1001.

Table 20: 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 {
        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;
      }
```

```
    }
}
```

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;
        }
    }
```

```
}
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 {
        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 {
      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;
      }
```

```
    }
  }
  [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 {
      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;
          }
        }
        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;
        }
    }
}
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;
}
```

Chapter 52

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 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 989.

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

Related Topics action-red-differential-delay, hello-timer

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 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 989.

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

Related Topics address, hello-timer

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 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 990.

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

Related Topics yellow-differential-delay

address

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 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 976; for a general discussion of address statement options, see the *JUNOS Network Interfaces Configuration Guide*.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* 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 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 976.

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

destination

Syntax destination *destination-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.

Options *destination-address*—Address of the remote side of the connection.

Usage Guidelines See "Multilink and Link Services Logical Interface Configuration Overview" on page 977.

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

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 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 979.

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 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 984; for general information about Frame Relay DLCIs, see the *JUNOS Network Interfaces Configuration Guide*.

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

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 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 979.

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

encapsulation

See the following sections:

- encapsulation (Logical Interface) on page 1017
- encapsulation (Physical Interface) on page 1018

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 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 used only on multilink and link services interfaces and their constituent T1 or E1 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 979; for information about encapsulation statement options used with other interface types, see the *JUNOS Network Interfaces Configuration Guide*.

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

encapsulation (Physical Interface)

Syntax encapsulation (multilink-frame-relay-uni-nni | ...);

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS 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.

Usage Guidelines See "Configuring Encapsulation for Link Services Physical Interfaces" on page 989; for information about encapsulation statement options used with other interface types, see the *JUNOS Network Interfaces Configuration Guide*.

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

family

Syntax

```
family family {
   address address {
      destination address;
   }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS 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 "Link and Multilink Services Configuration Guidelines" on page 973; for a general discussion of family statement options, see the *JUNOS Network Interfaces Configuration Guide* .

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

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 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 981.

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

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 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 989.

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

Related Topics address, acknowledge-timer

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS 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 Network Interfaces Configuration Guide*.

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

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 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 985.

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 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 991.

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

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 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 982.

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

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);
    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 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 989.

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

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 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 982.

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

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 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 982.

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

Related Topics *JUNOS Network Interfaces Configuration Guide*

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 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 984.

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

n391

Syntax n391 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before JUNOS 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 991.

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

Related Topics n392, n393, t391, and t392

n392

Syntax n392 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before JUNOS 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 991.

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

Related Topics n391, n393, t391, and t392

n393

Syntax n393 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For link services interfaces only, set the Frame Relay monitored event count.

Options *number*—Event count.
Range: 1 through 255
Default: 6

Usage Guidelines See "Configuring Keepalives on Link Services Physical Interfaces" on page 991.

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

Related Topics n391, n392, t391, and t392

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 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 990.

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

Related Topics action-red-differential-delay, yellow-differential-delay

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 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 983.

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 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 991.

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

Related Topics n391, n392, n393, and t392

t392

Syntax t392 *number*;

Hierarchy Level [edit interfaces ls-*fpc/pic/port*:*channel* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before JUNOS 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 991.

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

Related Topics n391, n392, n393, and t391

unit

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 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 Services Configuration Guidelines" on page 973; for a general discussion of logical interface properties, see the *JUNOS Network Interfaces Configuration Guide* .

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

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 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 990.

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

Related Topics action-red-differential-delay, red-differential-delay

Part 7

Real-Time Performance Monitoring Services

- Real-Time Performance Monitoring Services Overview on page 1035
- Real-Time Performance Monitoring Configuration Guidelines on page 1037
- Summary of Real-Time Performance Monitoring Configuration Statements on page 1055

Chapter 53

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
- 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.

Chapter 54

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];
        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;
        }
    }
    probe-server {
```

```
        tcp {
            destination-interface interface-name;
            port (RPM) number;
        }
        udp {
            destination-interface interface-name;
            port (RPM) number;
        }
    }
    probe-limit limit;
    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;
        }
    }
}
```

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 1045.

This chapter includes the following sections:

- Configuring BGP Neighbor Discovery Through RPM on page 1038
- Configuring Real-Time Performance Monitoring on page 1040
- Examples: Configuring BGP Neighbor Discovery Through RPM on page 1050
- Examples: Configuring Real-Time Performance Monitoring on page 1051

Configuring BGP Neighbor Discovery Through RPM

BGP neighbors can be configured at the following hierarchy levels:

- [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];
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 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 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 1040
- Configuring RPM Receiver Servers on page 1044
- Limiting the Number of Concurrent RPM Probes on page 1045
- Configuring RPM Timestamping on page 1045
- Configuring TWAMP on page 1048

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.

NOTE: If you configure the hardware timestamp feature (see "Configuring RPM Timestamping" on page 1045), 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 44 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 1045. 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.
  - 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.
  - 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 {
    tcp {
        destination-interface interface-name;
        port (RPM) number;
    }
```

```
    udp {
        port (RPM) 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, and MX Series routers with a Multiservices DPC, you can enable hardware timestamping of RPM probe messages. The timestamp is applied on both the RPM client router (the router 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;
```

Specify the RPM client 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;
```

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.

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, 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 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;
```

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 *JUNOS Policy Framework Configuration Guide* or the *JUNOS Routing Protocols Configuration Guide*.

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;
         }
      }
   }
}
```

```
interfaces sp-1/2/0 {
    unit 2 {
        rpm client;
        family inet {
            address 10.8.4.2/32;
            filter {
                input f1;
            }
        }
    }
}
```

For more information about firewall filters, see the *JUNOS Policy Framework Configuration Guide*; for more information about queuing, see the *JUNOS 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;
        inactivity-timeout seconds;
        maximum-connections count;
        maximum-connections-per-client count;
        maximum-sessions count;
        maximum-sessions-per-connection count;
        port number;
    }
}
```

The TWAMP configuration process includes the following tasks:

- Configuring TWAMP Interfaces on page 1049
- Configuring TWAMP Servers on page 1049

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;
   inactivity-timeout seconds;
   maximum-connections count;
   maximum-connections-per-client count;
   maximum-sessions count;
   maximum-sessions-per-connection count;
   port number;
}
```

- 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.
- 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.

For examples of TWAMP configuration, see "Examples: Configuring Real-Time Performance Monitoring" on page 1051.

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;
    }
}
```

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 MF 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]
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;
            forwarding-class network-control;
         }
      }
   }
   [edit interfaces]
   fe-5/0/0 {
      unit 0 {
         family inet {
            filter {
               input recos;
            }
            address 10.8.4.2/24;
         }
      }
   }
```

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 {
      inactivity-timeout 20;
      maximum-sessions 5;
      maximum-sessions-per-connection 2;
      maximum-connections 3;
      maximum-connections-per-client 1;
      port 10000;
      client-list LIST-1 {
        address {
          20.0.0.2/30;
        }
      }
    }
  }
}
```

Chapter 55

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 | encrypted | none);

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in JUNOS 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.
- encrypted—Data packets are encrypted.
- none—No authentication or encryption.

Usage Guidelines See "Configuring TWAMP" on page 1048.

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

client-list

Syntax
```
client-list list-name {
    address address;
}
```

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in JUNOS 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 1048.

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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

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 Release 7.4.

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 1045), 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 44 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 1038 or "Configuring RPM Probes" on page 1040.

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

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 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.

Options *interface-name*—Name of the adaptive services interface.

Usage Guidelines See "Configuring RPM Probes" on page 1040, "Configuring RPM Receiver Servers" on page 1044, or "Configuring RPM Timestamping" on page 1045.

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

Related Topics hardware-timestamp, rpm

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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

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 Release 7.4.

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

Usage Guidelines See “Configuring RPM Probes” on page 1040.

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 Release 8.1.
Statement applied to MX-Series routers in JUNOS Release 10.0.

Description On MX Series routers 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 1045.

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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 services rpm twamp server]

Release Information Statement introduced in JUNOS 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 1048.

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 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 1038.

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

maximum-connections

Syntax maximum-connections *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in JUNOS 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 1048.

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

maximum-connections-per-client

Syntax maximum-connections-per-client *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in JUNOS 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 1048.

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 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 1048.

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

maximum-sessions-per-connection

Syntax maximum-sessions-per-connection *count*;

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in JUNOS 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 1048.

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 probe *owner* test *test-name*]

Release Information Statement introduced in JUNOS Release 8.5.

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 1040.

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

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 Release 8.5.

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 Real-Time Performance Monitoring" on page 1040.

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

Related Topics destination-interface, hardware-timestamp

port

See the following sections:

- port (RPM) on page 1065
- port (TWAMP) on page 1065

port (RPM)

Syntax port *number*;

Hierarchy Level [edit services rpm probe-server (tcp | udp)]

Release Information Statement introduced before JUNOS Release 7.4.

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 1044.

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 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 1048.

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

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;
      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 Release 7.4.

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 1040.

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

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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

probe-limit

Syntax probe-limit *limit;*

Hierarchy Level [edit services rpm]

Release Information Statement introduced before JUNOS Release 7.4.

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 1045.

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 (RPM) number;
   }
   udp {
      destination-interface interface-name;
      port (RPM) number;
   }
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced before JUNOS Release 7.4.

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 1044.

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

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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

routing-instance

Syntax routing-instance *instance-name*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before JUNOS Release 7.4.

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 1040.

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 Release 7.6.

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 1038.

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

rpm

Syntax rpm *client*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in JUNOS Release 8.1.

Description Associate an RPM client (router that originates RPM probes) with a specified Adaptive Services interface.

Options *client*—Identifier for RPM client router.

Usage Guidelines See "Configuring RPM Timestamping" on page 1045.

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

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 Release 9.3.

Description TWAMP server configuration settings.

Options The remaining statements are described separately.

Usage Guidelines See “Configuring TWAMP” on page 1048.

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

services

Syntax services rpm { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS 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 Configuration Guidelines” on page 1037.

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 Release 7.4.

Description Specify the source IP address used for probes. 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.

Options *address*—Valid IP address.

Usage Guidelines See "Configuring RPM Probes" on page 1040.

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

target

Syntax target (url *url* | address *address*);

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the destination address 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.

Usage Guidelines See "Configuring RPM Probes" on page 1040.

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

tcp

Syntax

```
tcp {
    destination-interface interface-name;
    port (RPM) port;
}
```

Hierarchy Level [edit services rpm probe-server]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the port information for the TCP server.

The remaining statements are explained separately.

Usage Guidelines See “Configuring RPM Receiver Servers” on page 1044.

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

test

Syntax

```
test test-name {
    data-fill data;
    data-size size;
    destination-interface interface-name;
    destination-port port;
    dscp-code-point dscp-bits;
    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 Release 7.4.

Description Represents 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 1040.

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

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 Release 7.4.

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 1038 or "Configuring RPM Probes" on page 1040.

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

thresholds

Syntax thresholds *thresholds*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before JUNOS Release 7.4.

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 1040.

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

traps

Syntax traps *traps*;

Hierarchy Level [edit services rpm probe *owner* test *test-name*]

Release Information Statement introduced before JUNOS Release 7.4.

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.

Usage Guidelines See “Configuring RPM Probes” on page 1040.

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

twamp

Syntax

```
twamp {
    server {
        client-list list-name {
            [ address address ];
        }
        authentication-mode mode;
        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]

Release Information Statement introduced in JUNOS Release 9.3.

Description TWAMP configuration settings.

Options The remaining statements are described separately.

Usage Guidelines See “Configuring TWAMP” on page 1048.

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

twamp-server

Syntax

```
twamp-server;
```

Hierarchy Level [edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in JUNOS Release 9.3.

Description Specify the service PIC logical interface to provide the TWAMP service.

Usage Guidelines See “Configuring TWAMP” on page 1048.

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 (RPM) port;
}
```

Hierarchy Level [edit services rpm probe-server]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the port information for the UDP server.

The remaining statements are explained separately.

Usage Guidelines See “Configuring RPM Receiver Servers” on page 1044.

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

Part 8
Tunnel Services

- Tunnel Services Overview on page 1083
- Tunnel Interfaces Configuration Guidelines on page 1087
- Summary of Tunnel Services Configuration Statements on page 1103

Chapter 56

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 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 "Encryption Interfaces Configuration Guidelines" on page 763 and the *JUNOS System Basics Configuration Guide*. For information about VPNs, see the *JUNOS VPNs Configuration Guide*. For information about MPLS, see the *JUNOS MPLS Applications Configuration Guide*.

The Juniper Networks JUNOS Software supports the tunnel types shown in Table 21 on page 1083.

Table 21: 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 medium or PIC. You must configure the interface for it to perform GRE.
gre	Internally generated GRE interface. This interface is generated by the JUNOS Software to handle GRE. It is not a configurable interface.

Table 21: Tunnel Interface Types *(continued)*

Interface	Description
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. Generally, IP routing allows packets to be routed directly to a particular address. However, in some instances you might need to route an IP packet to one address and then encapsulate it for forwarding to a different address. In a mobile environment in which the location of the end device changes, a different IP address might be used as the end device migrates between networks. Within a router, packets are routed to this 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 medium or PIC. 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 Software 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 J Series routers, it has a different function: it provides class-of-service (CoS) support for real-time performance monitoring (RPM) probe packets. Within a J Series router, packets are routed to this internal interface for services. The lt interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to perform CoS for RPM services.
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 medium or PIC. If your router has a Tunnel Services PIC, the JUNOS Software 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 Software to handle multicast tunnel services. It is not a configurable interface.

Table 21: 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 medium or PIC. You must configure the interface for it to perform PIM de-encapsulation.
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 medium or PIC. You must configure the interface for it to perform PIM encapsulation.
pimd	Internally generated PIM de-encapsulation interface. This interface is generated by the JUNOS Software to handle PIM de-encapsulation. It is not a configurable interface.
pime	Internally generated PIM encapsulation interface. This interface is generated by the JUNOS Software 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 J Series Services Routers. 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.

Chapter 57

Tunnel Interfaces Configuration Guidelines

This chapter includes the following tunnel interface configuration tasks and examples:

- Configuring Unicast Tunnels on page 1087
- Restricting Tunnels to Multicast Traffic on page 1092
- Configuring Logical Tunnel Interfaces on page 1092
- Configuring Tunnel Interfaces for Routing Table Lookup on page 1094
- Configuring Virtual Loopback Tunnels for VRF Table Lookup on page 1094
- Configuring PIM Tunnels on page 1096
- Configuring IPv6-over-IPv4 Tunnels on page 1096
- Configuring Dynamic Tunnels on page 1097
- Configuring Tunnel Interfaces on MX Series Routers on page 1097
- Examples: Configuring Unicast Tunnels on page 1098
- Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 1099
- Example: Configuring an IPv6-over-IPv4 Tunnel on page 1100
- Example: Configuring Logical Tunnels on page 1100

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;
            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.

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.

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.

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 Class of Service Configuration Guide*.

To prevent an invalid configuration, the JUNOS Software 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 1098 and the *JUNOS 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 1089
- Enabling Fragmentation on GRE Tunnels on page 1090
- Specifying an MTU Setting for the Tunnel on page 1090
- Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 1091
- Configuring Packet Reassembly on page 1091

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.

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.

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 240.

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 1089 and "Enabling Fragmentation on GRE Tunnels" on page 1090), 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 Network Interfaces Configuration Guide*.

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 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.

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.

NOTE: If your router has a Tunnel Services PIC, the JUNOS Software 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 and T Series routers, logical tunnel interfaces allow you to connect logical systems, virtual routers, or VPN instances. The router must be equipped with a Tunnel Services PIC or an Adaptive Services Module (only available on M7i routers). For more information on the latter applications, see the *JUNOS VPNs Configuration Guide*.

- On J Series Services Routers, lt- interfaces have a completely different function; they enable you to change class-of-service (CoS) classifications and header bits.

For more information, see the following section or the *JUNOS Software Interfaces and Routing Configuration Guide*:

- Connecting Logical Systems on page 1093

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]

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).

For more information about configuring logical systems, see the *JUNOS 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 Software 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.

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 1094.

For more information about VPNs, see the *JUNOS 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 22 on page 1095 describes each method.

Table 22: 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 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.

- 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: For the virtual loopback tunnel interface, none of the logical interface statements are valid, except for the family statement; in particular, you cannot configure IPv4 or IPv6 addresses on these interfaces. Also, virtual loopback tunnels 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 Software creates two interfaces to handle PIM tunnels:

- 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 router, you can configure IPv6-over-IPv4 tunnels. To do this, you configure a unicast tunnel across an existing IPv4 network infrastructure. IPv6 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.

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 1087. For an IPv6-over-IPv4 tunnel configuration example, see "Example: Configuring an IPv6-over-IPv4 Tunnel" on page 1100.

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;
    tunnel-type type-of-tunnel;
}
```

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 Routing Protocols Configuration Guide*. For more information about VPNs, see the *JUNOS VPNs Configuration Guide*.

Configuring Tunnel Interfaces on MX Series Routers

The MX Series Ethernet Services Routers support Dense Port Concentrators (DPCs) with built-in Ethernet ports and does not support Tunnel Services PICs. To create tunnel interfaces on an MX Series router, you configure a DPC and the corresponding Packet Forwarding Engine by including the tunnel-services statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level. For *slot-number*, specify the number of the DPC, and for *pic-number*, specify the number of the Packet Forwarding Engine. Each DPC includes four Packet Forwarding Engines.

You can also specify the amount of bandwidth to allocate for tunnel traffic on each Packet Forwarding Engine by including the **bandwidth** statement at the [**edit chassis fpc** *slot-number* **pic** *pic-number*] hierarchy level. For more information, see the *JUNOS System Basics Configuration Guide*.

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;
    }
}
```

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 {
            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;
         }
         family inet6 {
            address 2001:DB8:2:1/126;
         }
      }
   }
}
```

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;
      }
      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;
        }
    }
}
```

Chapter 58

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 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 1091.

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

Related Topics reassemble-packets

backup-destination

Syntax backup-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 Release 7.4.

Description For tunnel interfaces, specify the remote address of the backup tunnel.

Options *address*—Address of the remote side of the connection.

Usage Guidelines See "Configuring IPsec Tunnel Redundancy" on page 771.

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

Related Topics destination

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 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 1091.

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

destination

See the following sections:

- destination (Tunnel Remote End) on page 1105
- destination (Routing Instance) on page 1105

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 Release 7.4.

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 1087, "Configuring Traffic Sampling" on page 791, and "Configuring Flow Monitoring" on page 799.

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 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 1094.

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

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*],
[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 Release 7.4.

Description Create a tunnel for routes in these destination networks.

Options *prefix*—Destination prefix of network.

Usage Guidelines See "Configuring Dynamic Tunnels" on page 1097.

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

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 Release 9.2.

Description Disable fragmentation of GRE-encapsulated packets.

Default By default, fragmentation is disabled.

Usage Guidelines See "Configuring Packet Reassembly" on page 1091.

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

Related Topics reassemble-packets

dynamic-tunnels

Syntax
```
dynamic-tunnels tunnel-name {
    destination-networks prefix;
    source-address address;
    tunnel-type type-of-tunnel;
}
```

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 Release 7.4.

Description Configure a dynamic tunnel between two provider edge (PE) routers.

Options *tunnel-name*—Name of the dynamic tunnel.

The statements are explained separately in this chapter.

Usage Guidelines See "Configuring Dynamic Tunnels" on page 1097.

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

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure interfaces on the br.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Usage Guidelines See the *JUNOS Network Interfaces Configuration Guide*.

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

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 Release 7.4.

Description For Adaptive Services and Multiservices interfaces on M Series and T Series routers, identify an individual traffic flow within a tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*.

Options *number*—Value of the key.
Range: 0 through 4,294,967,295

Usage Guidelines See "Configuring a Key Number on GRE Tunnels" on page 1089.

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

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 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 1092.

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

Related Topics tunnel

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 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 1092.

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 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 1091.

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

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 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 1094.

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 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 Routing Protocols Configuration Guide* and the *JUNOS Policy Framework Configuration Guide*.

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

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 Release 7.4.

Description Configure protocol-independent routing properties.

Usage Guidelines See the *JUNOS Routing Protocols Configuration Guide*.

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

source

Syntax source *source-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before JUNOS Release 7.4.

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.

Usage Guidelines See "Tunnel Interfaces Configuration Guidelines" on page 1087.

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 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 Release 7.4.

Description Configure the tunnel source address.

Options *address*—Name of the source address.

Usage Guidelines See "Configuring Dynamic Tunnels" on page 1097.

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 Release 7.4.

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

Usage Guidelines See "Tunnel Interfaces Configuration Guidelines" on page 1087.

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

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 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 "Encryption Interfaces Configuration Guidelines" on page 763 and "Tunnel Interfaces Configuration Guidelines" on page 1087.

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

Related Topics *JUNOS VPNs Configuration Guide*

tunnel-type

Syntax tunnel-type *type*;

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 Release 7.4.

Description Select the dynamic tunnel type.

Options *type*—Tunnel type. Generic routing encapsulation (GRE) is supported.

Usage Guidelines See "Configuring Dynamic Tunnels" on page 1097.

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

unit

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 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 "Tunnel Interfaces Configuration Guidelines" on page 1087.

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

Related Topics *JUNOS Network Interfaces Configuration Guide* for other statements that do not affect services interfaces.

Part 9
Index

- Index on page 1119
- Index of Statements and Commands on page 1139

Index

Symbols

#, comments in configuration statements.....liv
( ), in syntax descriptions.....liv
< >, in syntax descriptions.....liii
[ ], in configuration statements.....liv
{ }, in configuration statements.....liv
| (pipe), in syntax descriptions.....liv

A

AACL
    action statements.....728
    applications.....727
    example configuration.....729
    match conditions.....727
    rules.....729
aacl-fields statement.....750
aacl-statistics-profile statement.....751
accelerations statement.....622
accept
    action.....791
accounting statement
    flow monitoring.....842
    usage guidelines.....831
acknowledge-retries statement.....1011
    usage guidelines.....989
acknowledge-timer statement.....1012
    usage guidelines.....989
action-red-differential-delay statement.....1012
    usage guidelines.....990
actions statement.....621
adaptive-services-pics statement.....457
    usage guidelines.....904
address statement
    APPID
        usage guidelines.....694
    application rule.....701
    DFC.....949
        usage guidelines.....932
    encryption.....773
        usage guidelines.....763
    flow monitoring.....843
        usage guidelines.....791
    interfaces.....493
        usage guidelines.....481
    link services.....1013
        usage guidelines.....976
    NAT.....155
        usage guidelines.....130
    voice services.....405
        usage guidelines.....396
address-range statement
    NAT.....156
administrative statement
    BGF.....518
admission-control statement.....623, 624
aggregate-export-interval statement.....843
    usage guidelines.....831
aggregation statement.....191
    flow monitoring.....844
    usage guidelines.....184, 804
alert (system logging severity level).....299, 451, 483
algorithm statement.....519
ALGs
    application protocols.....59
    configuring.....60
    definition.....59
allow-fragmentation statement.....1103
    usage guidelines.....1091
allow-ip-options statement.....118
    usage guidelines.....110
allow-multicast statement.....458
    usage guidelines.....452
allowed-destinations statement.....950
    usage guidelines.....933
analyzer-address statement.....911
    usage guidelines.....902
analyzer-id statement.....911
    usage guidelines.....902
anomaly checklist.....42
anti-replay-window-size statement.....257, 459
    usage guidelines.....241, 447
any (system logging severity level).....299, 451, 483
any-any match condition
    Ipsec.....237
APPID
    example configuration.....699
application layer gateways *See* ALGs

application protocol
definition....59
application statement....97, 702, 703
APPID
usage guidelines....692
usage guidelines....60
application-aware-access-list-fields statement....752
application-data-inactivity-detection statement....520
application-group statement....703
APPID
usage guidelines....695
application-group-any statement....732
AACL
usage guidelines....727
application-groups statement....704, 732
AACL
usage guidelines....727
APPID
usage guidelines....695
application-profile statement....425
usage guidelines....420
application-protocol statement....98
usage guidelines....60
application-set statement....99
usage guidelines....70
application-sets statement
CoS....426
usage guidelines....418
IDS....192
usage guidelines....183
NAT....156
usage guidelines....136
stateful firewall....119
usage guidelines....109
application-system-cache-timeout statement....704
APPID
usage guidelines....696
applications....183
example configuration....93
applications statement
AACL....731
usage guidelines....727
APPID
usage guidelines....695
applications hierarchy....99
usage guidelines....59
applications identification....705
CoS....426
usage guidelines....418
IDS....192
usage guidelines....183
NAT....157
usage guidelines....136
stateful firewall....119
usage guidelines....109
applying service set to interface....440
archive-sites statement....912
usage guidelines....903
AS PIC
multicast traffic....452
redundancy....302, 487, 837
attack detection....179
audit-observed-events-returns statement....520
authentication statement....258
usage guidelines....219
authentication-algorithm statement
IKE....259
usage guidelines....223
IPsec....259
usage guidelines....231
authentication-method statement....260
usage guidelines....223
authentication-mode statement
RPM....1055
automatic statement....705
APPID
usage guidelines....697
autonomous-system-type statement....845
usage guidelines....804
auxiliary-spi statement....260
usage guidelines....219

B

backup AS PIC....487
backup Link Services IQ PIC....328
backup-destination statement....774, 1104
usage guidelines....771
backup-interface statement....774
usage guidelines....770
backup-remote-gateway statement....261
usage guidelines....240
bandwidth
and delay buffer allocation....344
guaranteed....344, 348
base-root statement....521
bgf-core statement....522
BGP
router identifier....1111
blacklist-period statement....624
braces, in configuration statements....liv
brackets
angle, in syntax descriptions....liii
square, in configuration statements....liv
bundle statement....406, 1013
usage guidelines....400, 976
by-destination statement....193
usage guidelines....184
by-pair statement....194
usage guidelines....184
by-source statement....195
usage guidelines....184

C

cancel-graceful statement....523
capture-group statement....951
  usage guidelines....931
cflowd statement....846
  usage guidelines....804
CIR....348
cisco-interoperability statement....383
  usage guidelines....327
cleanup-timeout statement....524
clear-dont-fragment-bit statement
  GRE tunnel....494
  IPsec....261
    usage guidelines....239
  service-set....460
  usage guidelines....240, 448, 484, 1090
clear-ike-sas-on-pic-restart statement....262
  usage guidelines....222
clear-ipsec-sas-on-pic-restart statement....262
  usage guidelines....222
client-list statement....1056
collector statement....912
  usage guidelines....903
collector-pic statement
  usage guidelines....904
comments, in configuration statements....liv
committed-burst-size statement....625
committed-information-rate statement....626
compression statement....406
  usage guidelines....397, 398
compression-device statement....407
  usage guidelines....400
configuration
  dynamic flow capture interface....937
  flow collector interface....905
  flow-tap application....946
content destinations
  DFC....929
  flow-tap....942
content-destination statement....952
  usage guidelines....932
context-indications statement....525
control source
  DFC....929
control-association-indications statement....526
control-source statement....953
  usage guidelines....933
controller-address statement....526
controller-failure statement....527
controller-port statement....527
conventions
  text and syntax....liii
copy-tos-to-outer-ip-header statement....1104
  usage guidelines....1091
core-dump statement....847
  usage guidelines....799
CoS
  action statements....419
  applications....418
  example configuration....421
  for tunnels
    GRE TOS bits....1091
  link services interfaces....341, 343, 992
  link services IQ interfaces....323
  match conditions....418
  rules....421
  scheduler map
    configuration example....986
critical (system logging severity level)....299, 451, 483
curly braces, in configuration statements....liv
customer support....liv
  contacting JTAC....liv

D

Data inactivity detection....589
data-fill statement....1056
data-format statement....913
  usage guidelines....902
data-inactivity-detection statement....528, 626
data-size statement....1057
  usage guidelines....1040
datastore statement....627
dead peer detection (DPD) protocol....240
default statement....529
default-media-realm statement....628
delay buffer
  calculating....344, 348
  shaping rate....344, 348
delay-buffer-rate statement
  usage guidelines....344
delivery-function statement....530
description statement
  IKE....263
    usage guidelines....229
  IPsec....263
    usage guidelines....231, 234
destination statement
  APPID
    usage guidelines....694
  application identification rule....706
  encryption....775
    usage guidelines....763, 771
  flow monitoring....848
    usage guidelines....791
  link services....1014
    usage guidelines....976
  tunnel....1105
    usage guidelines....1087, 1094

destination-address statement
AACL....733
usage guidelines....727
BGF....530
CoS....427
usage guidelines....418
IDS....196
usage guidelines....183
IPsec....263
usage guidelines....237
NAT....157
usage guidelines....136
stateful firewall....120
usage guidelines....109
destination-address-range statement
AACL....733
usage guidelines....727
IDS....196
usage guidelines....183
NAT....158
usage guidelines....136
stateful firewall....120
usage guidelines....109
destination-interface statement
RPM....1058
usage guidelines....1040, 1045
destination-networks statement
tunnel....1106
usage guidelines....1097
destination-pool statement....158
usage guidelines....137
destination-port statement
applications....99
BGF....531
RPM....100, 1058
usage guidelines....65
destination-prefix statement....159, 197
usage guidelines....184
destination-prefix-ipv6 statement....197
usage guidelines....184
destination-prefix-list statement
AACL....734
usage guidelines....727
CoS....427
IDS....198
NAT....159
stateful firewall....121
usage guidelines....109
destinations statement
flow collection....913
usage guidelines....901
detect statement....531
DFC
architecture....929
capture group....931
control source configuration....933
destination configuration....932
example configuration....937
interface configuration....934
system logging....935
threshold configuration....935
dh-group statement....264
usage guidelines....224
dial-options statement....495
interfaces
usage guidelines....300
dialogs statement....629
diffserv statement....532
direction statement....265
usage guidelines....217
disable statement
APPID
usage guidelines....692, 695
application....707
application-group....707
flow monitoring....848
port mapping....708
traffic sampling
usage guidelines....793
disable-all-instances statement
flow monitoring....849
disable-global-timeout-override statement....708
usage guidelines....692
disable-mlppp-inner-ppp-pfc statement....1014
usage guidelines....979
disable-session-mirroring statement....532
discard accounting
usage guidelines....831
disconnect statement....533
dlci statement....1015
usage guidelines....984
DLCIs
multicast-capable connections....984
point-to-point connections....984
do-not-fragment statement
tunnel....1106
usage guidelines....1091
documentation
comments on....liv
down statement....533
download statement
APPID....709
usage guidelines....697
drop-timeout statement....1016
usage guidelines....979
dscp statement....428
BGF....534
BSG....630
usage guidelines....419
dscp-code-point statement
RPM....1059
usage guidelines....1040

DTCP.....929, 941
duplicates-dropped-periodicity statement.....953
usage guidelines.....936
dynamic authentication.....243
dynamic flow capture *See* DFC
dynamic route insertion.....244
dynamic rules.....244
dynamic security associations
usage guidelines.....221, 222
dynamic statement.....266
usage guidelines.....221
Dynamic Tasking Control Protocol *See* DTCP
dynamic tunnels
destination.....1106
source.....1112
dynamic-flow-capture statement.....954
dynamic-tunnels statement.....1107
usage guidelines.....1097

E

egress-service-point statement.....630
embedded-spdf statement.....631
emergency (system logging severity level).....299, 451, 483
enable flow collection mode.....904
encapsulation statement.....407
link services.....1017
usage guidelines.....979
voice services
usage guidelines.....399
encoding statement.....534
encryption interface.....763
applying inbound filter.....769
example configuration.....769
applying outbound filter.....768
example configuration.....767, 768
configuring inbound filter.....768
example configuration.....769
configuring MTU.....764
encryption statement.....267
usage guidelines.....220
encryption-algorithm statement
IKE.....268
usage guidelines.....224
IPsec.....268
usage guidelines.....232
engine-id statement
flow monitoring.....849
engine-type statement.....850
error (system logging severity level).....299, 451, 483
ES interfaces
example configuration.....764
ES PIC
apply inbound filter.....769
PIC redundancy.....770
redundancy
example configuration.....771
tunnel redundancy.....771
es-options statement.....775
usage guidelines.....770
event policy
all (tracing flag).....454
APPID.....699
configuration (tracing flag).....454
database (tracing flag).....454
events (tracing flag).....454
policy (tracing flag).....454
event-timestamp-notification statement.....535
export-format statement.....850
usage guidelines.....801

F

f-max-period statement.....408
usage guidelines.....397
facility-override statement.....309, 461, 496
usage guidelines.....450
failover statement.....537
failover-cold statement.....535
failover-warm statement.....536
family statement
encryption.....776
usage guidelines.....763
flow monitoring.....851
usage guidelines.....791
interfaces.....497
usage guidelines.....481
link services.....1019
usage guidelines.....976
voice services.....409
fast-update-filters statement.....538
file statement.....854
BGF.....539
border signaling gateway.....632
L-PDF statistics.....753
traffic sampling.....854
traffic sampling output
usage guidelines.....794, 796
file-specification statement.....915
usage guidelines.....902, 903
filename statement.....855
filename-prefix statement.....914
usage guidelines.....903
files
logging information output file.....796
traffic sampling output files.....794
var/log/sampled file.....796
var/tmp/sampled.pkts file.....794
files statement.....855
usage guidelines.....794

filter statement
encryption....777
usage guidelines....769
flow monitoring....856
usage guidelines....791
filters
used with services....440
firewall filters
actions....791
in traffic sampling....791
service filters....486
flag statement....540, 633
flow aggregation....804
multiple flow servers....815
flow collector
analyzer configuration....902
destination configuration....901
example configuration....905
file format configuration....902
interface mapping....903
transfer log....903
flow limiting....450
flow monitoring
example configuration
multiple port mirroring....823
next-hop groups....823
load balancing....828
overview....783
redundancy....837
flow server
replicating flows to multiple servers....815
flow-active-timeout statement....857
usage guidelines....802
flow-collector statement....916
usage guidelines....899, 904
flow-export-destination statement....858
usage guidelines....801
flow-inactive-timeout statement....858
usage guidelines....802
flow-monitoring statement....859
flow-tap
application....941
architecture....942
interface....943
permissions statement....944
RADIUS configuration....944
restrictions....944
security....944
flow-tap application
example configuration....946
flow-tap statement....955
flow-tap-dtcp statement....944
flow-tap-lite statement....955
font conventions....liii
force-entry statement....198
usage guidelines....184
forward-manipulation statement....634
forwarding classes
fragmentation....341
forwarding-class statement....384, 428
usage guidelines....341, 419
forwarding-options statement....860
usage guidelines....787
fragment-threshold statement
link services....1020
usage guidelines....981
LSQ....385
usage guidelines....341
voice services....410
usage guidelines....398
fragmentation
forwarding classes....341
GRE tunnels....1090
multiclass MLPPP....343
fragmentation and reassembly....398, 985
example configuration....986
fragmentation-map statement....385
usage guidelines....341
fragmentation-maps statement....386
usage guidelines....341
Frame Relay connections
point-to-point connections....984
Frame Relay encapsulation
multicast-capable connections....984
framework statement....635
FRF.12....398
example configuration....372
LFI....985
LSQ....369
FRF.15 and FRF.16....973
FRF.16....359
configuration example....363
from statement
AACL....734
usage guidelines....726
border signaling gateway
new call usage policy....638
new transaction policy....639
service class....640
CoS....429
usage guidelines....416
IDS....199
usage guidelines....181, 183
IPsec....269
usage guidelines....235, 237
NAT....160
usage guidelines....133, 136
stateful firewall....121
usage guidelines....108, 109
ftp statement
flow collection....918
usage guidelines....901, 903

FTP traffic, sampling....798
full-cone NAT....135

G

g-duplicates-dropped-periodicity statement....956
usage guidelines....936
g-max-duplicates statement....956
usage guidelines....936
gateway statement
BGF....541
border signaling gateway....641
gateway-address statement....545
gateway-controller statement....546
gateway-port statement....547
graceful statement....548
graceful-restart statement....549
GRE tunnels
fragmentation....1090
key number....1089
guaranteed rate....348
guaranteed-rate statement
usage guidelines....348

H

H.248
properties.552, 570, 571, 572, 573, 574, 575, 576, 577, 580, 581
BFG....532
BGF....534
h248-options statement....550
h248-profile statement....551
h248-properties statement....552
h248-stack statement....555
h248-timers statement....556
hanging-termination-detection statement....556
hard-limit statement....957
usage guidelines....932
hard-limit-target statement....957
usage guidelines....932
hardware requirements....3
hardware-timestamp statement....1060
hello-interval statement
L2TP....310
usage guidelines....298
hello-timer statement
link services....1020
usage guidelines....989
hide-avps statement....310
usage guidelines....299
hint statement....161
history-size statement....1060
usage guidelines....1039, 1040
host statement....461, 498
L2TP....311
usage guidelines....299, 450, 483
hot-standby statement....386

I

icmp-code statement....100
usage guidelines....63
icmp-type statement....101
usage guidelines....63
icons defined, notice....lii
idle-timeout statement....709
APPID
usage guidelines....692
IDS
action statements....184
applications....183
example configurations....188
match conditions....183
rules....181
ids-rule-sets statement
usage guidelines....444
ids-rules statement....462
usage guidelines....444
ignore-entry statement....198
usage guidelines....184
IKE....46, 222
authentication algorithm
usage guidelines....223
authentication-method statement
usage guidelines....223
DH (Diffie-Hellman) group
usage guidelines....224
dynamic SAs....222
encryption-algorithm statement
usage guidelines....224
lifetime
usage guidelines....225
mode statement
usage guidelines....227
policy....225
example....230
policy statement
usage guidelines....225
pre-shared-key statement
usage guidelines....227
proposals statement
usage guidelines....227
IKE security associations
clearing....222
ike statement....270
usage guidelines....222
ike-access-profile statement....462
usage guidelines....246, 446
inactivity-delay statement....557
inactivity-duration statement....557, 645
inactivity-non-tcp-timeout statement....710
usage guidelines....692

inactivity-tcp-timeout statement.....710
usage guidelines.....692
inactivity-timeout statement.....101
BGF.....558
flow monitoring.....498
RPM.....1061
usage guidelines.....68, 482
inactivity-timer statement.....558
index statement.....711
APPID
usage guidelines.....692, 695
info (system logging severity level).....300, 451, 484
initial-average-ack-delay statement.....559
initiate-dead-peer-detection statement.....271
usage guidelines.....241
input statement
flow monitoring.....861
interfaces.....499
usage guidelines.....440, 485
input-interface-index statement.....862
input-packet-rate-threshold statement.....958
usage guidelines.....935
inside and outside interfaces.....442
inside-service-interface statement
usage guidelines.....443
instance statement
port mirroring.....863
sampling.....864
usage guidelines.....814
interchassis LSQ failover.....326
interface preservation.....331
interface statement
encryption
usage guidelines.....763
flow monitoring.....865
usage guidelines.....819
flow-tap.....958
usage guidelines.....943
service interface pool.....615
interface style service sets.....443
interface-map statement.....920
usage guidelines.....903
interface-service statement.....463
usage guidelines.....440
interfaces
naming.....480
interfaces statement
DFC.....959
usage guidelines.....934
encryption.....777
usage guidelines.....763
flow monitoring.....867
usage guidelines.....791
interfaces hierarchy.....499
usage guidelines.....479
link services.....1021
usage guidelines.....973
tunnel.....1107
usage guidelines.....1087
voice services.....410
interim-ah-scheme statement.....559
interleave-fragments statement.....1021
usage guidelines.....985
Internet Key Exchange *See* IKE
intrachassis LSQ failover.....328
intrusion detection
example configurations.....188
rule set.....187
tasks.....179
IP addresses
sampling traffic from single IP addresses.....797
ip statement
APPID
usage guidelines.....694
application identification.....711
ip-flow-stop-detection statement.....560
IPsec
action statements.....239
authentication statement
usage guidelines.....219
authentication-algorithm statement
usage guidelines.....231
direction
usage guidelines.....217
dynamic authentication.....243
dynamic endpoints interface configuration.....247
dynamic rules.....244
dynamic security associations
usage guidelines.....221
encryption
usage guidelines.....220
encryption-algorithm statement
usage guidelines.....232
ES PIC.....763
example configuration.....249
inbound traffic.....769
outbound traffic.....767
IKE.....46
lifetime of SA.....232
lifetime-seconds statement.....232
match conditions.....237
minimum configurations
dynamic SA.....215
manual SA.....215
overview.....46
perfect-forward-secrecy statement
usage guidelines.....234
policy
overview.....233
policy statement
usage guidelines.....233

proposal statement
usage guidelines....231
proposals statement
usage guidelines....234
protocol statement (dynamic SA)
usage guidelines....233
protocol statement (manual SA)
usage guidelines....218
rule sets....242
security associations....46
security parameter index
usage guidelines....219
service set dynamic endpoints
configuration....246
traffic....765
ipsec statement....271
usage guidelines....231
ipsec-inside-interface
usage guidelines....244
ipsec-inside-interface statement....272
usage guidelines....237
ipsec-interface-id statement
usage guidelines....247
ipsec-sa statement
encryption....778
usage guidelines....763
ipsec-transport-security-association statement....560
ipsec-vpn-options statement....463
usage guidelines....445
ipsec-vpn-rule-sets statement
usage guidelines....444
ipsec-vpn-rules statement....464
usage guidelines....444
ipv4-template statement....867
IPv6
transition
configured tunnel....1096
ipv6-multicast-filter statement
usage guidelines....133
ipv6-multicast-interfaces statement....161
IPv6-over-IPv4 tunnel
example configuration....1100
standards supported....1096
ipv6-template statement....867

K

key statement
tunnel....1108
usage guidelines....1089

L

L2TP
access profile....296, 297
attribute-value pairs....299
example configuration....304
redundancy....302
timers....298
l2tp statement
usage guidelines....291
l2tp-access-profile statement....311
usage guidelines....297
l2tp-interface-id statement
usage guidelines....300
l2tp-profile statement
usage guidelines....296
label-position statement....868
latch-deadlock-delay statement....561
lawful intercept architecture....942
learn-sip-register statement....102
usage guidelines....68
LFI....365, 369, 398, 985
example configuration....368, 372, 986
lifetime-seconds statement
IKE....272
usage guidelines....225
IPsec....272
usage guidelines....232
limiting flows per service set....450
link fragmentation and interleaving *See* LFI
link PIC redundancy....331
link services interfaces
CoS components....341, 343, 992
example configuration....993, 1000
interleave fragments....985
example configuration....986
link services IQ interfaces....368
CoS components....323
example configuration....357, 363
link state replication....331
link-layer overhead....343
link services protocols....969
link state replication
LSQ PICs....331
link-layer overhead
link services IQ interfaces....343
link-layer-overhead statement....387
usage guidelines....338, 343, 351
lmi-type statement....1022
usage guidelines....991
load balancing
on monitoring interfaces....828
local-certificate statement....273
usage guidelines....228
local-dump statement....868
usage guidelines....817
local-gateway address statement....312
usage guidelines....297
local-gateway statement....464
usage guidelines....445

local-id statement.....273
usage guidelines.....229
local-policy-decision-function statement.....754
log output
adaptive services.....453
APPID.....698
traffic sampling.....796
log-prefix statement.....465, 500
L2TP.....312
usage guidelines.....299, 450, 483
logging statement.....200, 465
usage guidelines.....184
logical interfaces
multicast-capable connections.....984
logical tunnels.....1092
example configuration.....1100
logical-system statement
RPM.....1061
usage guidelines.....1039
loopback tunnels.....1094
LSQ bandwidth
oversubscribing.....344
LSQ failover
interchassis.....326
stateful intrachassis.....328
stateless intrachassis.....328
LSQ PICs.....331
redundancy.....328
lsq-failure-options statement.....387
usage guidelines.....326

M

manipulation-rule statement.....646
manual security association.....217
manual statement.....274
usage guidelines.....217
manuals
comments on.....liv
match direction usage in service sets.....442
match-direction statement
AACL.....735
usage guidelines.....726
CoS.....429
usage guidelines.....418
IDS.....200
usage guidelines.....181
IPsec.....274
usage guidelines.....235
NAT.....162
usage guidelines.....133
stateful firewall.....122
usage guidelines.....109
max-burst-size statement.....562, 563
max-checked-bytes statement.....712
APPID
usage guidelines.....696
max-concurrent-calls statement.....564
max-duplicates statement.....959
usage guidelines.....936
max-flows statement.....466
usage guidelines.....450
max-packets-per-second statement.....869
usage guidelines.....792
maximum-age statement.....921
usage guidelines.....903
maximum-connections statement.....1062
maximum-connections-per-client statement.....1062
maximum-contexts statement.....411
usage guidelines.....397
maximum-fuf-percentage statement.....565
maximum-inactivity-time statement.....566
maximum-net-propagation-delay statement.....567
maximum-records-in-cache statement.....648
maximum-send-window statement.....313
usage guidelines.....298
maximum-sessions statement.....1063
maximum-sessions-per-connection statement.....1063
maximum-synchronization-mismatches statement.....567
maximum-terms statement.....568
maximum-time-in-cache statement.....649
maximum-waiting-delay statement.....568
media statement.....569
media-policy statement.....647
media-service statement.....569
media-type statement.....647
mediation devices
flow-tap.....942
message-manipulation statement.....648
message-manipulation-rules statement.....650
mg-maximum-pdu-size statement.....570
mg-originated-pending-limit statement.....571
mg-provisional-response-timer-value statement.....572
mg-segmentation-timer statement.....573
mgc-maximum-pdu-size statement.....574
mgc-originated-pending-limit statement.....575
mgc-provisional-response-timer-value statement.....576
mgc-segmentation-timer statement.....577
min-checked-bytes statement.....712
APPID
usage guidelines.....696
minimum links
link services interfaces.....982
multilink interfaces.....982
minimum statement
BGF.....651
minimum-links statement.....1022
usage guidelines.....982

minimum-priority statement....960
usage guidelines....933
MLFR and MLPPP....973
mlfr-uni-nni-bundle-options statement....1023
usage guidelines....988, 991
MLPPP....354, 365
configuration example....357
example configuration....368
mode statement....275
usage guidelines....227
monitor statement....578
monitoring statement....870
usage guidelines....800
moving-average-size statement....1064
usage guidelines....1040
MPLS
packets
passive flow monitoring....834
mpls-ipv4-template statement....871
mpls-template statement....871
mrru statement....1024
usage guidelines....982
mss statement....201
usage guidelines....184
mtu statement....1025
multicast filters
for IPv6 NAT....133
multicast traffic
AS PIC....452
multicast tunnels....1092
multicast-capable connections
Frame Relay encapsulation....984
multicast-dlci statement....1025
usage guidelines....984
multicast-only statement....1108
usage guidelines....1092
multiclass MLPPP
fragmentation....343
multilink bundles
fractional T1....365
example configuration....368, 369, 372
FRF.12....369
example configuration....372
MLPPP....365
example configuration....368
NxT1....354, 359
configuration example....357, 363
multilink interfaces
example configuration....997
minimum links....982
multilink-class statement....388
usage guidelines....343
multilink-max-classes statement....388
usage guidelines....343
multiservice-options statement....872
Multiservices PIC
hardware requirements....33

N

n391 statement....1026
usage guidelines....991
n392 statement....1026
usage guidelines....991
n393 statement....1027
usage guidelines....991
name-format statement....922
usage guidelines....902
name-resolution-cache statement....652
NAT
action statements....137
address configuration....130
applications....136
example configuration....139
match conditions....136
rule sets....139
twice NAT
description....45
type....135
nat-pool statement....578
nat-rule-sets statement
usage guidelines....444
nat-rules statement....466
usage guidelines....444
nat-type statement....162
usage guidelines....135
neighbor discovery using IPv6 NAT....133
network-operator-id statement....579
new-call-usage-policies statement....652
new-call-usage-policy statement....653
new-call-usage-policy-set statement....654
new-transaction-policies statement....654
new-transaction-policy statement....655
new-transaction-policy-set statement....656
next-hop groups....818
next-hop statement....872
border signaling gateway....657
next-hop groups
usage guidelines....819
usage guidelines....818
next-hop style service sets....443
next-hop-group statement
forwarding-options....873
port mirroring....874
usage guidelines....818, 819
next-hop-service statement....467
usage guidelines....441
no-anti-replay statement....275, 468
usage guidelines....241, 447

no-application-identification statement.....713
APPID
usage guidelines.....696
no-application-system-cache statement.....713
APPID
usage guidelines.....696
no-clear-application-system-cache statement.....713
APPID
usage guidelines.....696
no-core-dump statement.....847
usage guidelines.....799
no-dscp-bit-mirroring statement.....579
no-filter-check statement.....874
usage guidelines.....818
no-fragmentation statement.....389
usage guidelines.....341
no-ipsec-tunnel-in-traceroute statement.....276
usage guidelines.....248
no-local-dump statement.....868
usage guidelines.....817
no-rtcp-check statement.....580
no-signature-based statement.....714
APPID
usage guidelines.....696
no-stamp statement.....889
usage guidelines.....794
no-syslog statement
DFC.....960
flow monitoring.....890
usage guidelines.....935
no-termination-request statement.....389
usage guidelines.....326
no-translation statement.....163
usage guidelines.....137
no-world-readable statement
flow monitoring.....898
usage guidelines.....794
normal-mg-execution-time statement.....580
normal-mgc-execution-time statement.....581
notice (system logging severity level).....300, 451, 483
notice icons defined.....lii
Notification behavior.....582
notification-behavior statement.....582
notification-rate-limit statement.....582
notification-regulation statement.....583
notification-targets statement.....961
usage guidelines.....933
NxT1 bundles
FRF.16.....359
configuration example.....363
MLPPP.....354
configuration example.....357

O

one-way-hardware-timestamp statement.....1064
usage guidelines.....1046
open-timeout statement.....500
usage guidelines.....482
option-refresh-rate statement.....875
order statement.....714
APPID
usage guidelines.....694
output files
logging information output file.....796
traffic sampling output files.....794
output statement.....501
discard accounting.....876
flow monitoring.....877
port mirroring.....877
sampling.....878
usage guidelines.....440, 485
output-interface-index statement.....879
outside-service-interface statement
usage guidelines.....443
overload-control statement.....583
overload-pool statement.....163
usage guidelines.....137
overload-prefix statement.....164
usage guidelines.....137
oversubscription.....344

P

packet-based IPsec.....237
parentheses, in syntax descriptions.....liv
passive flow monitoring.....783
MPLS packets.....834
passive-mode-tunneling statement.....468
usage guidelines.....449
passive-monitor-mode statement.....879
usage guidelines.....833
password statement
flow collection.....924
usage guidelines.....901, 903
peak-data-rate statement.....584, 585
peer-unit statement
tunnel.....1109
usage guidelines.....1092
per-unit-scheduler statement.....390
usage guidelines.....344, 348, 354, 359
perfect-forward-secrecy statement.....276
usage guidelines.....234
performance, monitoring.....1040
pgcp statement
NAT.....164
pgcp-rule-sets statement
usage guidelines.....444

pgcp-rules statement
service-set....469
usage guidelines....444
PIC types for services....3
pic-memory-threshold statement....961
usage guidelines....935
PIM
tunnels....1096
PIR....344
platform statement....586
platforms, supported....4
point-to-point connections
Frame Relay encapsulation....984
policy statement
IKE....277
usage guidelines....225
IPsec....278
usage guidelines....233
policy-decision-statistics-profile statement....755
pool statement....165
service interface pool....616
usage guidelines....130
pop-all-labels statement....880
usage guidelines....834
port mirroring....818
disabling....848
disabling all instances....849
port statement
cflowd
usage guidelines....804
flow monitoring....880
NAT....166
usage guidelines....130
RPM....1065
TWAMP....1065
voice services....411
usage guidelines....397
port-mapping statement....715
port-mirror statement....881
port-mirroring statement
usage guidelines....818
port-range statement....715
APPID
usage guidelines....694
ports-per-session statement....166
post-service-filter statement....501
usage guidelines....440
ppp-access-profile statement....313
usage guidelines....297
ppp-profile statement
usage guidelines....296
pre-shared-key statement....278
usage guidelines....227
preserve-interface statement....390
usage guidelines....331
primary statement
link services....391
usage guidelines....329
services PIC....502
usage guidelines....487
probe statement
RPM....1066
usage guidelines....1040
probe-count statement....1067
usage guidelines....1040
probe-interval statement....1067
usage guidelines....1040
probe-limit statement....1068
usage guidelines....1045
probe-server statement....1068
usage guidelines....1044
probe-type statement....1069
usage guidelines....1040
probes, for monitoring traffic....1040
procedural overview....39
profile statement
APPID
usage guidelines....695
application identification....716
profile-name statement....586
profile-version statement....587
proposal statement
IKE....279
usage guidelines....222
IPsec....280
usage guidelines....231
proposals statement
IKE....280
usage guidelines....227
IPsec....280
usage guidelines....234
protocol statement
applications....103
usage guidelines....62
IPsec....281
usage guidelines....218, 233

Q

queue-limit-percentage statement....587
queues statement....412
usage guidelines....397

R

random-allocation statement....166
rate statement....882
usage guidelines....792, 818
Real-Time Performance Monitoring *See* RPM
reassemble-packets statement....1109
usage guidelines....1091

receive-options-packets statement.....882
usage guidelines.....833
receive-ttl-exceeded statement.....883
usage guidelines.....833
receive-window statement.....314
usage guidelines.....298
reconnect statement.....588
red-differential-delay statement.....1027
usage guidelines.....990
redundancy
AS PIC.....487
flow monitoring.....837
L2TP.....302
redundancy-options statement.....391, 502
usage guidelines.....487
reflexive | reverse statement.....430
usage guidelines.....420
reject-all-commands-threshold statement.....588
reject-new-calls-threshold statement.....589
remote-gateway statement.....281
usage guidelines.....240
remote-id statement.....282
usage guidelines.....229
remotely-controlled statement.....167
report-service-change statement.....589
request-timestamp statement.....590
request-uri statement.....658
required-depth statement.....883
usage guidelines.....834
retransmit-interval statement.....314
usage guidelines.....298
retry statement.....925
usage guidelines.....904
retry-delay statement.....925
usage guidelines.....904
reverse-manipulation statement.....659
RFC 2890.....1089
route statement.....660
route-record statement
usage guidelines.....804
router identifier.....1111
routing-instance statement
BGF.....590
RPM.....1069
tunnel.....1110
usage guidelines.....1094
routing-instances statement
RPM.....1070
usage guidelines.....1040
rpc-program-number statement.....104
usage guidelines.....69
RPM.....1035, 1037
example configuration.....1051
rpm statement.....1070
usage guidelines.....1045
rtp statement.....412, 591
usage guidelines.....397
rule statement
AACL.....736
usage guidelines.....726
APPID
usage guidelines.....694
application identification.....717
BGF.....592
CoS.....431
usage guidelines.....416
IDS.....202
usage guidelines.....181
IPsec.....283
usage guidelines.....235
NAT.....168
usage guidelines.....133
stateful firewall.....123
usage guidelines.....108
rule-set statement
AACL.....737
usage guidelines.....729
APPID
usage guidelines.....694
application identification.....718
BGF.....592
CoS.....432
usage guidelines.....421
IDS.....203
usage guidelines.....187
IPsec.....284
usage guidelines.....242
NAT.....169
usage guidelines.....139
stateful firewall.....124
usage guidelines.....112
run-length statement.....884
usage guidelines.....792, 818

S

sample (firewall filter action).....791
sample-once statement.....884
usage guidelines.....794
sampled file.....796
sampled.pkts file.....794
sampling
logical interface.....792
monitoring interface.....799
sampling rate.....792
sampling statement
flow monitoring.....887
usage guidelines.....791
sbc-utils statement.....593, 661

scheduler map
CoS
configuration example....986
secondary statement
link services....392
usage guidelines....329
services PIC....503
usage guidelines....487
security associations
clearing....222
segmentation statement....594
send cflowd records to flow collector....904
send-notification-on-delay statement....595
server statement....1071
service filters....486
service interface configuration....440
service packages....35
service rules configuration....444
service sets
example configuration....455
overview....34
service statement....503
usage guidelines....485
service-change statement....596
service-change-type statement....597
service-class statement....662
service-domain statement....504
usage guidelines....441
service-filter statement
firewall
usage guidelines....486
interfaces....504
usage guidelines....440
service-interface statement....315, 469
BGF....597
border signaling gateway
gateway....663
service point....663
usage guidelines....297, 440
service-interface-pools statement....616
service-point statement....664
service-point-type statement....664
service-policies statement....665
service-port statement....962
usage guidelines....933
service-set statement....470, 505
usage guidelines....439, 485
service-state statement
virtual BGF....598
virtual interface in BGF....599
services configuration overview....39
services PICs....3
services statement
AACL
usage guidelines....725
APPID
usage guidelines....691
BGF....599
border signaling gateway....665
CoS....432
usage guidelines....415
DFC....962
usage guidelines....930
flow monitoring
usage guidelines....790
flow-monitoring....887
IDS....203
usage guidelines....179
interfaces....506
usage guidelines....483
IPsec....284
usage guidelines....213
L2TP....316
usage guidelines....299
NAT....169
usage guidelines....129
RPM....1071
usage guidelines....1037
service sets....472
usage guidelines....450
stateful firewall....124
usage guidelines....107
services-options statement....507
usage guidelines....299, 482, 483
session-limit statement....204
usage guidelines....184
session-mirroring statement....600
session-timeout statement....720, 721
usage guidelines....692
session-trace statement....666
shaping-rate statement
usage guidelines....339, 344, 348
shared-key statement....963
usage guidelines....933
short-sequence statement....1028
usage guidelines....983
signaling statement....667
SIP configuration....68
sip statement....669
sip-call-hold-timeout statement....104
usage guidelines....68
sip-header statement....671
sip-stack statement....673
sip-text statement....433
usage guidelines....420
sip-video statement....433
usage guidelines....420

sip-voice statement....434
usage guidelines....420
size statement....888
usage guidelines....796
snmp-command statement....105
usage guidelines....69
soft-limit statement....963
usage guidelines....932
soft-limit-clear statement....964
usage guidelines....932
SONET interfaces
sampling SONET interfaces....796
source statement
APPID
usage guidelines....694
application identification rule....721
encryption....778
tunnel....1111
usage guidelines....771, 1087
source-address statement
AACL....738
usage guidelines....727
BGF....600
CoS....434
usage guidelines....418
flow monitoring....889
usage guidelines....800
IDS....205
usage guidelines....183
IPsec....285
usage guidelines....237
NAT....170
usage guidelines....136
RPM....1072
usage guidelines....1040
stateful firewall....125
usage guidelines....109
tunnel....1112
tunnel services
usage guidelines....1097
source-address-range statement
AACL....738
usage guidelines....727
IDS....205
usage guidelines....183
NAT....170
usage guidelines....136
stateful firewall....125
usage guidelines....109
source-addresses statement
DFC....964
usage guidelines....933
source-pool statement....171
usage guidelines....137
source-port statement
BGF....601
RPM....105
usage guidelines....65
source-prefix statement....171, 206
usage guidelines....184
source-prefix-ipv6 statement....206
usage guidelines....184
source-prefix-list statement
AACL....739
usage guidelines....727
CoS....435
IDS....207
NAT....172
stateful firewall....126
usage guidelines....109
spi statement....285
usage guidelines....219
stamp option....795
stamp statement....889
usage guidelines....794
state-loss statement....601
stateful firewall
action statements....110
anomalies....42
applications....109
example configuration....112
match conditions....109
rules....112
stateful-firewall-rule-sets statement
usage guidelines....444
stateful-firewall-rules statement....474
usage guidelines....444
statement
flow monitoring
usage guidelines....796
IPsec
usage guidelines....248
L2TP
usage guidelines....303
services
usage guidelines....904
statistics statement
L-PDF....756
stop-detection-on-drop statement....602
support, technical *See* technical support
sustained-data-rate statement....603
gate in packet gateway....604
syn-cookie statement....207
usage guidelines....184
syntax conventions....liii
syslog statement
CoS....435
usage guidelines....419
flow monitoring....890

IDS....208
usage guidelines....184
interfaces....508
usage guidelines....483
IPsec....286
usage guidelines....239, 242
L2TP....318
usage guidelines....299
NAT....172
usage guidelines....137
service sets....474
usage guidelines....450
stateful firewall....126
usage guidelines....110

T

t391 statement....1028
usage guidelines....991
t392 statement....1029
usage guidelines....991
target statement....1072
target-url statement
usage guidelines....1040
tcp statement
RPM....1073
tcp-mss statement....475
technical support
contacting JTAC....liv
template statement
flow monitoring....891
template-refresh-rate statement....893
term statement
AACL....740
usage guidelines....726
border signaling gateway
new call usage policy....675
new transaction policy....676
service-class....677
CoS....436
usage guidelines....416
IDS....209
usage guidelines....181
IPsec....287
usage guidelines....235
NAT....173
usage guidelines....133
stateful firewall....127
usage guidelines....108
test statement
RPM....1074
usage guidelines....1040
test-interval statement....1075
usage guidelines....1040
then statement
AACL....741
usage guidelines....726
AACLl
usage guidelines....726
border signaling gateway
new call usage policy....678
new transaction policy....679
service-class....680
CoS....437
usage guidelines....416
IDS....211
usage guidelines....181
IPsec....288
usage guidelines....235
NAT....174
usage guidelines....133
stateful firewall....128
usage guidelines....108, 110
threshold statement....212
usage guidelines....184
thresholds statement
RPM....1076
usage guidelines....1040
time-to-live threshold....69
timer-c statement....680
timers statement....681
timerx statement....605
timestamp option....795
tmax-retransmission-delay statement....605
trace-options
server (tracing flag)....454
timer-events (tracing flag)....454
traceoptions statement....682
application identification....722
BGF....606
flow monitoring....893
IPsec....289
L-PDF....757
L2TP....319
services....476
tracing flags
event policy
all....454, 699
configuration....454
database....454
events....454
policy....454
server....454
timer-events....454
tracing operations
adaptive services....452
APPID....697
traffic....765
inbound (decryption)....769
IPsec, configuring....765

monitoring....1040
outbound (encryption)....767
traffic sampling
configuring....791
disabling....793, 848
example configurations....796
flow aggregation....804
FTP traffic....798
logging information output file....796
output files....794
SONET interfaces....796
traffic from single IP addresses....797
traffic-control-profiles statement
usage guidelines....344, 348
traffic-management statement....607
transactions statement....683
transfer statement....926
usage guidelines....902
transfer-log-archive statement....926
usage guidelines....903
translated statement....175
usage guidelines....137
translation-type statement....176
usage guidelines....137
transport statement
NAT....177
transport-details statement....684
traps statement....1077
usage guidelines....1040
trigger-link-failure statement....392
usage guidelines....326
trusted-ca statement....477
usage guidelines....447
ttl statement
DFC....965
usage guidelines....932
tunnel....1112
usage guidelines....1087
ttl-threshold statement....106
usage guidelines....69
tunnel interfaces
configuration statements....1087, 1092, 1094
dynamic tunnels....1097
example configuration....1098
logical tunnels....1092
loopback tunnels....1094
multicast tunnels....1092
PIM tunnels....1096
unicast tunnels....1087
tunnel statement....1113
encryption....779
usage guidelines....763
redundancy
usage guidelines....771
unicast
usage guidelines....1087
tunnel-group statement....321
usage guidelines....296
tunnel-mtu statement....290, 478
usage guidelines....242, 449
tunnel-timeout statement....322
usage guidelines....298
tunnel-type statement....1114
usage guidelines....1097
tunnels
definition....1083
GRE
fragmentation of....1090
key number....1089
interface types....1083
twamp statement....1078
twamp-server statement....1078
twice NAT....45
type statement....723
APPID
usage guidelines....692
type-of-service statement....723
APPID
usage guidelines....692

U

udp statement
RPM....1079
unicast tunnels....1087
unit statement
encryption....780
usage guidelines....763
flow monitoring....894
usage guidelines....791
interfaces....509
usage guidelines....479
link services....413, 1030
usage guidelines....973
tunnel....1115
usage guidelines....1087
Universal Unique Identifier....70
up statement
BGF....608
url statement....724
APPID
usage guidelines....697
use-lower-case statement....608
use-wildcard-response statement....609
username statement
flow collection....927
usage guidelines....903
uuid statement....106
usage guidelines....70

V

var/log/sampled file....796
var/tmp/sampled.pkts file....794
variant statement....927
    usage guidelines....902
version statement
    flow monitoring....895
    usage guidelines....804
version9 statement....896
virtual loopback tunnel
    configuration guidelines....1094
    VRF table lookup
        example configuration....1099
virtual-interface statement....609
virtual-interface-down statement....610
virtual-interface-indications statement....610
virtual-interface-up statement....611
voice services
    bundles....400
    configuration....395
    encapsulation....399
    example configuration....400
    interface type....396
voice services interfaces
    interleave fragments....398

W

warm standby
    AS PIC....487
    LSQ PIC....328
warm statement....611
warm-standby statement....393
warning (system logging severity
  level)....300, 451, 483
world-readable statement
    flow monitoring....898
    usage guidelines....794

Y

yellow-differential-delay statement....1031
    usage guidelines....990

Index of Statements and Commands

A

aacl-fields statement....750
aacl-statistics-profile statement....751
accelerations statement....622
accounting statement
    flow monitoring....842
acknowledge-retries statement....1011
acknowledge-timer statement....1012
action-red-differential-delay statement....1012
actions statement....621
adaptive-services-pics statement....457
address statement
    application rule....701
    DFC....949
    encryption....773
    flow monitoring....843
    interfaces....493
    link services....1013
    NAT....155
    voice services....405
address-range statement
    NAT....156
administrative statement
    BGF....518
admission-control statement....623, 624
aggregate-export-interval statement....843
aggregation statement....191
    flow monitoring....844
algorithm statement....519
allow-fragmentation statement....1103
allow-ip-options statement....118
allow-multicast statement....458
allowed-destinations statement....950
analyzer-address statement....911
analyzer-id statement....911
anti-replay-window-size statement....257, 459
application statement....97, 702, 703
application-aware-access-list-fields statement....752
application-data-inactivity-detection statement....520
application-group statement....703
application-group-any statement....732
application-groups statement....704, 732
application-profile statement....425
application-protocol statement....98
application-set statement....99
application-sets statement
    CoS....426
    IDS....192
    NAT....156
    stateful firewall....119
application-system-cache-timeout statement....704
applications statement
    AACL....731
    applications hierarchy....99
    applications identification....705
    CoS....426
    IDS....192
    NAT....157
    stateful firewall....119
archive-sites statement....912
audit-observed-events-returns statement....520
authentication statement....258
authentication-algorithm statement
    IKE....259
    IPsec....259
authentication-method statement....260
authentication-mode statement
    RPM....1055
automatic statement....705
autonomous-system-type statement....845
auxiliary-spi statement....260

B

backup-destination statement....774, 1104
backup-interface statement....774
backup-remote-gateway statement....261
base-root statement....521
bgf-core statement....522
blacklist-period statement....624
bundle statement....406, 1013
by-destination statement....193
by-pair statement....194
by-source statement....195

C

cancel-graceful statement....523
capture-group statement....951
cflowd statement....846

cisco-interoperability statement....383
cleanup-timeout statement....524
clear-dont-fragment-bit statement
GRE tunnel....494
IPsec....261
service-set....460
clear-ike-sas-on-pic-restart statement....262
clear-ipsec-sas-on-pic-restart statement....262
client-list statement....1056
collector statement....912
committed-burst-size statement....625
committed-information-rate statement....626
compression statement....406
compression-device statement....407
content-destination statement....952
context-indications statement....525
control-association-indications statement....526
control-source statement....953
controller-address statement....526
controller-failure statement....527
controller-port statement....527
copy-tos-to-outer-ip-header statement....1104
core-dump statement....847

D

data-fill statement....1056
data-format statement....913
data-inactivity-detection statement....528, 626
data-size statement....1057
datastore statement....627
default statement....529
default-media-realm statement....628
delivery-function statement....530
description statement
IKE....263
IPsec....263
destination statement
application identification rule....706
encryption....775
flow monitoring....848
link services....1014
tunnel....1105
destination-address statement
AACL....733
BGF....530
CoS....427
IDS....196
IPsec....263
NAT....157
stateful firewall....120
destination-address-range statement
AACL....733
IDS....196
NAT....158
stateful firewall....120
destination-networks statement
tunnel....1106
destination-pool statement....158
destination-port statement
applications....99
BGF....531
RPM....100, 1058
destination-prefix statement....159, 197
destination-prefix-ipv6 statement....197
destination-prefix-list statement
AACL....734
CoS....427
IDS....198
NAT....159
stateful firewall....121
destinations statement
flow collection....913
detect statement....531
dh-group statement....264
dial-options statement....495
dialogs statement....629
diffserv statement....532
direction statement....265
disable statement
application....707
application-group....707
flow monitoring....848
port mapping....708
disable-all-instances statement
flow monitoring....849
disable-global-timeout-override statement....708
disable-mlppp-inner-ppp-pfc statement....1014
disable-session-mirroring statement....532
disconnect statement....533
dlci statement....1015
do-not-fragment statement
tunnel....1106
down statement....533
download statement
APPID....709
drop-timeout statement....1016
dscp statement....428
BGF....534
BSG....630
dscp-code-point statement
RPM....1059
duplicates-dropped-periodicity statement....953
dynamic route insertion....244
dynamic statement....266
dynamic-flow-capture statement....954
dynamic-tunnels statement....1107

E

egress-service-point statement....630
embedded-spdf statement....631

encapsulation statement....407
link services....1017
encoding statement....534
encryption statement....267
encryption-algorithm statement
IKE....268
IPsec....268
engine-id statement
flow monitoring....849
engine-type statement....850
es-options statement....775
event-timestamp-notification statement....535
export-format statement....850

F

f-max-period statement....408
facility-override statement....309, 461, 496
failover statement....537
failover-cold statement....535
failover-warm statement....536
family statement
encryption....776
flow monitoring....851
interfaces....497
link services....1019
voice services....409
fast-update-filters statement....538
file statement....854
BGF....539
border signaling gateway....632
L-PDF statistics....753
traffic sampling....854
file-specification statement....915
filename statement....855
filename-prefix statement....914
files statement....855
filter statement
encryption....777
flow monitoring....856
flag statement....540, 633
flow-active-timeout statement....857
flow-collector statement....916
flow-export-destination statement....858
flow-inactive-timeout statement....858
flow-monitoring statement....859
flow-tap statement....955
flow-tap-lite statement....955
force-entry statement....198
forward-manipulation statement....634
forwarding-class statement....384, 428
forwarding-options statement....860
fragment-threshold statement
link services....1020
LSQ....385
voice services....410
fragmentation-map statement....385
fragmentation-maps statement....386
framework statement....635
from statement
AACL....734
border signaling gateway
new call usage policy....638
new transaction policy....639
service class....640
CoS....429
IDS....199
IPsec....269
NAT....160
stateful firewall....121
ftp statement
flow collection....918

G

g-duplicates-dropped-periodicity statement....956
g-max-duplicates statement....956
gateway statement
BGF....541
border signaling gateway....641
gateway-address statement....545
gateway-controller statement....546
gateway-port statement....547
graceful statement....548
graceful-restart statement....549

H

h248-options statement....550
h248-profile statement....551
h248-properties statement....552
h248-stack statement....555
h248-timers statement....556
hanging-termination-detection statement....556
hard-limit statement....957
hard-limit-target statement....957
hardware-timestamp statement....1060
hello-interval statement
L2TP....310
hello-timer statement
link services....1020
hide-avps statement....310
hint statement....161
history-size statement....1060
host statement....461, 498
L2TP....311
hot-standby statement....386

I

icmp-code statement....100
icmp-type statement....101

idle-timeout statement....709
ids-rules statement....462
ignore-entry statement....198
ike statement....270
ike-access-profile statement....462
inactivity-delay statement....557
inactivity-duration statement....557, 645
inactivity-non-tcp-timeout statement....710
inactivity-tcp-timeout statement....710
inactivity-timeout statement....101
BGF....558
flow monitoring....498
RPM....1061
inactivity-timer statement....558
index statement....711
initial-average-ack-delay statement....559
initiate-dead-peer-detection statement....271
input statement
flow monitoring....861
interfaces....499
input-interface-index statement....862
input-packet-rate-threshold statement....958
instance statement
port mirroring....863
sampling....864
interface statement
flow monitoring....865
flow-tap....958
service interface pool....615
interface-map statement....920
interface-service statement....463
interfaces statement
DFC....959
encryption....777
flow monitoring....867
interfaces hierarchy....499
link services....1021
tunnel....1107
voice services....410
interim-ah-scheme statement....559
interleave-fragments statement....1021
ip statement
application identification....711
ip-flow-stop-detection statement....560
ipsec statement....271
ipsec-inside-interface statement....272
ipsec-sa statement
encryption....778
ipsec-transport-security-association statement....560
ipsec-vpn-options statement....463
ipsec-vpn-rules statement....464
ipv4-template statement....867
ipv6-multicast-interfaces statement....161
ipv6-template statement....867

K

key statement
tunnel....1108

L

l2tp-access-profile statement....311
label-position statement....868
latch-deadlock-delay statement....561
learn-sip-register statement....102
lifetime-seconds statement
IKE....272
IPsec....272
link-layer-overhead statement....387
lmi-type statement....1022
local-certificate statement....273
local-dump statement....868
local-gateway address statement....312
local-gateway statement....464
local-id statement....273
local-policy-decision-function statement....754
log-prefix statement....465, 500
L2TP....312
logging statement....200, 465
logical-system statement
RPM....1061
lsq-failure-options statement....387

M

manipulation-rule statement....646
manual statement....274
match-direction statement
AACL....735
CoS....429
IDS....200
IPsec....274
NAT....162
stateful firewall....122
max-burst-size statement....562, 563
max-checked-bytes statement....712
max-concurrent-calls statement....564
max-duplicates statement....959
max-flows statement....466
max-packets-per-second statement....869
maximum-age statement....921
maximum-connections statement....1062
maximum-connections-per-client statement....1062
maximum-contexts statement....411
maximum-fuf-percentage statement....565
maximum-inactivity-time statement....566
maximum-net-propagation-delay statement....567
maximum-records-in-cache statement....648
maximum-send-window statement....313
maximum-sessions statement....1063
maximum-sessions-per-connection statement....1063

maximum-synchronization-mismatches statement....567
maximum-terms statement....568
maximum-time-in-cache statement....649
maximum-waiting-delay statement....568
media statement....569
media-policy statement....647
media-service statement....569
media-type statement....647
message-manipulation statement....648
message-manipulation-rules statement....650
mg-maximum-pdu-size statement....570
mg-originated-pending-limit statement....571
mg-provisional-response-timer-value statement....572
mg-segmentation-timer statement....573
mgc-maximum-pdu-size statement....574
mgc-originated-pending-limit statement....575
mgc-provisional-response-timer-value statement....576
mgc-segmentation-timer statement....577
min-checked-bytes statement....712
minimum statement
  BGF....651
minimum-links statement....1022
minimum-priority statement....960
mlfr-uni-nni-bundle-options statement....1023
mode statement....275
monitor statement....578
monitoring statement....870
moving-average-size statement....1064
mpls-ipv4-template statement....871
mpls-template statement....871
mrru statement....1024
mss statement....201
mtu statement....1025
multicast-dlci statement....1025
multicast-only statement....1108
multilink-class statement....388
multilink-max-classes statement....388
multiservice-options statement....872

N

n391 statement....1026
n392 statement....1026
n393 statement....1027
name-format statement....922
name-resolution-cache statement....652
nat-pool statement....578
nat-rules statement....466
nat-type statement....162
network-operator-id statement....579
new-call-usage-policies statement....652
new-call-usage-policy statement....653
new-call-usage-policy-set statement....654
new-transaction-policies statement....654
new-transaction-policy statement....655
new-transaction-policy-set statement....656
next-hop statement....872
  border signaling gateway....657
next-hop-group statement
  forwarding-options....873
  port mirroring....874
next-hop-service statement....467
no-anti-replay statement....275, 468
no-application-identification statement....713
no-application-system-cache statement....713
no-clear-application-system-cache statement....713
no-core-dump statement....847
no-dscp-bit-mirroring statement....579
no-filter-check statement....874
no-fragmentation statement....389
no-ipsec-tunnel-in-traceroute statement....276
no-local-dump statement....868
no-rtcp-check statement....580
no-signature-based statement....714
no-stamp statement....889
no-syslog statement
  DFC....960
  flow monitoring....890
no-termination-request statement....389
no-translation statement....163
no-world-readable statement
  flow monitoring....898
normal-mg-execution-time statement....580
normal-mgc-execution-time statement....581
notification-behavior statement....582
notification-rate-limit statement....582
notification-regulation statement....583
notification-targets statement....961

O

one-way-hardware-timestamp statement....1064
open-timeout statement....500
option-refresh-rate statement....875
order statement....714
output statement....501
  discard accounting....876
  flow monitoring....877
  port mirroring....877
  sampling....878
output-interface-index statement....879
overload-control statement....583
overload-pool statement....163
overload-prefix statement....164

P

passive-mode-tunneling statement....468
passive-monitor-mode statement....879
password statement
  flow collection....924

peak-data-rate statement.....584, 585
peer-unit statement
tunnel.....1109
per-unit-scheduler statement.....390
perfect-forward-secrecy statement.....276
pgcp statement
NAT.....164
pgcp-rules statement
service-set.....469
pic-memory-threshold statement.....961
platform statement.....586
policy statement
IKE.....277
policy-decision-statistics-profile statement.....755
pool statement.....165
service interface pool.....616
pop-all-labels statement.....880
port statement
flow monitoring.....880
NAT.....166
RPM.....1065
TWAMP.....1065
voice services.....411
port-mapping statement.....715
port-mirror statement.....881
port-range statement.....715
ports-per-session statement.....166
post-service-filter statement.....501
ppp-access-profile statement.....313
pre-shared-key statement.....278
preserve-interface statement.....390
primary statement
link services.....391
services PIC.....502
probe statement
RPM.....1066
probe-count statement.....1067
probe-interval statement.....1067
probe-limit statement.....1068
probe-server statement.....1068
probe-type statement.....1069
profile statement
application identification.....716
profile-name statement.....586
profile-version statement.....587
proposal statement
IKE.....279
IPsec.....280
proposals statement
IKE.....280
IPsec.....280
protocol statement
applications.....103
IPsec.....281

Q

queue-limit-percentage statement.....587
queues statement.....412

R

random-allocation statement.....166
rate statement.....882
reassemble-packets statement.....1109
receive-options-packets statement.....882
receive-ttl-exceeded statement.....883
receive-window statement.....314
reconnect statement.....588
red-differential-delay statement.....1027
redundancy-options statement.....391, 502
reflexive | reverse statement.....430
reject-all-commands-threshold statement.....588
reject-new-calls-threshold statement.....589
remote-gateway statement.....281
remote-id statement.....282
remotely-controlled statement.....167
report-service-change statement.....589
request-timestamp statement.....590
request-uri statement.....658
required-depth statement.....883
retransmit-interval statement.....314
retry statement.....925
retry-delay statement.....925
reverse-manipulation statement.....659
route statement.....660
routing-instance statement
BGF.....590
RPM.....1069
tunnel.....1110
routing-instances statement
RPM.....1070
rpc-program-number statement.....104
rpm statement.....1070
rtp statement.....412, 591
rule statement
AACL.....736
application identification.....717
BGF.....592
CoS.....431
IDS.....202
IPsec.....283
NAT.....168
stateful firewall.....123
rule-set statement
AACL.....737
application identification.....718
BGF.....592
CoS.....432
IDS.....203
IPsec.....284

NAT....169
stateful firewall....124
run-length statement....884

S

sample-once statement....884
sampling statement
flow monitoring....887
sbc-utils statement....593, 661
secondary statement
link services....392
services PIC....503
segmentation statement....594
send-notification-on-delay statement....595
server statement....1071
service statement....503
service-change statement....596
service-change-type statement....597
service-class statement....662
service-domain statement....504
service-filter statement
interfaces....504
service-interface statement....315, 469
BGF....597
border signaling gateway
gateway....663
service point....663
service-interface-pools statement....616
service-point statement....664
service-point-type statement....664
service-policies statement....665
service-port statement....962
service-set statement....470, 505
service-state statement
virtual BGF....598
virtual interface in BGF....599
services statement
BGF....599
border signaling gateway....665
CoS....432
DFC....962
flow-monitoring....887
IDS....203
interfaces....506
IPsec....284
L2TP....316
NAT....169
RPM....1071
service sets....472
stateful firewall....124
services-options statement....507
session-limit statement....204
session-mirroring statement....600
session-timeout statement....720, 721
session-trace statement....666
shared-key statement....963
short-sequence statement....1028
signaling statement....667
sip statement....669
sip-call-hold-timeout statement....104
sip-header statement....671
sip-stack statement....673
sip-text statement....433
sip-video statement....433
sip-voice statement....434
size statement....888
snmp-command statement....105
soft-limit statement....963
soft-limit-clear statement....964
source statement
application identification rule....721
encryption....778
tunnel....1111
source-address statement
AACL....738
BGF....600
CoS....434
flow monitoring....889
IDS....205
IPsec....285
NAT....170
RPM....1072
stateful firewall....125
tunnel....1112
source-address-range statement
AACL....738
IDS....205
NAT....170
stateful firewall....125
source-addresses statement
DFC....964
source-pool statement....171
source-port statement
BGF....601
RPM....105
source-prefix statement....171, 206
source-prefix-ipv6 statement....206
source-prefix-list statement
AACL....739
CoS....435
IDS....207
NAT....172
stateful firewall....126
spi statement....285
stamp statement....889
state-loss statement....601
stateful-firewall-rules statement....474
statistics statement
L-PDF....756
stop-detection-on-drop statement....602

sustained-data-rate statement....603
gate in packet gateway....604
syn-cookie statement....207
syslog statement
CoS....435
flow monitoring....890
IDS....208
interfaces....508
IPsec....286
L2TP....318
NAT....172
service sets....474
stateful firewall....126

T

t391 statement....1028
t392 statement....1029
target statement....1072
tcp statement
RPM....1073
tcp-mss statement....475
template statement
flow monitoring....891
template-refresh-rate statement....893
term statement
AACL....740
border signaling gateway
new call usage policy....675
new transaction policy....676
service-class....677
CoS....436
IDS....209
IPsec....287
NAT....173
stateful firewall....127
test statement
RPM....1074
test-interval statement....1075
then statement
AACL....741
border signaling gateway
new call usage policy....678
new transaction policy....679
service-class....680
CoS....437
IDS....211
IPsec....288
NAT....174
stateful firewall....128
threshold statement....212
thresholds statement
RPM....1076
timer-c statement....680
timers statement....681
timerx statement....605
tmax-retransmission-delay statement....605
traceoptions statement....682
application identification....722
BGF....606
flow monitoring....893
IPsec....289
L-PDF....757
L2TP....319
services....476
traffic-management statement....607
transactions statement....683
transfer statement....926
transfer-log-archive statement....926
translated statement....175
translation-type statement....176
transport statement
NAT....177
transport-details statement....684
traps statement....1077
trigger-link-failure statement....392
trusted-ca statement....477
ttl statement
DFC....965
tunnel....1112
ttl-threshold statement....106
tunnel statement....1113
encryption....779
tunnel-group statement....321
tunnel-mtu statement....290, 478
tunnel-timeout statement....322
tunnel-type statement....1114
twamp statement....1078
twamp-server statement....1078
type statement....723
type-of-service statement....723

U

udp statement
RPM....1079
unit statement
encryption....780
flow monitoring....894
interfaces....509
link services....413, 1030
tunnel....1115
up statement
BGF....608
url statement....724
use-lower-case statement....608
use-wildcard-response statement....609
username statement
flow collection....927
uuid statement....106

V

variant statement....927
version statement
    flow monitoring....895
version9 statement....896
virtual-interface statement....609
virtual-interface-down statement....610
virtual-interface-indications statement....610
virtual-interface-up statement....611

W

warm statement....611
warm-standby statement....393
world-readable statement
    flow monitoring....898

Y

yellow-differential-delay statement....1031