



JUNOS® Software

Services Interfaces Configuration Guide

Release 9.3

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Part Number: 530-027201-01, Revision 1

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

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

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

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

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

Juniper Networks, 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.

JUNOS® Software Services Interfaces Configuration Guide
Release 9.3

Copyright © 2008, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Alan Twigg, Sarah Lesway-Ball, Justine Kangas, Myron Weintraub
Editing: Sonia Saruba, Benjamin Mann
Illustration: Nathaniel Woodward
Cover Design: Edmonds Design

Revision History
10 October 2008—Revision 1

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 présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés 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

xliii

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	27
Chapter 4	Applications Configuration Guidelines	53
Chapter 5	Summary of Applications Configuration Statements	91
Chapter 6	Stateful Firewall Services Configuration Guidelines	101
Chapter 7	Summary of Stateful Firewall Configuration Statements	111
Chapter 8	Network Address Translation Services Configuration Guidelines	123
Chapter 9	Summary of Network Address Translation Configuration Statements	145
Chapter 10	Intrusion Detection Service Configuration Guidelines	169
Chapter 11	Summary of Intrusion Detection Service Configuration Statements	181
Chapter 12	IPSec Services Configuration Guidelines	203
Chapter 13	Summary of IPSec Services Configuration Statements	245
Chapter 14	Layer 2 Tunneling Protocol Services Configuration Guidelines	279
Chapter 15	Summary of Layer 2 Tunneling Protocol Configuration Statements	297
Chapter 16	Link Services IQ Interfaces Configuration Guidelines	311
Chapter 17	Summary of Link Services IQ Configuration Statements	371
Chapter 18	Voice Services Configuration Guidelines	383
Chapter 19	Summary of Voice Services Configuration Statements	397
Chapter 20	Class-of-Service Configuration Guidelines	409
Chapter 21	Summary of Class-of-Service Configuration Statements	419
Chapter 22	Service Set Configuration Guidelines	433
Chapter 23	Summary of Service Set Configuration Statements	451
Chapter 24	Interface Configuration Guidelines	467
Chapter 25	Summary of Interface Configuration Statements	481
Chapter 26	Packet Gateway Configuration Guidelines	499
Chapter 27	Summary of Packet Gateway Configuration Statements	521
Chapter 28	Service Interface Pools Configuration Guidelines	603
Chapter 29	Summary of Service Interface Pools Statements	605

Part 3	Data Link Switching	
Chapter 30	Data Link Switching Overview	609
Chapter 31	Data Link Switching Configuration Guidelines	611
Chapter 32	Summary of Data Link Switching Configuration Statements	623
Part 4	Encryption Services	
Chapter 33	Encryption Overview	645
Chapter 34	Encryption Interfaces Configuration Guidelines	647
Chapter 35	Summary of Encryption Configuration Statements	657
Part 5	Flow Monitoring and Discard Accounting Services	
Chapter 36	Flow Monitoring and Discard Accounting Overview	667
Chapter 37	Flow Monitoring and Discard Accounting Configuration Guidelines	671
Chapter 38	Summary of Flow-Monitoring Configuration Statements	719
Chapter 39	Flow Collection Configuration Guidelines	775
Chapter 40	Summary of Flow Collection Configuration Statements	787
Chapter 41	Dynamic Flow Capture Configuration Guidelines	805
Chapter 42	Flow-Tap Configuration Guidelines	817
Chapter 43	Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements	823
Part 6	Link and Multilink Services	
Chapter 44	Link and Multilink Services Overview	841
Chapter 45	Link and Multilink Services Configuration Guidelines	845
Chapter 46	Summary of Link Services Configuration Statements	885
Part 7	Real-Time Performance Monitoring Services	
Chapter 47	Real-Time Performance Monitoring Services Overview	907
Chapter 48	Real-Time Performance Monitoring Configuration Guidelines	909
Chapter 49	Summary of Real-Time Performance Monitoring Configuration Statements	927
Part 8	Tunnel Services	
Chapter 50	Tunnel Services Overview	955
Chapter 51	Tunnel Interfaces Configuration Guidelines	959
Chapter 52	Summary of Tunnel Services Configuration Statements	975
Part 9	Index	
	Index	991
	Index of Statements and Commands	1009

Table of Contents

	About This Guide	xliii
	Objectives	xliii
	Audience	xliii
	Supported Routing Platforms	xliv
	Using the Indexes	xliv
	Using the Examples in This Manual	xliv
	Merging a Full Example	xlvi
	Merging a Snippet	xlvi
	Documentation Conventions	xlvi
	List of Technical Publications	xlvi
	Documentation Feedback	lv
	Requesting Technical Support	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 protocols] Hierarchy Level	11
	[edit services] Hierarchy Level	12

Part 2**Adaptive Services**

Chapter 3**Adaptive Services Overview****27**

Enabling Service Packages	29
Layer 2 Service Package Capabilities and Interfaces	32
Services Configuration Procedure	33
Packet Flow Through the Adaptive Services or MultiServices PIC	34
Stateful Firewall Overview	35
Stateful Firewall Support for Application Protocols	36
Stateful Firewall Anomaly Checking	36
Network Address Translation Overview	38
Traditional NAT	38
Twice NAT	39
IPSec Overview	39
IPSec	40
Security Associations	40
IKE	40
Comparison of IPSec Services and ES Interface Configuration	41
Layer 2 Tunneling Protocol Overview	42
Voice Services Overview	42
Class of Service Overview	43
Packet Overview	43
Examples: Services Interfaces Configuration	43

Chapter 4**Applications Configuration Guidelines****53**

Configuring Application Protocol Properties	54
Configuring an Application Protocol	54
Configuring the Network Protocol	56
Configuring the ICMP Code and Type	57
Configuring Source and Destination Ports	59
Configuring the Inactivity Timeout Period	62
Configuring SIP	62
Configuring an SNMP Command	63
Configuring an RPC Program Number	63
Configuring the TTL Threshold	63
Configuring a Universal Unique Identifier	64
Configuring Application Sets	64
ALG Descriptions	64
Basic TCP ALG	65
Basic UDP ALG	65
BOOTP	66
DCE RPC Services	66
FTP	66
H323	67
ICMP	67
IIOP	67
NetShow	68

RealAudio	68
RPC and RPC Portmap Services	68
RTSP	69
SMB	70
SNMP	70
SQLNet	70
TFTP	70
Traceroute	71
UNIX Remote-Shell Services	71
WinFrame	71
Verification	72
FTP Example	72
Sample Output	72
FTP System Log Messages	73
Analysis	74
Troubleshooting Questions	74
RTSP ALG Example	75
Sample Output	75
Analysis	75
Troubleshooting Questions	76
System Log Messages	77
System Log Configuration	78
System Log Output	79
JUNOS Default Groups	79
Examples: Referencing the Preset Statement from the JUNOS Default Group	85
Examples: Configuring Applications	87

Chapter 5

Summary of Applications Configuration Statements	91
application	91
application-protocol	92
application-set	93
applications	93
destination-port	94
icmp-code	94
icmp-type	95
inactivity-timeout	95
learn-sip-register	96
protocol	97
rpc-program-number	98
sip-call-hold-timeout	98
snmp-command	99
source-port	99
ttl-threshold	100
uuid	100

Chapter 6	Stateful Firewall Services Configuration Guidelines	101
	Configuring Stateful Firewall Properties	102
	Configuring the Stateful Firewall Rule Set	102
	Configuring Stateful Firewall Rule Content	102
	Configuring Match Direction	103
	Configuring Stateful Firewall Match Conditions	103
	Configuring Stateful Firewall Actions	104
	Examples: Configuring Stateful Firewall Properties	106
Chapter 7	Summary of Stateful Firewall Configuration Statements	111
	allow-ip-option	112
	application-sets	113
	applications	113
	destination-address	114
	destination-address-range	114
	destination-prefix-list	115
	from	115
	match-direction	116
	rule	117
	rule-set	118
	services	118
	source-address	119
	source-address-range	119
	source-prefix-list	120
	syslog	120
	term	121
	then	122
Chapter 8	Network Address Translation Services Configuration Guidelines	123
	Configuring Network Address Translation Properties	124
	Configuring Address and Port Information	124
	Configuring a Pool	124
	Specifying the Destination or Source Prefix	126
	Addressing Information	126
	Configuring IPv6 Multicast Filters	127
	Configuring the NAT Rule Set	127
	Configuring NAT Rule Content	128
	Configuring NAT Match Direction	129
	Configuring NAT Type	129
	Configuring NAT Match Conditions	130
	Configuring NAT Actions	131
	Examples: Configuring Network Address Translation Properties	133
	Dynamic Source Translation	134
	Static Source Translation	134
	Dynamic and Static Source Translation	134
	Oversubscribed Pool with No Fallback	135
	Oversubscribed Pool with a Fallback to NAT	136

Multiple Prefixes and Address Ranges in Static Source Translation	136
Assigning Addresses from a Dynamic Pool for Static Use	137
Configuring NAT Rules Without Defining a Pool	137
Configuring Specific Addresses That Are Not to Be Translated	138
Configuring Network Address Translation for Multicast Traffic	139
Rendezvous Point Configuration	139
Router 1 Configuration	142
Configuring Twice NAT	143
Configuring Full-Cone NAT	143

Chapter 9

Summary of Network Address Translation Configuration Statements

145

address	145
address-range	146
application-sets	146
applications	147
destination-address	147
destination-address-range	148
destination-pool	148
destination-prefix	149
destination-prefix-list	149
from	150
hint	151
ipv6-multicast-interfaces	151
match-direction	152
nat-type	152
no-translation	153
overload-pool	153
overload-prefix	154
pgcp	154
pool	155
port	156
ports-per-session	156
remotely-controlled	157
rule	158
rule-set	159
services	159
source-address	160
source-address-range	160
source-pool	161
source-prefix	161
source-prefix-list	162
syslog	162
term	163
then	164

translated	165
translation-type	166
translation-type (Traditional NAT)	166
translation-type (Twice NAT)	166
transport	167

Chapter 10**Intrusion Detection Service Configuration Guidelines 169**

Configuring Intrusion Detection Properties	171
Configuring the IDS Rule Set	171
Configuring IDS Rule Content	171
Configuring Match Direction	173
Configuring IDS Match Conditions	173
Configuring IDS Actions	174
Examples: Configuring Intrusion Detection Properties	178

Chapter 11**Summary of Intrusion Detection Service Configuration Statements 181**

aggregation	181
application-sets	182
applications	182
by-destination	183
by-pair	184
by-source	185
destination-address	186
destination-address-range	186
destination-prefix	187
destination-prefix-ipv6	187
destination-prefix-list	188
force-entry	188
from	189
ignore-entry	189
logging	190
match-direction	190
mss	191
rule	192
rule-set	193
services	193
session-limit	194
source-address	195
source-address-range	195
source-prefix	196
source-prefix-ipv6	196
source-prefix-list	197
syn-cookie	197
syslog	198
term	199
then	201
threshold	202

Chapter 12**IPSec Services Configuration Guidelines****203**

Minimum Security Association Configurations	205
Manual SA Configuration	205
Dynamic SA Configuration	205
Configuring Security Associations	206
Configuring Manual Security Associations	207
Configuring Direction	207
Configuring the Protocol	208
Configuring the Security Parameter Index	209
Configuring the Auxiliary Security Parameter Index	209
Configuring Authentication	209
Configuring Encryption	210
Configuring Dynamic Security Associations	211
Clearing Security Associations	211
Configuring an IKE Proposal	212
Configuring an IKE Authentication Algorithm	212
Configuring an IKE Authentication Method	213
Configuring an IKE Diffie-Hellman Group	213
Configuring an IKE Encryption Algorithm	213
Configuring the Lifetime for an IKE SA	214
Example: Configuring an IKE Proposal	214
Configuring an IKE Policy	215
Configuring the IKE Policy Mode	216
Configuring IKE Policy Proposals	216
Configuring a Preshared Key	216
Configuring a Local Certificate	217
Configuring a Certificate Revocation List	217
Configuring an IKE Policy Description	218
Configuring Local and Remote IDs	218
Example: Configuring an IKE Policy	219
Configuring an IPSec Proposal	220
Configuring an Authentication Algorithm	220
Configuring an IPSec Proposal Description	220
Configuring an Encryption Algorithm	220
Configuring the Lifetime for an IPSec SA	221
Configuring the Protocol for the Dynamic SA	222
Configuring an IPSec Policy	222
Configuring an IPSec Policy Description	223
Configuring Perfect Forward Secrecy	223
Configuring IPSec Policy Proposals	223
Example: IPSec Policy Configuration	223
Configuring IPSec Service Rules	224
Configuring the IPSec Rule Set	224
Configuring IPSec Rule Content	225
Configuring IPSec Match Conditions	226
Configuring IPSec Actions	227
Configuring Dynamic Endpoints	230
Authentication Process	231
Dynamic Implicit Rules	231
Reverse Route Insertion	232

Configuring an IKE Access Profile	232
Configuring the Service Set	234
Configuring the Interface Identifier	234
Default IKE and IPSec Proposals	235
Tracing IPSec Operations	236
Examples: Configuring IPSec Services	236
Statically Assigned Tunnels	236
Dynamically Assigned Tunnels	239

Chapter 13**Summary of IPSec Services Configuration Statements 245**

authentication	245
authentication-algorithm	246
authentication-algorithm (IKE)	246
authentication-algorithm (IPSec)	246
authentication-method	247
auxiliary-spi	247
backup-remote-gateway	248
clear-dont-fragment-bit	248
clear-ike-sas-on-pic-restart	249
clear-ipsec-sas-on-pic-restart	249
description	250
destination-address	250
dh-group	251
direction	252
dynamic	253
encryption	254
encryption-algorithm	255
from	256
ike	257
initiate-dead-peer-detection	258
ipsec	258
ipsec-inside-interface	259
lifetime-seconds	259
local-certificate	260
local-id	260
manual	261
match-direction	261
mode	262
no-anti-replay	262
perfect-forward-secrecy	263
policy	264
policy (IKE)	264
policy (IPSec)	265
pre-shared-key	265
proposal	266
proposal (IKE)	266
proposal (IPSec)	267
proposals	267
protocol	268

remote-gateway	268
remote-id	269
rule	270
rule-set	271
services	271
source-address	272
spi	272
syslog	273
term	274
then	275
traceoptions	276
tunnel-mtu	277

Chapter 14**Layer 2 Tunneling Protocol Services Configuration Guidelines 279**

L2TP Services Components	280
L2TP Minimum Configuration	281
Configuring L2TP Group Properties	283
Configuring a Tunnel Group	284
Configuring Access Profiles	284
Configuring Addressing	285
Configuring Window Size	285
Configuring Timers	286
Hiding Attribute-Value Pairs	286
Configuring System Log Properties	286
Configuring the Logical Interface Identifier	288
Configuring Multilink on Shared Interfaces	288
L2TP Redundancy	289
Tracing L2TP Operations	290
Examples: Configuring L2TP Services	292

Chapter 15**Summary of Layer 2 Tunneling Protocol Configuration Statements 297**

facility-override	297
hello-interval	298
hide-avps	298
host	299
l2tp-access-profile	299
local-gateway address	300
log-prefix	300
maximum-send-window	301
ppp-access-profile	301
receive-window	302
retransmit-interval	302
service-interface	303
services	304
services (Hierarchy)	304
services (L2TP System Logging)	305
syslog	306

traceoptions	307
tunnel-group	309
tunnel-timeout	310

Chapter 16**Link Services IQ Interfaces Configuration Guidelines 311**

Layer 2 Service Package Capabilities and Interfaces	312
LSQ Failure Recovery	313
Interchassis LSQ Failover	314
Configuring Failover Options	314
Configuring Interoperability	315
Configuration Restrictions	315
Stateless Intrachassis LSQ Failover	316
Stateful Intrachassis LSQ Failover	316
Configuring Intrachassis LSQ Redundancy	317
LSQ Failure Conditions and Constraints	317
Link PIC Redundancy	319
Examples: Configuring Intrachassis LSQ Redundancy	320
Link Services IQ CoS Components	324
Scheduler Buffer Size	325
Scheduler Priority	326
Scheduler Shaping Rate	326
Drop Profiles	326
Configuring Fragmentation by Forwarding Class	328
Configuring Link-Layer Overhead	329
Configuring Multiclass MLPPP	330
Oversubscribing Interface Bandwidth	331
Example: Oversubscribing an LSQ Interface	334
Providing a Guaranteed Minimum Rate	335
Example: Providing a Guaranteed Minimum Rate	337
Configuring Link Services and CoS	338
Link Services CoS on J-series Services Routers	341
Common Uses for the Link Services IQ Interface	342
Configuring an NxT1 Bundle Using MLPPP	343
Example: Configuring an NxT1 Bundle Using MLPPP	346
Configuring an NxT1 Bundle Using FRF.16	348
Examples: Configuring an NxT1 Bundle Using FRF.16	351
Configuring a Single Fractional T1 Interface Using MLPPP and LFI	353
Example: Configuring a Single Fractional T1 Using MLPPP and LFI	356
Configuring a Single Fractional T1 Interface Using FRF.12	357
Examples: Configuring a Single Fractional T1 Interface Using FRF.12	360
Configuring an NxT1 Bundle Using FRF.15	364
Configuring a T3 Link for Compressed RTP over MLPPP	365
Configuring a T3 or OC3 Interface Using FRF.12	366
Configuring an ATM2 IQ Interface Using MLPPP	368

Chapter 17	Summary of Link Services IQ Configuration Statements	371
	cisco-interopability	371
	forwarding-class	372
	fragment-threshold	373
	fragmentation-map	373
	fragmentation-maps	374
	hot-standby	374
	link-layer-overhead	375
	lsq-failure-options	375
	multilink-class	376
	multilink-max-classes	376
	no-fragmentation	377
	no-termination-request	377
	per-unit-scheduler	378
	preserve-interface	378
	primary	379
	redundancy-options	379
	secondary	380
	trigger-link-failure	380
	warm-standby	381
Chapter 18	Voice Services Configuration Guidelines	383
	Configuring Voice Services Properties	385
	Configuring the Interface Address	385
	Configuring Compression	386
	Configuring Delay-Sensitive Packet Interleaving	387
	Example: Configuring Compression	387
	Configuring Encapsulation	388
	Configuring the Network Interface	389
	Configuring Voice Services Bundles	389
	Configuring the Compression Interface	389
	Configuring VoIP Routing on J-series Services Routers	390
	Functional Components	390
	Configuring the VoIP Interface	390
	Configuring the Media Gateway Controller List	391
	Configuring Dynamic Call Admission Control	392
	Examples: Configuring Voice Services	393
Chapter 19	Summary of Voice Services Configuration Statements	397
	activation-priority	397
	address	398
	bearer-bandwidth-limit	399
	bundle	399
	compression	400
	compression-device	400
	dynamic-call-admission-control	401
	encapsulation	402

f-max-period	402
family	403
fragment-threshold	404
interfaces	404
maximum-contexts	405
port	405
queues	406
rtp	406
unit	407

Chapter 20**Class-of-Service Configuration Guidelines 409**

Configuring Class-of-Service Properties	410
Configuring the Class-of-Service Rule Set	410
Configuring Class-of-Service Rule Content	410
Configuring Class-of-Service Match Direction	411
Configuring Class-of-Service Match Conditions	412
Configuring Class-of-Service Actions	413
CoS Configuration Restrictions	415
Output Packet Rewriting	415
CoS Hierarchy Configuration Examples	415
Examples: Configuring Class-of-Service Properties	417

Chapter 21**Summary of Class-of-Service Configuration Statements 419**

application-profile	419
application-sets	420
applications	420
destination-address	421
destination-prefix-list	421
dscp	422
forwarding-class	422
from	423
match-direction	423
(reflexive reverse)	424
rule	425
rule-set	426
services	426
sip-text	427
sip-video	427
sip-voice	428
source-address	428
source-prefix-list	429
syslog	429
term	430
then	431

Chapter 22 **Service Set Configuration Guidelines** **433**

Configuring Service Sets	434
Configuring Services Interfaces	434
Configuring an Interface Service Set	434
Configuring a Next-Hop Service Set	435
Determining Traffic Direction	436
Configuring Service Rules	438
Configuring IPSec Options	438
Configuring the Local Gateway Address	439
Configuring an IKE Access Profile	440
Configuring Certification Authorities	440
Configuring the Maximum Number of Flows	441
Configuring System Log Properties	441
Allowing Multicast Traffic	442
Service Set Extensions for JUNOS SDK Applications	443
Interface Service Sets for JUNOS SDK Applications	443
Next-Hop Service Sets for JUNOS SDK Applications	444
Limitations and Constraints for SDK Service Sets	445
Enabling JUNOS SDK Services	445
Tracing Adaptive Services or MultiServices PIC Operations	446
Configuring the Adaptive Services Log Filename	447
Configuring the Number and Size of Adaptive Services Log Files	447
Configuring Access to the Log File	448
Configuring a Regular Expression for Lines to Be Logged	448
Configuring the Trace Operations	448
Example: Configuring Service Sets	449

Chapter 23 **Summary of Service Set Configuration Statements** **451**

adaptive-services-pics	451
allow-multicast	452
extension-service	452
facility-override	453
host	453
ids-rules	454
ike-access-profile	454
interface-service	455
ipsec-vpn-options	455
ipsec-vpn-rules	456
local-gateway	456
log-prefix	457
logging	457
max-flows	458
nat-rules	458
next-hop-service	459
pgcp-rules	460
service-interface	460

service-set	461
services	462
services (Hierarchy)	462
services (System Logging)	463
stateful-firewall-rules	463
syslog	464
traceoptions	465
trusted-ca	466

Chapter 24**Interface Configuration Guidelines 467**

Naming Services Interfaces	468
Configuring Interface Properties	469
Configuring the Interface Address and Domain	470
Configuring Default Timeout Settings	470
Configuring Default System Log Properties	471
Enabling Fragmentation on GRE Tunnels	472
Applying Filters and Services to an Interface	473
Configuring Service Filters	474
Configuring AS or MultiServices PIC Redundancy	475
Examples: Configuring a Services Interface	478

Chapter 25**Summary of Interface Configuration Statements 481**

address	481
clear-dont-fragment-bit	482
dial-options	483
facility-override	484
family	485
host	486
inactivity-timeout	486
input	487
interfaces	487
log-prefix	488
open-timeout	488
output	489
post-service-filter	489
primary	490
redundancy-options	490
secondary	491
service	491
service-domain	492
service-filter	492
service-set	493
services	494
services-options	495
syslog	496
unit	497

Chapter 26	Packet Gateway Configuration Guidelines	499
	Configuring Virtual Packet Gateways	503
	Configuring a Connection to the Packet Gateway Controller	506
	Configuring NAT Pools for the Packet Gateway	507
	Configuring a Media Service	507
	Configuring a Virtual Interface	507
	Configuring H.248 Timers	508
	Configuring Default Values for H.248 Properties	508
	Limiting FUF Terms on VPGs	510
	Limiting the Rate of Messages Sent from the PIC	511
	Enabling Application Layer Gateways for RTP and RTCP for Media Flows	511
	Configuring Graceful Restart	511
	Detecting Latch Deadlocks and Inactivity Delays	511
	Configuring the Service State of a VPG or a Virtual Interface	512
	Configuring the Service State of a VPG	512
	Configuring the Service State of a Virtual Interface	512
	Configuring the Packet Gateway Rule Content	513
	Configuring the Packet Gateway Rule Set	513
	Configuring Session Mirroring	514
	Configuring Overload Control	514
	Enabling Wildcards for Service Changes	514
	Enabling a History of Media Inactivity Notifications	514
	Disabling Bit Mirroring of DSCP Marking	515
	Configuring ServiceChange Method and Reason	515
	Tracing PGCP Operations	516
	Example: Configuring a Packet Gateway	516
 Chapter 27	 Summary of Packet Gateway Configuration Statements	 521
	administrative	522
	administrative (Control Association Down)	522
	administrative (Virtual Interface Down)	523
	algorithm	523
	application-data-inactivity-detection	524
	audit-observed-events-returns-history	524
	base-root	525
	cancel-graceful	526
	cancel-graceful (Control Association Up)	526
	cancel-graceful (Virtual Interface Up)	527
	cleanup-timeout	527
	context-indications	528
	control-association-indications	529
	controller-address	530
	controller-failure	530
	controller-port	531
	data-inactivity-detection	531
	delivery-function	532
	destination-address	532
	destination-port	533

detect	533
diffserv	534
disable-session-mirroring	534
disconnect	535
down	536
dscp	537
encoding	537
event-timestamp-notification	538
failover-cold	538
failover-warm	539
failure	540
failure (Control Association Down)	540
failure (Virtual Interface Down)	541
fast-update-filters	541
gateway	542
gateway-address	545
gateway-controller	546
gateway-port	546
graceful	547
graceful (Control Association Down)	547
graceful (Virtual Interface Down)	548
graceful-restart	548
h248-options	549
h248-properties	551
h248-timers	552
hanging-termination-detection	553
inactivity-timeout	553
inactivity-delay	554
inactivity-duration	554
inactivity-timer	555
initial-average-ack-delay	555
interface	556
interim-ah-scheme	556
ip-flow-stop-detection	557
latch-deadlock-delay	557
link-loss	558
max-burst-size	559
max-burst-size (All Streams)	559
max-burst-size (RTCP Streams)	560
max-concurrent-calls	561
maximum-fuf-percentage	562
maximum-inactivity-time	562
maximum-net-propagation-delay	563
maximum-synchronization-mismatches	563
maximum-synchronization-time	564
maximum-terms	564
maximum-waiting-delay	565
media	565
media-service	566
mg-maximum-pdu-size	566
mg-originated-pending-limit	567

mg-provisional-response-timer-value	567
mg-segmentation-timer	568
mgc-maximum-pdu-size	568
mgc-originated-pending-limit	569
mgc-provisional-response-timer-value	569
mgc-segmentation-timer	570
monitor	570
nat-pool	571
network-operator-id	571
normal-mg-execution-time	572
normal-mgc-execution-time	572
notification-behavior	573
notification-rate-limit	573
notification-regulation	574
no-dscp-bit-mirroring	574
overload-control	575
peak-data-rate	576
peak-data-rate (All Streams)	576
peak-data-rate (RTCP Streams)	577
queue-limit-percentage	578
reconnect	578
report-service-change	579
request-timestamp	579
routing-instance	580
rtcp	580
rtp	581
rule	581
rule-set	582
segmentation	582
send-notification-on-delay	583
service-change	584
service-change-type	585
service-interface	585
service-state	586
service-state (VPG)	586
service-state (Virtual Interface)	587
services	587
session-mirroring	588
source-address	588
source-port	589
state-loss	589
stop-detection-on-drop	590
sustained-data-rate	591
sustained-data-rate (All Streams)	591
sustained-data-rate (RTCP Streams)	592
timerx	593
traffic-management	594
tmax-retransmission-delay	595
traceoptions	596
up	597
virtual-interface	598

	virtual-interface-down	599
	virtual-interface-indications	600
	virtual-interface-up	601
	warm	601
	wildcard-response-service-change	602
Chapter 28	Service Interface Pools Configuration Guidelines	603
	Configuring Service Interface Pools	603
Chapter 29	Summary of Service Interface Pools Statements	605
	interface	605
	pool	606
	service-interface-pools	606
Part 3	Data Link Switching	
Chapter 30	Data Link Switching Overview	609
	Overview	609
	DLSw Standards	609
Chapter 31	Data Link Switching Configuration Guidelines	611
	Configuring DLSw Properties	612
	Minimum DLSw Configuration	612
	Configuring the Remote Peer	612
	Configuring Load Balancing	613
	Configuring DLSw Timers	613
	Configuring the Local Peer	614
	Examples: Configuring DLSw Peers	614
	Configuring the Initial Pacing Window	615
	Configuring the Idle Timeout	615
	Configuring the Multicast Address	615
	Configuring Class of Service	615
	Example: Configuring CoS for a DLSw Connection	616
	Tracing DLSw Protocol Traffic	617
	Configuring Logical Link Control on an Interface	617
	Example: Configuring LLC Options on an Interface	618
	Configuring DLSw Ethernet Redundancy Using LLC2 Properties	618
	Example: Configuring DLSw Ethernet Redundancy	620

Chapter 32 Summary of Data Link Switching Configuration Statements 623

advertise-interval	623
circuit-weight	624
connection-idle-timeout	624
cost	625
destination	625
destination-interface	626
dls w	626
dls w-cos	627
explorer-wait-time	627
hold-time	628
interface	629
load-balance	629
local-mac	630
local-peer	630
map	631
multicast-address	631
no-preempt	631
peer	632
preempt	633
priority	634
promiscuous	634
protocols	635
reachability-cache-timeout	635
receive-initial-pacing	636
redundancy-group	637
remote-mac	637
remote-peer	638
traceoptions	639
track	641
type-of-service	642

Part 4 Encryption Services

Chapter 33 Encryption Overview 645

Chapter 34 Encryption Interfaces Configuration Guidelines 647

Configuring an Encryption Interface	647
Specifying the Security Association Name	648
Configuring the MTU for an Encryption Interface	648
Example: Configuring an Encryption Interface	648
Configuring Traffic	649
Traffic Overview	649
Configuring the Security Association	650
Configuring an Outbound Traffic Filter	651
Example: Configuring an Outbound Traffic Filter	651

	Applying the Outbound Traffic Filter	652
	Example: Applying the Outbound Traffic Filter	652
	Configuring an Inbound Traffic Filter	652
	Example: Configuring an Inbound Traffic Filter	653
	Applying the Inbound Traffic Filter to the Encryption Interface	653
	Example: Applying the Inbound Traffic Filter to the Encryption Interface	653
	Configuring an ES Tunnel Interface for a Layer 3 VPN	654
	Configuring ES PIC Redundancy	654
	Example: Configuring ES PIC Redundancy	655
	Configuring IPSec Tunnel Redundancy	655
Chapter 35	Summary of Encryption Configuration Statements	657
	address	657
	backup-destination	658
	backup-interface	658
	destination	659
	es-options	659
	family	660
	filter	661
	interfaces	661
	ipsec-sa	662
	source	662
	tunnel	663
	unit	664
Part 5	Flow Monitoring and Discard Accounting Services	
Chapter 36	Flow Monitoring and Discard Accounting Overview	667
	Passive Flow Monitoring	667
	Active Flow Monitoring	668
Chapter 37	Flow Monitoring and Discard Accounting Configuration Guidelines	671
	Minimum Traffic Sampling or Forwarding Configuration	675
	Configuring Traffic Sampling	675
	Configuring Traffic Sampling Properties	676
	Disabling Traffic Sampling	677
	Configuring Traffic Sampling Output	677
	Traffic Sampling Output Files	679

Tracing Traffic Sampling Operations	679
Examples: Configuring Traffic Sampling	680
Sampling a Single SONET Interface	680
Sampling All Traffic from a Single IP Address	681
Sampling All FTP Traffic	682
Configuring Flow Monitoring	683
Configuring the Flow-Monitoring Interface	683
Configuring Flow-Monitoring Properties	684
Configuring the Flow-Monitoring Interface	684
Exporting Flows	685
Configuring Timers	685
Example: Configuring Flow Monitoring	686
Configuring Flow Aggregation	687
Configuring Version 5 or Version 8 cflowd	688
Configuring Version 9 Flow Templates	690
Configuring the Traffic to be Sampled	690
Configuring the Version 9 Template Properties	691
Restrictions	692
Fields Included in Each Template Type	692
MPLS Sampling Behavior	693
Verification	694
Examples: Configuring Version 9 Flow Templates	694
Replicating Flows to Multiple Flow Servers	695
Replicating Routing Engine–Based Sampling to Multiple Flow Servers	696
Replicating Version 9 Flow Aggregation to Multiple Flow Servers	696
Debugging cflowd Flow Aggregation	697
Configuring Port Mirroring	698
Configuring Tunnels	700
Filter-Based Forwarding with Multiple Monitoring Interfaces	701
Restrictions	701
Configuring Port Mirroring on Services Interfaces	702
Examples: Configuring Port Mirroring	703
Load Balancing Among Multiple Monitoring Interfaces	707
Configuring Discard Accounting	710
Enabling Passive Flow Monitoring	711
Passive Flow Monitoring for MPLS Encapsulated Packets	712
Removing MPLS Labels from Incoming Packets	713
Example: Enabling Passive Flow Monitoring	714
Flow-Monitoring Redundancy	716

Chapter 38

Summary of Flow-Monitoring Configuration Statements **719**

accounting	720
address	721
aggregate-export-interval	721
aggregation	722

autonomous-system-type	723
cflowd	724
cflowd (Discard Accounting and Sampling)	724
cflowd (Flow Monitoring)	725
core-dump	725
destination	726
disable	726
engine-id	727
engine-type	727
export-format	728
family	729
family (Interfaces)	729
family (Monitoring)	730
family (Port Mirroring)	731
family (Sampling)	731
file	732
file (Sampling)	732
file (Trace Options)	733
filename	733
files	734
filter	735
flow-active-timeout	736
flow-export-destination	737
flow-inactive-timeout	737
flow-monitoring	738
forwarding-options	738
input	739
input (Port Mirroring)	739
input (Sampling)	739
input-interface-index	740
interface	741
interface (Accounting or Sampling)	741
interface (Monitoring)	742
interface (Port Mirroring)	742
interfaces	743
ipv4-template	743
label-position	744
local-dump	744
max-packets-per-second	745
monitoring	746
mpls-ipv4-template	747
mpls-template	747
multiservice-options	748
next-hop	748
next-hop-group	749
no-core-dump	749
no-filter-check	750
no-local-dump	750
no-stamp	750
no-syslog	750
no-world-readable	750

option-refresh-rate	751
output	752
output (Accounting)	752
output (Monitoring)	753
output (Port Mirroring)	753
output (Sampling)	754
output-interface-index	755
passive-monitor-mode	755
pop-all-labels	756
port	757
port-mirroring	758
rate	759
receive-options-packets	759
receive-ttl-exceeded	760
required-depth	760
run-length	761
sampling	762
sampling (Forwarding Options)	763
sampling (Interfaces)	764
services	764
size	765
source-address	766
stamp	766
syslog	767
template	768
template (Forwarding Options)	768
template (Services)	769
template-refresh-rate	769
traceoptions	770
unit	771
version	772
version9	773
version9 (Forwarding Options)	773
version9 (Services)	774
world-readable	774

Chapter 39

Flow Collection Configuration Guidelines 775

Configuring Flow Collection Properties	777
Configuring Flow Collector Destinations	777
Configuring a Packet Analyzer	777
Configuring File Formats	778
Configuring Interface Mappings	778
Configuring Transfer Logs	779
Configuring Retry Attempts	779
Sending cflowd Records to the Flow Collector Interface	780
Enabling Flow Collection Mode and Interface	780
Example: Flow Collector Interface Configuration	780

Chapter 40	Summary of Flow Collection Configuration Statements	787
	analyzer-address	787
	analyzer-id	787
	archive-sites	788
	collector	788
	data-format	789
	destinations	789
	filename-prefix	790
	file-specification	791
	file-specification (File Format)	791
	file-specification (Interface Mapping)	791
	flow-collector	792
	ftp	794
	ftp (Flow Collector Files)	795
	ftp (Transfer Log Files)	796
	interface-map	796
	maximum-age	797
	name-format	798
	password	800
	password (Flow Collector File Servers)	800
	password (Transfer Log File Servers)	800
	retry	801
	retry-delay	801
	transfer	802
	transfer-log-archive	802
	username	803
	variant	803
Chapter 41	Dynamic Flow Capture Configuration Guidelines	805
	Dynamic Flow Capture Architecture	805
	Liberal Sequence Windowing	806
	Configuring Dynamic Flow Capture Properties	806
	Configuring the Capture Group	807
	Configuring the Content Destination	808
	Configuring the Control Source	809
	Configuring the DFC PIC Interface	810
	Configuring System Logging	811
	Configuring Thresholds	811
	Configuring Maximum Duplicates	812
	Example: Dynamic Flow Capture Configuration	813

Chapter 42	Flow-Tap Configuration Guidelines	817
	Flow-Tap Architecture	818
	Configuring Flow-Tap Properties	819
	Configuring the Flow-Tap Interface	819
	Configuring Security Properties	820
	Restrictions	820
	Example: Flow-Tap Configuration	821
 Chapter 43	 Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements	 823
	address	823
	allowed-destinations	824
	capture-group	825
	content-destination	826
	control-source	827
	duplicates-dropped-periodicity	827
	dynamic-flow-capture	828
	flow-tap	829
	g-duplicates-dropped-periodicity	829
	g-max-duplicates	830
	hard-limit	830
	hard-limit-target	831
	input-packet-rate-threshold	831
	interface	832
	interfaces	832
	max-duplicates	833
	minimum-priority	833
	no-syslog	834
	notification-targets	834
	pic-memory-threshold	835
	service-port	835
	services	836
	shared-key	836
	soft-limit	837
	soft-limit-clear	837
	source-addresses	838
	ttl	838

Part 6**Link and Multilink Services****Chapter 44****Link and Multilink Services Overview****841****Chapter 45****Link and Multilink Services Configuration Guidelines****845**

Configuring Multilink and Link Services Logical Interface Properties	846
Default Settings for Multilink and Link Services Logical Interfaces	847
Configuring a Link Services Point-to-Point DLCI	848
Configuring a Link Services Multicast-Capable DLCI	848
Configuring a Drop Timeout Period	849
Configuring Logical Interface Encapsulation	850
Configuring a Fragmentation Threshold	851
Configuring Link Services Delay-Sensitive Packet Interleaving	851
Configuring LFI with DLCI Scheduling	853
Configuring Minimum Links	854
Configuring the MRRU and MTU Values	855
Configuring the Sequence Format	856
Configuring Compressed RTP with MLPPP Encapsulation	856
Example: Configuring Compressed RTP with MLPPP Encapsulation	857
Configuring Compressed RTP with PPP Encapsulation	857
Example: Configuring Compressed RTP with PPP Encapsulation	858
Configuring Link Services Physical Interface Properties	859
Default Settings for Link Services Interfaces	859
Configuring the Link Services Physical Interface Encapsulation	860
Configuring Link Services Acknowledgment Timers	860
Configuring the Link Services Differential Delay	861
Configuring Link Services Keepalive Settings on Frame Relay LMI	862
Multilink and Link Services Interface Structure	863
Multilink Services and Link Services PIC Capacities	863
Link Services PIC Capabilities	864
Configuring Bundles	864
Configuring CoS Components on Link Services PICs	866
Link Services CoS on J-series Services Routers	866
Link Services CoS on M-series and T-series Platforms	866
Example: Configuring Link Services CoS Components	868
Examples: Configuring Multilink Interfaces	871
Examples: Configuring Link Services Interfaces	875

Chapter 46**Summary of Link Services Configuration Statements****885**

acknowledge-retries	885
acknowledge-timer	886
action-red-differential-delay	886
address	887
bundle	887
compression-device	888
destination	888

disable-mlppp-inner-ppp-pfc	889
dlci	889
drop-timeout	890
encapsulation	891
encapsulation (Logical Interface)	891
encapsulation (Physical Interface)	892
family	893
fragment-threshold	894
hello-timer	894
interfaces	895
interleave-fragments	895
lmi-type	896
minimum-links	896
mlfr-uni-nni-bundle-options	897
mrru	898
mtu	898
multicast-dlci	899
n391	899
n392	900
n393	900
red-differential-delay	901
short-sequence	901
t391	902
t392	902
unit	903
yellow-differential-delay	904

Part 7

Real-Time Performance Monitoring Services

Chapter 47

Real-Time Performance Monitoring Services Overview 907

Chapter 48

Real-Time Performance Monitoring Configuration Guidelines 909

Configuring BGP Neighbor Discovery Through RPM	910
Configuring Real-Time Performance Monitoring Properties	912
Configuring the Probe	912
Configuring the Server	917
Configuring the Maximum Number of Probes	917
Configuring RPM Timestamping	917
Timestamps on M-series and T-series Routing Platforms	918
Timestamps on J-series Services Routers	920
Configuring TWAMP	920
Configuring the TWAMP Interface	921
Configuring the TWAMP Server Properties	921
Examples: Configuring BGP Neighbor Discovery Through RPM	922
Examples: Configuring Real-Time Performance Monitoring	923

Chapter 49
Summary of Real-Time Performance Monitoring Configuration Statements
927

client-list	927
data-fill	928
data-size	928
destination-interface	929
destination-port	930
dscp-code-point	931
hardware-timestamp	932
history-size	932
inactivity-timeout	933
logical-system	933
maximum-connections	934
maximum-connections-per-client	934
maximum-sessions	935
maximum-sessions-per-connection	935
moving-average-size	936
one-way-hardware-timestamp	936
port	937
port (RPM)	937
port (TWAMP)	937
probe	938
probe-count	939
probe-interval	939
probe-limit	940
probe-server	940
probe-type	941
routing-instance	941
routing-instances	942
rpm	942
server	943
services	943
source-address	944
target	944
tcp	945
test	946
test-interval	947
thresholds	948
traps	949
twamp	950
twamp-server	950
udp	951

Part 8**Tunnel Services**

Chapter 50	Tunnel Services Overview	955
Chapter 51	Tunnel Interfaces Configuration Guidelines	959
	Configuring a Unicast Tunnel	959
	Configuring a Key Number on GRE Tunnels	961
	Enabling Fragmentation on GRE Tunnels	962
	Specifying an MTU Setting for the Tunnel	962
	Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header	963
	Configuring Packet Reassembly	963
	Restricting a Tunnel to Multicast Traffic	964
	Configuring a Logical Tunnel Interface	964
	Connecting Logical Systems	965
	Configuring Logical Tunnels on J-series Platforms	966
	Configuring a Tunnel Interface for Routing Table Lookup	966
	Configuring a Virtual Loopback Tunnel for VRF Table Lookup	967
	Configuring PIM Tunnels	968
	Configuring an IPv6-over-IPv4 Tunnel	969
	Configuring a Dynamic Tunnel	969
	Configuring Tunnel Interfaces on MX-series Routers	970
	Example: Configuring Unicast Tunnels	970
	Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup	971
	Example: Configuring an IPv6-over-IPv4 Tunnel	972
	Example: Configuring a Logical Tunnel	972
Chapter 52	Summary of Tunnel Services Configuration Statements	975
	allow-fragmentation	975
	backup-destination	976
	copy-tos-to-outer-ip-header	976
	destination	977
	destination (Address)	977
	destination (Routing Instance)	977
	destination-networks	978
	do-not-fragment	978
	dynamic-tunnels	979
	interfaces	979
	key	980
	multicast-only	980
	peer-unit	981
	reassemble-packets	981
	routing-instance	982
	routing-instances	982
	routing-options	983
	source	983
	source-address	984
	ttl	984

tunnel985

tunnel-type986

unit987

Part 9

Index

Index991

Index of Statements and Commands1009

List of Figures

Figure 1: Packet Flow Through the Adaptive Services or MultiServices PIC	35
Figure 2: Configuring NAT for Multicast Traffic	139
Figure 3: IPSec Dynamic Endpoint Tunneling Topology	240
Figure 4: DLSw Ethernet Redundancy Topology	620
Figure 5: Example: IPSec Tunnel Connecting Security Gateways	649
Figure 6: IPSec Tunnel Redundancy	656
Figure 7: Passive Monitoring Application Topology	668
Figure 8: Active Monitoring Configuration Topology	670
Figure 9: Configure Sampling Rate	677
Figure 10: Flow Collector Interface Topology Diagram	781
Figure 11: Dynamic Flow Capture Topology	806
Figure 12: Flow-Tap Topology	819
Figure 13: Multilink Interface Configuration	865

List of Tables

Table 1: Notice Icons	xlvi
Table 2: Text and Syntax Conventions	xlvi
Table 3: Technical Documentation for Supported Routing Platforms	xlvi
Table 4: JUNOS Software Network Operations Guides	lii
Table 5: JUNOS Software with Enhanced Services Documentation	liii
Table 6: Additional Books Available Through http://www.juniper.net/books	liv
Table 7: AS and MultiServices PIC Services by Service Package, PIC, and Platform	30
Table 8: Statement Equivalents for ES and AS Interfaces	41
Table 9: Application Protocols Supported by Services Interfaces	55
Table 10: Network Protocols Supported by Services Interfaces	56
Table 11: ICMP Codes and Types Supported by Services Interfaces	58
Table 12: Port Names Supported by Services Interfaces	59
Table 13: Supported RPC Services	68
Table 14: IP Option Values	106
Table 15: Default IKE and IPSec Proposals for Dynamic Negotiations	235
Table 16: System Log Message Severity Levels	287
Table 17: System Log Message Severity Levels	441
Table 18: Adaptive Services Tracing Flags	449
Table 19: System Log Message Severity Levels	471
Table 20: Multilink and Link Services Logical Interface Statements	847
Table 21: Link Services Physical Interface Statements for MLFR FRF.16	859
Table 22: Multilink Services PIC Capacities	864
Table 23: Link Services CoS Queues	867
Table 24: Link Services Bundle	875
Table 25: Tunnel Interface Types	955
Table 26: Methods for Configuring Egress Filtering	967

About This Guide

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

- Objectives on page xliii
- Audience on page xliii
- Supported Routing Platforms on page xliv
- Using the Indexes on page xliv
- Using the Examples in This Manual on page xliv
- Documentation Conventions on page xlvi
- List of Technical Publications on page xlviii
- Documentation Feedback on page lv
- Requesting Technical Support on page lv

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: This guide documents Release 9.3 of the JUNOS software. For additional information about the 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 routing platform.

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 Routing Platforms

For the features described in this manual, the JUNOS software currently supports the following routing platforms:

- J-series
- M-series
- MX-series
- T-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.xml; }
```

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

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

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

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

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

Documentation Conventions

Table 1 on page xlvii 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 xlvii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>JUNOS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric <i>metric</i> >;
(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 (<i>string1</i> <i>string2</i> <i>string3</i>)
# (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 [<i>community-ids</i>]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; 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.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.

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

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

List of Technical Publications

Table 3 on page xlviii lists the software and hardware guides and release notes for Juniper Networks J-series, M-series, MX-series, and T-series routing platforms and describes the contents of each document. Table 4 on page lii lists the books included in the *Network Operations Guide* series. Table 5 on page liii lists the manuals and release notes supporting JUNOS software with enhanced services. All documents are available at <http://www.juniper.net/techpubs/>.

Table 6 on page liv lists additional books on Juniper Networks solutions that you can order through your bookstore. A complete list of such books is available at <http://www.juniper.net/books>.

Table 3: Technical Documentation for Supported Routing Platforms

Book	Description
JUNOS Software for Supported Routing Platforms	
<i>Access Privilege</i>	Explains how to configure access privileges in user classes by using permission flags and regular expressions. Lists the permission flags along with their associated command-line interface (CLI) operational mode commands and configuration statements.
<i>Class of Service</i>	Provides an overview of the class-of-service (CoS) functions of the JUNOS software and describes how to configure CoS features, including configuring multiple forwarding classes for transmitting packets, defining which packets are placed into each output queue, scheduling the transmission service level for each queue, and managing congestion through the random early detection (RED) algorithm.
<i>CLI User Guide</i>	Describes how to use the JUNOS command-line interface (CLI) to configure, monitor, and manage Juniper Networks routing platforms. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Feature Guide</i>	Provides a detailed explanation and configuration examples for several of the most complex features in the JUNOS software.
<i>High Availability</i>	Provides an overview of hardware and software resources that ensure a high level of continuous routing platform operation and describes how to configure high availability (HA) features such as nonstop active routing (NSR) and graceful Routing Engine switchover (GRES).
<i>MPLS Applications</i>	Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols.

Table 3: Technical Documentation for Supported Routing Platforms (continued)

Book	Description
<i>Multicast Protocols</i>	Provides an overview of multicast concepts and describes how to configure multicast routing protocols.
<i>Multiplay Solutions</i>	Describes how you can deploy IPTV and voice over IP (VoIP) services in your network.
<i>MX-series Layer 2 Configuration Guide</i>	Provides an overview of the Layer 2 functions of the MX-series routers, including configuring bridging domains, MAC address and VLAN learning and forwarding, and spanning-tree protocols. It also details the routing instance types used by Layer 2 applications. All of this material was formerly covered in the <i>JUNOS Routing Protocols Configuration Guide</i> .
<i>MX-series Solutions Guide</i>	Describes common configuration scenarios for the features supported on the MX-series routers, including basic bridged VLANs with normalized VLAN tags, aggregated Ethernet links, bridge domains, Multiple Spanning Tree Protocol (MSTP), and integrated routing and bridging (IRB).
<i>Network Interfaces</i>	Provides an overview of the network interface functions of the JUNOS software and describes how to configure the network interfaces on the routing platform.
<i>Network Management</i>	Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP and accounting options.
<i>Policy Framework</i>	Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, and forwarding options.
<i>Protected System Domain</i>	Provides an overview of the JCS 1200 platform and the concept of Protected System Domains (PSDs). The JCS 1200 platform, which contains up to 12 Routing Engines running JUNOS software, can be connected to up to three T-series routing platforms. To configure a PSD, you assign any number of Flexible PIC concentrators (FPCs) on a T-series routing platform to a pair of Routing Engines on the JCS 1200 platform. Each PSD has the same capabilities and functionality as a physical router, with its own control plane, forwarding plane, and administration.
<i>Routing Protocols</i>	Provides an overview of routing concepts and describes how to configure routing, routing instances, and unicast routing protocols.
<i>Secure Configuration Guide for Common Criteria and JUNOS-FIPS</i>	Provides an overview of secure Common Criteria and JUNOS-FIPS protocols for the JUNOS software and describes how to install and configure secure Common Criteria and JUNOS-FIPS on a routing platform.
<i>Services Interfaces</i>	Provides an overview of the services interfaces functions of the JUNOS software and describes how to configure the services interfaces on the router.
<i>Software Installation and Upgrade Guide</i>	Describes the JUNOS software components and packaging and explains how to initially configure, reinstall, and upgrade the JUNOS system software. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>Subscriber Access</i>	Provides an overview of the subscriber access features of the JUNOS software and describes how to configure subscriber access support on the router, including dynamic profiles, class of service, AAA, and access methods.
<i>System Basics</i>	Describes Juniper Networks routing platforms and explains how to configure basic system parameters, supported protocols and software processes, authentication, and a variety of utilities for managing your router on the network.
<i>VPNs</i>	Provides an overview and describes how to configure Layer 2 and Layer 3 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 2 circuits. Provides configuration examples.
JUNOS References	
<i>Hierarchy and RFC Reference</i>	Describes the JUNOS configuration mode commands. Provides a hierarchy reference that displays each level of a configuration hierarchy, and includes all possible configuration statements that can be used at that level. This material was formerly covered in the <i>JUNOS System Basics Configuration Guide</i> .
<i>Interfaces Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot interfaces.
<i>Routing Protocols and Policies Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot routing policies and protocols, including firewall filters.
<i>System Basics and Services Command Reference</i>	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot system basics, including commands for real-time monitoring and route (or path) tracing, system software management, and chassis management. Also describes commands for monitoring and troubleshooting services such as class of service (CoS), IP Security (IPsec), stateful firewalls, flow collection, and flow monitoring.
<i>System Log Messages Reference</i>	Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message.
J-Web User Guide	
<i>J-Web Interface User Guide</i>	Describes how to use the J-Web graphical user interface (GUI) to configure, monitor, and manage Juniper Networks routing platforms.
JUNOS API and Scripting Documentation	
<i>JUNOScript API Guide</i>	Describes how to use the JUNOScript application programming interface (API) to monitor and configure Juniper Networks routing platforms.
<i>JUNOS XML API Configuration Reference</i>	Provides reference pages for the configuration tag elements in the JUNOS XML API.

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>JUNOS XML API Operational Reference</i>	Provides reference pages for the operational tag elements in the JUNOS XML API.
<i>NETCONF API Guide</i>	Describes how to use the NETCONF API to monitor and configure Juniper Networks routing platforms.
<i>JUNOS Configuration and Diagnostic Automation Guide</i>	Describes how to use the commit script and self-diagnosis features of the JUNOS software. This guide explains how to enforce custom configuration rules defined in scripts, how to use commit script macros to provide simplified aliases for frequently used configuration statements, and how to configure diagnostic event policies.
Hardware Documentation	
<i>Hardware Guide</i>	Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide.
<i>PIC Guide</i>	Describes the routing platform's Physical Interface Cards (PICs). Each platform has its own PIC guide.
<i>DPC Guide</i>	Describes the Dense Port Concentrators (DPCs) for all MX-series routers.
JUNOScope Documentation	
<i>JUNOScope Software User Guide</i>	Describes the JUNOScope software graphical user interface (GUI), how to install and administer the software, and how to use the software to manage routing platform configuration files and monitor routing platform operations.
Advanced Insight Solutions (AIS) Documentation	
<i>Advanced Insight Solutions Guide</i>	Describes the Advanced Insight Manager (AIM) application, which provides a gateway between JUNOS devices and Juniper Support Systems (JSS) for case management and intelligence updates. Explains how to run AI-Scripts on Juniper Networks devices.
J-series Routing Platform Documentation	
<i>Getting Started Guide</i>	Provides an overview, basic instructions, and specifications for J-series routing platforms. The guide explains how to prepare your site for installation, unpack and install the router and its components, install licenses, and establish basic connectivity. Use the <i>Getting Started Guide</i> for your router model.
<i>Basic LAN and WAN Access Configuration Guide</i>	Explains how to configure the interfaces on J-series Services Routers for basic IP routing with standard routing protocols, ISDN backup, and digital subscriber line (DSL) connections.
<i>Advanced WAN Access Configuration Guide</i>	Explains how to configure J-series Services Routers in virtual private networks (VPNs) and multicast networks, configure data link switching (DLSw) services, and apply routing techniques such as policies, stateless and stateful firewall filters, IP Security (IPsec) tunnels, and class-of-service (CoS) classification for safer, more efficient routing.

Table 3: Technical Documentation for Supported Routing Platforms (*continued*)

Book	Description
<i>Administration Guide</i>	Shows how to manage users and operations, monitor network performance, upgrade software, and diagnose common problems on J-series Services Routers.
Release Notes	
<i>JUNOS Release Notes</i>	Summarize new features and known problems for a particular software release, provide corrections and updates to published JUNOS, JUNOScript, and NETCONF manuals, provide information that might have been omitted from the manuals, and describe upgrade and downgrade procedures.
<i>Hardware Release Notes</i>	Describe the available documentation for the routing platform and summarize known problems with the hardware and accompanying software. Each platform has its own release notes.
<i>JUNOScope Release Notes</i>	Contain corrections and updates to the published JUNOScope manual, provide information that might have been omitted from the manual, and describe upgrade and downgrade procedures.
<i>AIS Release Notes</i>	Summarize AIS new features and guidelines, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide initial setup, upgrade, and downgrade procedures.
<i>AIS AI-Scripts Release Notes</i>	Summarize AI-Scripts new features, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide instructions for automatic and manual installation, including deleting and rolling back.
<i>J-series Services Router Release Notes</i>	Briefly describe Services Router features, identify known hardware problems, and provide upgrade and downgrade instructions.

Table 4: JUNOS Software Network Operations Guides

Book	Description
<i>Baseline</i>	Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling JUNOS software, gathering basic system management information, verifying your network topology, and searching log messages.
<i>Interfaces</i>	Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms.
<i>MPLS</i>	Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routing platforms in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network.

Table 4: JUNOS Software Network Operations Guides (continued)

Book	Description
<i>MPLS Log Reference</i>	Describes MPLS status and error messages that appear in the output of the <code>show mpls lsp extensive</code> command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network.
<i>MPLS Fast Reroute</i>	Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing.
<i>Hardware</i>	Describes tasks for monitoring M-series and T-series routing platforms.

To configure and operate a J-series Services Router running JUNOS software with enhanced services, you must also use the configuration statements and operational mode commands documented in JUNOS configuration guides and command references. To configure and operate a WX Integrated Services Module, you must also use WX documentation.

Table 5: JUNOS Software with Enhanced Services Documentation

Book	Description
All Platforms	
<i>JUNOS Software Interfaces and Routing Configuration Guide</i>	Explains how to configure J-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification.
<i>JUNOS Software Security Configuration Guide</i>	Explains how to configure and manage security services such as stateful firewall policies, IP Security (IPsec) virtual private networks (VPNs), firewall screens, Network Address Translation (NAT), Public Key Cryptography, and Application Layer Gateways (ALGs).
<i>JUNOS Software Administration Guide</i>	Shows how to monitor J-series devices and routing operations, firewall and security services, system alarms and events, and network performance. This guide also shows how to administer user authentication and access, upgrade software, and diagnose common problems.
<i>JUNOS Software CLI Reference</i>	Provides the complete JUNOS software with enhanced services configuration hierarchy and describes the configuration statements and operational mode commands not documented in the standard JUNOS manuals.
J-series Only	
<i>JUNOS Software with Enhanced Services Design and Implementation Guide</i>	Provides guidelines and examples for designing and implementing IPsec VPNs, firewalls, and routing on J-series Services Routers running JUNOS software with enhanced services.

Table 5: JUNOS Software with Enhanced Services Documentation (continued)

Book	Description
<i>JUNOS Software with Enhanced Services Quick Start</i>	Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity.
<i>JUNOS Software with Enhanced Services J-series Services Router Hardware Guide</i>	Provides an overview, basic instructions, and specifications for J-series Services Routers. This guide explains how to prepare a site, unpack and install the router, replace router hardware, and establish basic router connectivity. This guide contains hardware descriptions and specifications.
<i>JUNOS Software with Enhanced Services Migration Guide</i>	Provides instructions for migrating an SSG device running ScreenOS software or a J-series Services Router running the JUNOS software to JUNOS software with enhanced services.
<i>WXC Integrated Services Module Installation and Configuration Guide</i>	Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration.
<i>JUNOS Software with Enhanced Services Release Notes</i>	Summarizes new features and known problems for a particular release of JUNOS software with enhanced services on J-series Services Routers, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for JUNOS software with enhanced services.

Table 6: Additional Books Available Through <http://www.juniper.net/books>

Book	Description
<i>Interdomain Multicast Routing</i>	Provides background and in-depth analysis of multicast routing using Protocol Independent Multicast sparse mode (PIM SM) and Multicast Source Discovery Protocol (MSDP); details any-source and source-specific multicast delivery models; explores multiprotocol BGP (MBGP) and multicast IS-IS; explains Internet Gateway Management Protocol (IGMP) versions 1, 2, and 3; lists packet formats for IGMP, PIM, and MSDP; and provides a complete glossary of multicast terms.
<i>JUNOS Cookbook</i>	Provides detailed examples of common JUNOS software configuration tasks, such as basic router configuration and file management, security and access control, logging, routing policy, firewalls, routing protocols, MPLS, and VPNs.
<i>MPLS-Enabled Applications</i>	Provides an overview of Multiprotocol Label Switching (MPLS) applications (such as Layer 3 virtual private networks [VPNs], Layer 2 VPNs, virtual private LAN service [VPLS], and pseudowires), explains how to apply MPLS, examines the scaling requirements of equipment at different points in the network, and covers the following topics: point-to-multipoint label switched paths (LSPs), DiffServ-aware traffic engineering, class of service, interdomain traffic engineering, path computation, route target filtering, multicast support for Layer 3 VPNs, and management and troubleshooting of MPLS networks.
<i>OSPF and IS-IS: Choosing an IGP for Large-Scale Networks</i>	Explores the full range of characteristics and capabilities for the two major link-state routing protocols: Open Shortest Path First (OSPF) and IS-IS. Explains architecture, packet types, and addressing; demonstrates how to improve scalability; shows how to design large-scale networks for maximum security and reliability; details protocol extensions for MPLS-based traffic engineering, IPv6, and multipoint topology routing; and covers troubleshooting for OSPF and IS-IS networks.

Table 6: Additional Books Available Through <http://www.juniper.net/books> (continued)

Book	Description
<i>Routing Policy and Protocols for Multivendor IP Networks</i>	Provides a brief history of the Internet, explains IP addressing and routing (Routing Information Protocol [RIP], OSPF, IS-IS, and Border Gateway Protocol [BGP]), explores ISP peering and routing policies, and displays configurations for both Juniper Networks and other vendors' routers.
<i>The Complete IS-IS Protocol</i>	Provides the insight and practical solutions necessary to understand the IS-IS protocol and how it works by using a multivendor, real-world approach.

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 name
- Document part number
- Page number
- Software release version (not required for *Network Operations Guides [NOGs]*)

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 located at <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 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 Internet Protocol 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 27.
- 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 647.
- 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 671.

- 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 845.
- 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 Multiprotocol Label Switching (MPLS). For more information about tunnel interfaces, see “Tunnel Interfaces Configuration Guidelines” on page 959.

Supported Platforms

For information about which platforms support Adaptive Services and MultiServices PICs and their features, see Table 7 on page 30.

For information about PIC support on a specific M-series or T-series platform, see the appropriate *PIC Guide* for the platform.

For information about services supported on J-series Services Routers, see the J-series documentation, in particular the *J-series Services Router Advanced WAN Access 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 protocols] Hierarchy Level on page 11
- [edit services] Hierarchy Level on page 12

[edit applications] Hierarchy Level

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

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

```

}
application-set application-set-name {
  [ application application-names ];
}

```

[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 [ group-names ] {
    interface interface-name {
        next-hop [ addresses ];
    }
}
port-mirroring {
    input {
        family (inet | inet6) {
            rate rate;
            run-length number;
        }
    }
    output {
        interface interface-name {
            next-hop address;
        }
        no-filter-check;
    }
    traceoptions {
        file filename {
            files number;
            size bytes;
            (world-readable | no-world-readable);
        }
    }
}
sampling {
    disable;
    input {
        family (inet | mpls) {
            max-packets-per-second number;
            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.

```

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 {
            maximum-contexts number <force>;
            f-max-period number;
            queues [ queue-numbers ];
            port {
                minimum port-number;
                maximum port-number;
            }
        }
    }
}
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);
}
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
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-names <service-filter filter-name> ];
            post-service-filter filter-name;
        }
        output {
            [ service-set service-set-names <service-filter filter-name> ];
        }
    }
}
fragment-threshold bytes;
interleave-fragments {
    queues [ queue-numbers ];
}
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;
        }
    }
}

```



```

reachability-cache-timeout seconds;
receive-initial-pacing number;
remote-peer peer-address {
    circuit-weight weight;
    cost cost;
}
traceoptions {
    file name <replace> <size size> <files number> <no-stamp> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}

```

[edit services] Hierarchy Level

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

```

adaptive-services-pics {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
        flag flag;
    }
}
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 {
            applications [ application-names ];
            application-sets [ set-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-names ];
}
}
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 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;
            mpls-template {
                label-position [ positions ];
            }
            mpls-ipv4-template {
                label-position [ positions ];
            }
            option-refresh-rate packets packets seconds seconds;
            template-refresh-rate packets packets seconds seconds;
        }
    }
}
flow-tap {
    interface interface-name;
}
ids {
    rule rule-name {
        match-direction (input | output | input-output);
        term term-name {
            from {
                applications [ application-names ];
                application-sets [ set-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-names ];
}
}
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 [ values ];
        ipv6_addr [ values ];
        key_id [ values ];
    }
    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 {
            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 ];
}
traceoptions {
    file {
        files number;
        size bytes;
    }
    flag flag;
}
}
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);
  pool nat-pool-name {
    address (address | address-range low value high value | prefix);
    pgcp {
      hint hint-string;
      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 {
      applications [ application-names ];
      application-sets [ set-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-names ];
}
}

```

```

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;
    maximum-synchronization-time seconds;
  }
  data-inactivity-detection {
    inactivity-delay;
    latch-deadlock-delay seconds;
    send-notification-on-delay;
    inactivity-duration seconds;
    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 {
    normal-mg-execution-time default milliseconds;
    normal-mgc-execution-time default milliseconds;
    mg-provisional-response-timer-value default milliseconds;
    mgc-provisional-response-timer-value default milliseconds;
    mg-originated-pending-limit default number-of-pendings-received;
    mgc-originated-pending-limit default number-of-pendings-received;
  }
  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;
  mgc-segmentation-timer default milliseconds;
  mg-maximum-pdu-size default bytes;
  mgc-maximum-pdu-size default bytes;
}
traffic-management {
  peak-data-rate {
    default bytes-per-second;
    rtcp {
      fixed-value bytes-per-second;
      percentage percentage;
    }
  }
  sustained-data-rate {
    default bytes-per-second;
    rtcp {
      fixed-value bytes-per-second;
      percentage percentage;
    }
  }
  max-burst-size {
    default bytes;
    rtcp {
      fixed-value bytes;
      percentage percentage;
    }
  }
}
inactivity-timer {
  inactivity-timeout {
    detect;
    maximum-inactivity-time {
      default 10-millisecond-units;
    }
  }
}
}
h248-options {
  audit-observed-events-returns-history;
  encoding {
    no-dscp-bit-mirroring;
  }
  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);
        failure (forced-904 | forced-906 | none);
        graceful (graceful-905 | none);
        link-loss (forced-906 | none);
    }
    virtual-interface-up {
        cancel-graceful (none | restart-918);
        warm (none | restart-900);
    }
}
context-indications {
    state-loss (forced-910 | forced-915 | none);
}
}
wildcard-response-service-change;
}
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 <files number> <match regex> <size size> <world-readable |
            no-world-readable>;
        flag flag;
    }
    virtual-interface interface-number {
        media-service media-service-name;
        interface interface-identifier;
        routing-instance instance-name {
            service-interface name.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 {
        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 {
    ([ 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 {
        ike-access-profile profile-name;
        local-gateway address;
        trusted-ca [ ca-profile-name ];
    }
    max-flows number;
    next-hop-service {
        inside-service-interface name.number;
        outside-service-interface name.number;
        service-interface-pool name;
    }
}

```

```

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 {
        applications [ application-names ];
        application-sets [ set-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-option [ values ];
        syslog;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
}

```

Part 2

Adaptive Services

- Adaptive Services Overview on page 27
- Applications Configuration Guidelines on page 53
- Summary of Applications Configuration Statements on page 91
- Stateful Firewall Services Configuration Guidelines on page 101
- Summary of Stateful Firewall Configuration Statements on page 111
- Network Address Translation Services Configuration Guidelines on page 123
- Summary of Network Address Translation Configuration Statements on page 145
- Intrusion Detection Service Configuration Guidelines on page 169
- Summary of Intrusion Detection Service Configuration Statements on page 181
- IPSec Services Configuration Guidelines on page 203
- Summary of IPSec Services Configuration Statements on page 245
- Layer 2 Tunneling Protocol Services Configuration Guidelines on page 279
- Summary of Layer 2 Tunneling Protocol Configuration Statements on page 297
- Link Services IQ Interfaces Configuration Guidelines on page 311
- Summary of Link Services IQ Configuration Statements on page 371
- Voice Services Configuration Guidelines on page 383
- Summary of Voice Services Configuration Statements on page 397
- Class-of-Service Configuration Guidelines on page 409
- Summary of Class-of-Service Configuration Statements on page 419
- Service Set Configuration Guidelines on page 433
- Summary of Service Set Configuration Statements on page 451
- Interface Configuration Guidelines on page 467
- Summary of Interface Configuration Statements on page 481
- Packet Gateway Configuration Guidelines on page 499
- Summary of Packet Gateway Configuration Statements on page 521
- Service Interface Pools Configuration Guidelines on page 603
- Summary of Service Interface Pools Statements on page 605

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 PIC with 256 megabytes (MB) of memory is supported on all M-series routing platforms except the M320 router.
- The Adaptive Services II PIC with 512 MB of memory is supported on all M-series and T-series routing platforms, including 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 routing platform. 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 routing platforms except M20 routers.

The MultiServices DPC is available for MX-series routers; it includes a subset of the functionality supported on the MultiServices PIC. For JUNOS Release 9.3, 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). For more information about supported packages, see “Enabling Service Packages” on page 29.



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.
- Internet Protocol 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 29
- Services Configuration Procedure on page 33
- Packet Flow Through the Adaptive Services or MultiServices PIC on page 34
- Stateful Firewall Overview on page 35
- Network Address Translation Overview on page 38
- IPSec Overview on page 39
- Layer 2 Tunneling Protocol Overview on page 42
- Voice Services Overview on page 42
- Class of Service Overview on page 43
- Packet Overview on page 43
- Examples: Services Interfaces Configuration on page 43

Enabling Service Packages

For AS PICs, MultiServices PICs, and the internal Adaptive Services Module (ASM) in the M7i platform, 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 routing platform, 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 except the ES PIC. It is supported on all M-series and T-series platforms except for TX Matrix installations. 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 7 on page 30 lists the services supported within each service package for each PIC and platform. Some of these services are also available on J-series Services 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 *J-series Services Router Advanced WAN Access 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 32 and “Link Services IQ Interfaces Configuration Guidelines” on page 311.



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 7: 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 7: 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	No
LNS Services:					
■ L2TP LNS	Yes	Yes (M7i and M10i only)	Yes (M120 only)	No	No
Voice Services:					

Table 7: 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
■ Packet gateway controller function	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 (<i>gr-fpc/pic/port</i>)	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 (<i>ip-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
■ Logical tunnels (<i>lt-fpc/pic/port</i>)	No	No	No	No	No
■ Multicast tunnels (<i>mt-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
■ PIM de-encapsulation (<i>pd-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
■ PIM encapsulation (<i>pe-fpc/pic/port</i>)	Yes	Yes	Yes	Yes	Yes
■ Virtual tunnels (<i>vt-fpc/pic/port</i>)	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 311 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 7 on page 30.

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 311 and “Link and Multilink Services Configuration Guidelines” on page 845.



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 35. (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.



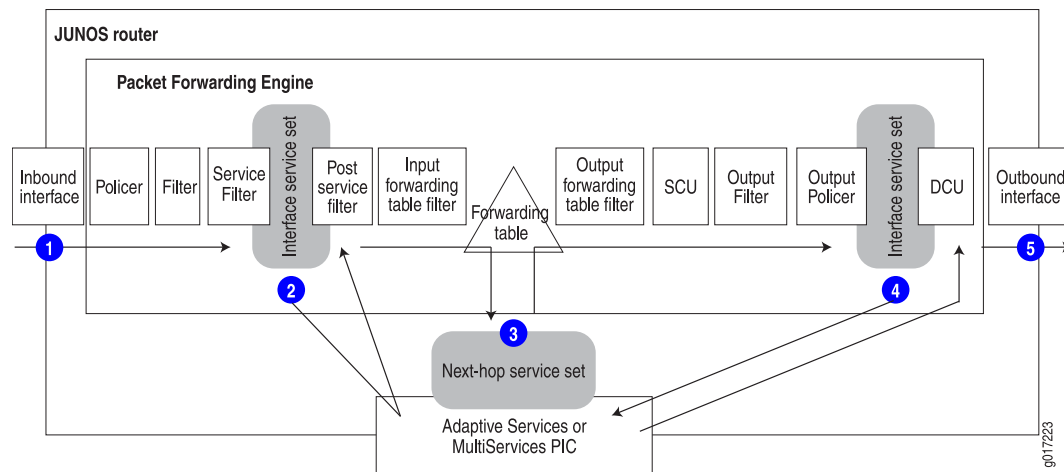
NOTE: For NAT, the next-hop service set can only be applied to the VRF table. For all other services, the next-hop service set can be applied to either the VRF table or to **inet.0**.

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 101.

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.

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

- Traditional NAT on page 38
- Twice NAT on page 39

Traditional NAT

Traditional NAT, specified in RFC 3022, *Traditional IP Network Address Translator*, is fully supported by 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 123.

Twice NAT

Twice NAT, specified in RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, is fully supported by 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 123.

IPSec Overview

The JUNOS software supports IPSec. This section discusses the following topics, which provide background information about configuring IPSec:

- IPSec on page 40
- Security Associations on page 40
- IKE on page 40
- Comparison of IPSec Services and ES Interface Configuration on page 41

For a list of IPSec- and IKE-supported standards, 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 8 on page 41 compares the top-level configuration of IPSec features on the ES PIC interfaces and on the AS or MultiServices PIC interfaces.

Table 8: 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 <i>rule-name</i>] term <i>term-name</i> match-conditions {...} then dynamic {...}]
[edit security ipsec] security-association sa-manual {...}	[edit services ipsec-vpn rule <i>rule-name</i>] term <i>term-name</i> 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 <i>es-fpc/pic/port</i>] tunnel source <i>address</i>	[edit services ipsec-vpn service-set <i>set-name</i> ipsec-vpn local-gateway <i>address</i>]
[edit interfaces <i>es-fpc/pic/port</i>] tunnel destination <i>address</i>	[edit services ipsec-vpn rule <i>rule-name</i>] remote-gateway <i>address</i>

For more information about configuring IPSec services on an AS or MultiServices PIC, see “IPSec Services Configuration Guidelines” on page 203. For more information about configuring encryption services on an ES PIC, see “Encryption Interfaces Configuration Guidelines” on page 647.



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 routing platforms only:

- M7i (AS PICs)
- M10i (AS and MultiServices 100 PICs)
- M120 (AS, MultiServices 100, and MultiServices 400 PICs)

For more information about configuring L2TP services, see “Layer 2 Tunneling Protocol Services Configuration Guidelines” on page 279.

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.



NOTE: On J-series Services Routers, you can configure CRTP with MLPPP or PPP logical interface encapsulation on link services (ls-) interfaces. For more information, see “Configuring Compressed RTP with MLPPP Encapsulation” on page 856 and “Configuring Compressed RTP with PPP Encapsulation” on page 857.

Voice services do not require a separate service rules configuration.

Voice services also support LFI on M-series routers, except the M320 router. For more information about configuring voice services, see “Voice Services Configuration Guidelines” on page 383.

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 page 338.

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 409.

Packet Overview

For overview information about the packet gateway feature, see the *JUNOS Multiplay Solutions Guide*.

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:

```
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-names ];  
}
```

This chapter includes the following sections:

- Configuring Application Protocol Properties on page 54
- Configuring Application Sets on page 64
- ALG Descriptions on page 64
- Verification on page 72
- JUNOS Default Groups on page 79
- Examples: Configuring Applications on page 87

Configuring Application Protocol Properties

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

```
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 64.

This section includes the following tasks for configuring applications:

- Configuring an Application Protocol on page 54
- Configuring the Network Protocol on page 56
- Configuring the ICMP Code and Type on page 57
- Configuring Source and Destination Ports on page 59
- Configuring the Inactivity Timeout Period on page 62
- Configuring SIP on page 62
- Configuring an SNMP Command on page 63
- Configuring an RPC Program Number on page 63
- Configuring the TTL Threshold on page 63
- Configuring a Universal Unique Identifier on page 64

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:

```
application-protocol protocol-name;
```

Table 9 on page 55 shows the list of supported protocols. For more information about specific protocols, see “ALG Descriptions” on page 64.

Table 9: 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.
Internet Protocol (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.
Remote Procedure Call (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 62.
Shell	shell	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.

Table 9: Application Protocols Supported by Services Interfaces (*continued*)

Protocol Name	CLI Value	Comments
Simple Network Management Protocol (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 123.

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:

```
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 10 on page 56 shows the list of the supported protocols.

Table 10: Network Protocols Supported by Services Interfaces

Network Protocol Type	CLI Value	Comments
IP Security (IPSec) authentication header (AH)	ah	–

Table 10: Network Protocols Supported by Services Interfaces *(continued)*

Network Protocol Type	CLI Value	Comments
External Gateway Protocol (EGP)	egp	–
IPSec Encapsulating Security Payload (ESP)	esp	–
Generic routing encapsulation (GR)	gre	–
ICMP	icmp	Requires an <code>application-protocol</code> value of <code>icmp</code> .
Internet Group Management Protocol (IGMP)	igmp	–
IP in IP	ipip	–
Open Shortest Path First (OSPF)	ospf	–
Protocol Independent Multicast (PIM)	pim	–
Resource Reservation Protocol (RSVP)	rsvp	–
TCP	tcp	Requires a <code>destination-port</code> or <code>source-port</code> value unless you specify <code>application-protocol rcp</code> or <code>dce-rcp</code> .
UDP	udp	Requires a <code>destination-port</code> or <code>source-port</code> value unless you specify <code>application-protocol rcp</code> or <code>dce-rcp</code> .
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: Internet Protocol 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 123.

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:

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 11 on page 58 shows the list of supported ICMP values.

Table 11: ICMP Codes and Types Supported by Services Interfaces

CLI Statement	Description
icmp-code	<p>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 <i>JUNOS Policy Framework Configuration Guide</i>.</p> <p>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:</p> <p>parameter-problem: ip-header-bad (0), required-option-missing (1)</p> <p>redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2)</p> <p>time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)</p> <p>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)</p>
icmp-type	<p>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 <i>JUNOS Policy Framework Configuration Guide</i>.</p> <p>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).</p>



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:

```
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 9 on page 55.

You can specify either a numeric value or one of the text synonyms listed in Table 12 on page 59.

Table 12: 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 12: 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
mobileip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518

Table 12: 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
xmcp	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:

```
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” on page 470.

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 54. 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:

```
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:

```
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

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:

```
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 54.

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:

```
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 54.

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:

```
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 54.

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:

```
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 54. For more information on UUID numbers, see <http://www.opengroup.org/onlinepubs/9629399/apdx.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:

```
application-set application-set-name {
  application application1;
  application application2;
}
```

For an example of a typical application set, see “Examples: Configuring Applications” on page 87.

ALG Descriptions

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

- Basic TCP ALG on page 65
- Basic UDP ALG on page 65
- BOOTP on page 66
- DCE RPC Services on page 66
- FTP on page 66
- H323 on page 67
- ICMP on page 67
- IIOP on page 67
- NetShow on page 68
- RealAudio on page 68
- RPC and RPC Portmap Services on page 68
- RTSP on page 69
- SMB on page 70
- SNMP on page 70
- SQLNet on page 70

- TFTP on page 70
- Traceroute on page 71
- UNIX Remote-Shell Services on page 71
- WinFrame on page 71

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 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 13 on page 68:

Table 13: Supported RPC Services

Name	Description	Comments
rpc.mountd	Network File Server (NFS) mount daemon for details, see the UNIX man page for <code>rpc.mountd(8)</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050).
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 RFC1050).
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 RFC1050).
rpc.nlockmgr	Network lock manager.	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc.nlockmgr</code> service can be allowed or blocked based on RPC program 100021.
rpc.pcnfsd	Kernel statistics server. For details, see the UNIX man pages for <code>rstatd</code> and <code>rpc.rstatd</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc.rstat</code> service can be allowed or blocked based on RPC program 150001.

Table 13: Supported RPC Services (continued)

Name	Description	Comments
rpc-rwall	Used to write a message to users; for details, see the UNIX man page for <code>rpc.rwalld</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc-rwall</code> service can be allowed or blocked based on RPC program 150008.
rpc-ybind	NIS binding process. For details, see the UNIX man page for <code>ybind</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc-ybind</code> service can be allowed or blocked based on RPC program 100007.
rpc-yppasswd	NIS password server. For details, see the UNIX man page for <code>yppasswd</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc-yppasswd</code> service can be allowed or blocked based on RPC program 100009.
rpc-ypserv	NIS server. For details, see the UNIX man page for <code>ypserv</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc-ypserv</code> 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 RFC1050). Once the RPC program table is built, <code>rpc-ypupdated</code> 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 <code>rpc.ypxfrd</code> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC1050). Once the RPC program table is built, <code>rpc-ypxfrd</code> 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 (**smb**s) 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: $+ nhops - 1$. 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.

Verification

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 72
- RTSP ALG Example on page 75
- System Log Messages on page 77

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 72
- FTP System Log Messages on page 73
- Analysis on page 74
- Troubleshooting Questions on page 74

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 77.

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

- Rule Accept system log:


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


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


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

Analysis

Control Flows

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

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

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

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

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

Data Flows

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

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

- 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.
- 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.
- 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 75
- Analysis on page 75
- Troubleshooting Questions on page 76

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 78
- System Log Output on page 79

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 routing platform. 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#
groups {
  junos-defaults {
    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 54; for details about a specific protocol, see “ALG Descriptions” on page 64.

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:

```
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.

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

```
    }
}
```

Examples: Configuring Applications

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:

```
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 54.

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

application-protocol

Syntax	<code>application-protocol <i>protocol-name</i>;</code>
Hierarchy Level	<code>[edit applications application <i>application-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4. login and sip options added in JUNOS Release 7.4. ip option added in JUNOS Release 8.2.
Description	Identify the application protocol name. Application protocols are also called application layer gateways (ALGs).
Options	<p><i>protocol-name</i>—Name of the protocol. The following protocols are supported:</p> <ul style="list-style-type: none"> 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 54.

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-names] {
 [application application-names];
}`

Hierarchy Level [edit applications]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure which applications to include in an application set.

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

Usage Guidelines See “Configuring Application Sets” on page 64.

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 53.

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 <i>port-value</i> ;
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number.
Options	<i>port-value</i> —Identifier for the port. For a complete list, see “Configuring Source and Destination Ports” on page 59.
Usage Guidelines	See “Configuring Source and Destination Ports” on page 59.
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 <i>value</i> ;
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Internet Control Message Protocol (ICMP) code value.
Options	<i>value</i> —The ICMP code value. For a complete list, see “Configuring the ICMP Code and Type” on page 57.
Usage Guidelines	See “Configuring the ICMP Code and Type” on page 57.
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 <i>value</i> ;
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	ICMP packet type value.
Options	<i>value</i> —The ICMP type value, such as <code>echo</code> or <code>echo-reply</code> . For a complete list, see “Configuring the ICMP Code and Type” on page 57.
Usage Guidelines	See “Configuring the ICMP Code and Type” on page 57.
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 <i>seconds</i> ;
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Inactivity timeout period, in seconds.
Options	<i>seconds</i> —Length of time the application is inactive before it times out. Default: 60 seconds
Usage Guidelines	See “Configuring the Inactivity Timeout Period” on page 62.
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 <i>application-name</i>]
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 62.
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: Internet Protocol version 6 (IPv6) is not supported as a network protocol in application definitions.

Usage Guidelines See “Configuring the Network Protocol” on page 56.

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

rpc-program-number

Syntax	<code>rpc-program-number <i>number</i>;</code>
Hierarchy Level	<code>[edit applications application <i>application-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Remote procedure call (RPC) or Distributed Computing Environment (DCE) value.
Options	<i>number</i> —RPC or DCE program value. Range: 100,000 through 400,000
Usage Guidelines	See “Configuring an RPC Program Number” on page 63.
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	<code>sip-call-hold-timeout <i>seconds</i>;</code>
Hierarchy Level	<code>[edit applications application <i>application-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Timeout period for SIP calls placed on hold, in seconds.
Options	<i>seconds</i> —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 62.
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 <i>command</i> ;
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Simple Network Management Protocol (SNMP) command format.
Options	<i>command</i> —Supported commands are SNMP <i>get</i> , <i>get-next</i> , <i>set</i> , and <i>trap</i> .
Usage Guidelines	See “Configuring an SNMP Command” on page 63.
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 <i>port-number</i> ;
Hierarchy Level	[edit applications application <i>application-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Source port identifier.
Options	<i>port-value</i> —Identifier for the port. For a complete list, see “Configuring Source and Destination Ports” on page 59.
Usage Guidelines	See “Configuring Source and Destination Ports” on page 59.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ttl-threshold

Syntax	<code>ttl-threshold <i>number</i>;</code>
Hierarchy Level	<code>[edit applications application <i>application-name</i>]</code>
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	<i>number</i> —TTL threshold value.
Usage Guidelines	See “Configuring the TTL Threshold” on page 63.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

uuid

Syntax	<code>uuid <i>hex-value</i>;</code>
Hierarchy Level	<code>[edit applications application <i>application-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the Universal Unique Identifier (UUID) for DCE RPC objects.
Options	<i>hex-value</i> —Hexadecimal value.
Usage Guidelines	See “Configuring a Universal Unique Identifier” on page 64.
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` statements at the `[edit services]` hierarchy level:

```
stateful-firewall {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        applications [ application-names ];
        application-sets [ set-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-option [ values ];
        syslog;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
}
```

This chapter contains the following sections:

- Configuring Stateful Firewall Properties on page 102
- Examples: Configuring Stateful Firewall Properties on page 106

Configuring Stateful Firewall Properties

This section describes the following tasks for configuring stateful firewalls:

- Configuring the Stateful Firewall Rule Set on page 102
- Configuring Stateful Firewall Rule Content on page 102

Configuring the Stateful Firewall Rule Set

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:

```
rule-set rule-set-name {
    rule rule-name1;
    rule rule-name2;
    rule rule-name3;
    ...
}
```

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 Stateful Firewall Rule Content

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

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            applications [ application-names ];
            application-sets [ set-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-option [ 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 describe stateful firewall rule content in more detail:

- Configuring Match Direction on page 103
- Configuring Stateful Firewall Match Conditions on page 103
- Configuring Stateful Firewall Actions on page 104

Configuring Match Direction

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:

```
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 Services Interfaces” on page 434.

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 Stateful Firewall Match Conditions

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

```
from {
  applications [ application-names ];
```

```

application-sets [ set-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 Properties” on page 106.

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 53.

- 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 Stateful Firewall Actions

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

```

then {
    (accept | discard | reject);
    allow-ip-option [ 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-option** 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-option** statement. If you do not configure **allow-ip-option**, 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 14 on page 106 lists the possible **allow-ip-option** values. 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 14: 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	–

Examples: Configuring Stateful Firewall Properties

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 content at the `[edit policy-options]` hierarchy level:

```

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:

```

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:

```

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`.

```

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 43.

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-option

Syntax allow-ip-option [*values*];

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 Stateful Firewall Actions” on page 104.

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

application-sets

Syntax	<code>applications-sets [<i>set-name</i>];</code>
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more target application sets.
Options	<i>set-name</i> —Name of the target application set.
Usage Guidelines	See “Configuring Stateful Firewall Match Conditions” on page 103.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

applications

Syntax	<code>applications [<i>application-name</i>];</code>
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more applications to which the stateful firewall services apply.
Options	<i>application-name</i> —Name of the target application.
Usage Guidelines	See “Configuring Stateful Firewall Match Conditions” on page 103.
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 (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> 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 Stateful Firewall Match Conditions” on page 103.
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 <i>minimum-value</i> high <i>maximum-value</i> <except>;
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the destination address range for rule matching.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. except—(Optional) Exclude the specified address range from rule matching.
Usage Guidelines	See “Configuring Stateful Firewall Match Conditions” on page 103.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-prefix-list

Syntax	<code>destination-prefix-list <i>list-name</i> <except>;</code>
Hierarchy Level	<code>[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from]</code>
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the destination prefix list for rule matching. You configure the prefix list using the <code>prefix-list</code> statement at the <code>[edit policy-options]</code> hierarchy level.
Options	<p><i>list-name</i>—Destination prefix list.</p> <p><code>except</code>—(Optional) Exclude the specified prefix list from rule matching.</p>
Usage Guidelines	See “Configuring Stateful Firewall Match Conditions” on page 103.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

from

Syntax	<pre>from { applications [<i>application-names</i>]; application-sets [<i>set-names</i>]; destination-address (<i>address</i> any-unicast) <except>; destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; destination-prefix-list <i>list-name</i> <except>; source-address (<i>address</i> any-unicast) <except>; source-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>; source-prefix-list <i>list-name</i> <except>; }</pre>
Hierarchy Level	<code>[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify input conditions for a stateful firewall term.
Options	<p>For information on match conditions, see the description of firewall filter match conditions in the <i>JUNOS Policy Framework Configuration Guide</i>.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Stateful Firewall Rule Content” on page 102.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

match-direction

Syntax	match-direction (input output input-output);
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the direction in which the rule match is applied.
Options	<p>input—Apply the rule match on the input side of the interface.</p> <p>output—Apply the rule match on the output side of the interface.</p> <p>input-output—Apply the rule match bidirectionally.</p>
Usage Guidelines	See “Configuring Stateful Firewall Rule Content” on page 102.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

rule

Syntax

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            applications [ application-names ];
            application-sets [ set-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 Rule Content” on page 102.

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

rule-set

Syntax	<code>rule-set <i>rule-set-name</i> { [rule <i>rule-names</i>]; }</code>
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	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
Usage Guidelines	See “Configuring the Stateful Firewall Rule Set” on page 102.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

services

Syntax	<code>services stateful-firewall { ... }</code>
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the service rules to be applied to traffic.
Options	<i>stateful-firewall</i> —Identifies the stateful firewall set of rules statements.
Usage Guidelines	See “Configuring Stateful Firewall Properties” on page 102.
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 (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> 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 Stateful Firewall Match Conditions” on page 103.
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 <i>minimum-value</i> high <i>maximum-value</i> <except>;
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Source address range for rule matching.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —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 Stateful Firewall Match Conditions” on page 103.
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 <i>list-name</i> <except>;
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the source prefix list for rule matching. You configure the prefix list using the <code>prefix-list</code> statement at the [edit <code>policy-options</code>] hierarchy level.
Options	<i>list-name</i> —Destination prefix list. except—(Optional) Exclude the specified prefix list from rule matching.
Usage Guidelines	See “Configuring Stateful Firewall Match Conditions” on page 103.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

syslog

Syntax	syslog;
Hierarchy Level	[edit services stateful-firewall rule <i>rule-name</i> term <i>term-name</i> 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 <code>/var/log</code> directory. This setting overrides any <code>syslog</code> statement setting included in the service set or interface default configuration.
Usage Guidelines	See “Configuring Stateful Firewall Actions” on page 104.
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 {`
 `applications [application-names];`
 `application-sets [set-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 Rule Content” on page 102.

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 Stateful Firewall Actions” on page 104.

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:

```
nat {
  ipv6-multicast-interfaces (all interface-name);
  pool nat-pool-name {
    address [ prefix ];
    address-range [ low minimum-value high maximum-value];
    pgcp {
      hint hint-string;
      ports-per-session ports;
      remotely-controlled;
      transport [ transport-protocol ];
    }
    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 {
      applications [ application-names ];
      application-sets [ set-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 Network Address Translation Properties on page 124
- Examples: Configuring Network Address Translation Properties on page 133

Configuring Network Address Translation Properties

This section describes the following tasks for configuring network address translation services:

- Configuring Address and Port Information on page 124
- Configuring the NAT Rule Set on page 127
- Configuring NAT Rule Content on page 128

Configuring Address and Port Information

To specify translated addresses, you can perform the tasks in either of the following sections:

- Configuring a Pool on page 124
- Specifying the Destination or Source Prefix on page 126
- Addressing Information on page 126
- Configuring IPv6 Multicast Filters on page 127

For additional addressing information, see “Addressing Information” on page 126.

Configuring a Pool

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 **pool** statement to specify that NAT pool is used exclusively by the packet gateway. To configure the information, include the **pool** statement at the [edit **services nat**] hierarchy level:


```

pool nat-pool-name {
  address [ prefix ];
  pgcp;
  address-range [ low minimum-value high maximum-value];
  port (automatic | range low minimum-value high maximum-value) {
    random-allocation;
  }
}

```

To configure pools for traditional NAT, you specify either a destination pool or a source pool. To configure pools for twice NAT, you 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 NAT Actions” on page 131.

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.

For more information about configuring a NAT pool for the packet gateway by including the `pgcp` statement, see “Configuring NAT Pools for the Packet Gateway” on page 507.

Specifying the Destination or Source Prefix

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:

```
rule rule-name {
  term term-name {
    then {
      translated {
        destination-prefix prefix;
        source-prefix prefix;
      }
    }
  }
}
```

Addressing Information

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 Network Address Translation Properties” on page 133.

- 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.

```
ipv6-multicast-interfaces (all interface-name);
```

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.

Configuring the NAT Rule Set

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:

```
rule-set rule-set-name {
  rule rule-name1;
  rule rule-name2;
  rule rule-name3;
  ...
}
```

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 NAT Rule Content

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

```
rule rule-name {
  match-direction (input | output);
  term term-name {
    nat-type (full-cone | symmetric);
    from {
      applications [ application-names ];
      application-sets [ set-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 describe NAT rule content in more detail:

- Configuring NAT Match Direction on page 129
- Configuring NAT Type on page 129
- Configuring NAT Match Conditions on page 130
- Configuring NAT Actions on page 131

Configuring NAT Match Direction

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:

```
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 Services Interfaces” on page 434.

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

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:

```
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 NAT Actions” on page 131.
- It supports IPv4 addresses on Juniper Networks J-series Services Routers only. It is not supported on M-series or T-series routing platforms.
- 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 “Configuring Full-Cone NAT” on page 143.

Configuring NAT Match Conditions

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

```
from {
  applications [ application-names ];
  application-sets [ set-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 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 NAT rule. For an example, see “Examples: Configuring Stateful Firewall Properties” on page 106.

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 53:

- 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 NAT Actions

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

```
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 Address and Port Information” on page 124.

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*.

Examples: Configuring Network Address Translation Properties

This section provides the following configuration examples:

- Dynamic Source Translation on page 134
- Static Source Translation on page 134
- Dynamic and Static Source Translation on page 134
- Oversubscribed Pool with No Fallback on page 135
- Oversubscribed Pool with a Fallback to NAPT on page 136
- Multiple Prefixes and Address Ranges in Static Source Translation on page 136
- Assigning Addresses from a Dynamic Pool for Static Use on page 137
- Configuring NAT Rules Without Defining a Pool on page 137
- Configuring Specific Addresses That Are Not to Be Translated on page 138
- Configuring Network Address Translation for Multicast Traffic on page 139
- Configuring Twice NAT on page 143
- Configuring Full-Cone NAT on page 143

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 43.

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;
      }
    }
  }
}
```

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;
      }
    }
  }
}
```

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
      }
    }
  }
}

```

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;
      }
    }
  }
}

```

Oversubscribed Pool with a Fallback to NAT

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 NATP 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;
      }
    }
  }
}
```

Multiple Prefixes and Address Ranges in Static Source Translation

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;
      }
    }
  }
}
```

```

    }
}

```

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;
        }
    }
}

```

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;
    }
  }
}
```

Configuring Specific Addresses That Are Not to Be Translated

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;
      }
    }
  }
}
```

```


    }
  }
}

```

Configuring Network Address Translation for Multicast Traffic

Figure 2 on page 139 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 Open Shortest Path First (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;
}
```

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;
        }
      }
    }
  }
}
```

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;  
        }  
      }  
    }  
  }  
}
```

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 [<i>prefix</i>];
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. <i>prefix</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the NAT pool prefix value.
Options	<i>prefix</i> —Specify an IPv4 or IPv6 prefix value.
Usage Guidelines	See “Configuring Address and Port Information” on page 124.
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 <i>minimum-value</i> high <i>maximum-value</i>];
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the NAT pool address range.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.
Usage Guidelines	See “Configuring Address and Port Information” on page 124.
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 [<i>set-names</i>];
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more target application sets.
Options	<i>set-name</i> —Name of the target application set.
Usage Guidelines	See “Configuring NAT Match Conditions” on page 130.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

applications

Syntax	<code>applications [<i>application-names</i>];</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more application protocols to which the NAT services apply.
Options	<i>application-name</i> —Name of the target application.
Usage Guidelines	See “Configuring NAT Match Conditions” on page 130.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-address

Syntax	<code>destination-address (<i>address</i> any-unicast) <except>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> 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	<i>address</i> —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 NAT Match Conditions” on page 130.
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 <i>minimum-value</i> high <i>maximum-value</i> <except>;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the destination address range for rule matching.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. except—(Optional) Prevent the specified address range from being translated.
Usage Guidelines	See “Configuring NAT Match Conditions” on page 130.
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 <i>nat-pool-name</i> ;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the destination address pool for translated traffic.
Options	<i>nat-pool-name</i> —Destination pool name.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-prefix

Syntax	<code>destination-prefix <i>destination-prefix</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in JUNOS Release 7.6. <i>destination-prefix</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the destination prefix for translated traffic.
Options	<i>destination-prefix</i> —IPv4 or IPv6 destination prefix value.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-prefix-list

Syntax	<code>destination-prefix-list <i>list-name</i> <except>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the destination prefix list for rule matching. You configure the prefix list using the <i>prefix-list</i> statement at the [edit policy-options] hierarchy level.
Options	<i>list-name</i> —Destination prefix list. except—(Optional) Exclude the specified prefix list from rule matching.
Usage Guidelines	See “Configuring NAT Match Conditions” on page 130.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

from

Syntax from {
 application-sets [*set-names*];
 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 Rule Content” on page 128.

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

hint

Syntax	<code>hint <i>hint-string</i> [hint <i>hint-string</i>];</code>
Hierarchy Level	<code>[edit services nat pool <i>nat-pool-name</i> pgcp]</code>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Configure a hint that enables the packet gateway to choose a NAT pool by direction rather than by virtual interface. The packet gateway 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 packet gateway can choose any NAT pool associated with the virtual interface.
Options	<i>hint-string</i> —Alphanumeric string of up to 3 characters that the packet gateway 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: <code>[hint xx hint yy]</code> .
Usage Guidelines	See “Configuring NAT Pools for the Packet Gateway” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ipv6-multicast-interfaces

Syntax	<code>ipv6-multicast-interfaces (all <i>interface-name</i>);</code>
Hierarchy Level	<code>[edit services nat]</code>
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	<i>all</i> —Enable filters on all interfaces. <i>interface-name</i> —Enable filters on a specific interface only.
Usage Guidelines	See “Configuring IPv6 Multicast Filters” on page 127.
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 <i>rule-name</i>]
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 Rule Content” on page 128.
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 <i>rule-name</i> term <i>term-name</i>]
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” on page 129
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 <i>rule-name</i> term <i>term-name</i> 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 NAT Actions” on page 131.
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 <i>overload-pool-name</i> ;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> 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	<i>overload-pool-name</i> —Name of the overload pool.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

overload-prefix

Syntax	<code>overload-prefix <i>overload-prefix</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> 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	<i>overload-prefix</i> —Prefix value.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pgcp

Syntax	<pre>pgcp { hint <i>hint-string</i>; ports-per-session <i>ports</i>; remotely-controlled; transport [<i>transport-protocol</i>]; }</pre>
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4. <code>remotely-controlled</code> and <code>ports-per-session</code> statements added in JUNOS Release 8.5. <code>hint</code> statement added in JUNOS Release 9.0.
Description	Specify that the NAT pool is used exclusively by the packet gateway.
Usage Guidelines	See “Configuring NAT Pools for the Packet Gateway” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pool

Syntax `pool nat-pool-name {
 address [prefix];
 address-range [low minimum-value high maximum-value];
 pgcp {
 hint hint-string;
 ports-per-session ports;
 remotely-controlled;
 }
 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 Address and Port Information” on page 124.

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 <i>minimum-value</i> high <i>maximum-value</i>) { random-allocation; }
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
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. <i>minimum-value</i> —Lower boundary for the port range. <i>maximum-value</i> —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 Address and Port Information” on page 124.
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 <i>ports</i> ;
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i> 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	<i>number-of-ports</i> —Number of ports to enable: 2 or 4 for combined voice and video services. Default: 2
Usage Guidelines	See “Configuring NAT Pools for the Packet Gateway” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

remotely-controlled

Syntax	remotely-controlled;
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i> 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 packet gateway controller.
Usage Guidelines	See “Configuring NAT Pools for the Packet Gateway” on page 507.
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);
  term term-name {
    nat-type (full-cone | symmetric);
    from {
      applications [ application-names ];
      application-sets [ set-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 Rule Content” on page 128.

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

rule-set

Syntax	<code>rule-set <i>rule-set-name</i> { [rule <i>rule-names</i>]; }</code>
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	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
Usage Guidelines	See “Configuring the NAT Rule Set” on page 127.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

services

Syntax	<code>services nat { ... }</code>
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the service rules to be applied to traffic.
Options	<i>nat</i> —Identifies the NAT set of rules statements.
Usage Guidelines	See “Configuring Network Address Translation Properties” on page 124.
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 (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> 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 NAT Match Conditions” on page 130.
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 <i>minimum-value</i> high <i>maximum-value</i> <except>;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the source address range for rule matching.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. except—(Optional) Prevent the specified address range from being translated.
Usage Guidelines	See “Configuring NAT Match Conditions” on page 130.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

source-pool

Syntax	<code>source-pool nat-pool-name;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the source address pool for translated traffic.
Options	<i>nat-pool-name</i> —Source pool name.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

source-prefix

Syntax	<code>source-prefix source-prefix;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in JUNOS Release 7.6. <i>source-prefix</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the source prefix for translated traffic.
Options	<i>source-prefix</i> —IPv4 or IPv6 source prefix value.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
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 <i>list-name</i> <except>;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the source prefix list for rule matching. You configure the prefix list using the <code>prefix-list</code> statement at the [edit <code>policy-options</code>] hierarchy level.
Options	<i>list-name</i> —Destination prefix list. except—(Optional) Exclude the specified prefix list from rule matching.
Usage Guidelines	See “Configuring NAT Match Conditions” on page 130.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

syslog

Syntax	syslog;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> 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 <code>/var/log</code> directory.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
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 {`
 `applications [application-names];`
 `application-sets [set-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 Rule Content” on page 128.

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 Rule Content” on page 128.

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 NAT Actions” on page 131.

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 166
- translation-type (Twice NAT) on page 166

translation-type (Traditional NAT)

Syntax	translation-type (destination type source type)
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the NAT types for traditional NAT.
Options	<i>type</i> —You can specify source dynamic, source static, or destination static.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
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 <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Specify the NAT types for twice NAT.
Options	<i>type</i> —You must specify destination static and either source dynamic or source static.
Usage Guidelines	See “Configuring NAT Actions” on page 131.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

transport

Syntax	transport [<i>transport-protocols</i>]
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i> pgcp]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the packet gateway to select a NAT pool based on transport protocol type.
Options	[<i>transport-protocol</i>]—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.
Usage Guidelines	See “Configuring NAT Pools for the Packet Gateway” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

Chapter 10

Intrusion Detection Service Configuration Guidelines

The Adaptive Services (AS) or MultiServices PIC supports a limited set of intrusion detection services 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.

The intrusion detection service (IDS) configuration allows you to focus the 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:

```
ids {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      rule {
        applications [ application-names ];
        application-sets [ set-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 describes the following tasks for configuring intrusion detection service:

- Configuring Intrusion Detection Properties on page 171
- Examples: Configuring Intrusion Detection Properties on page 178

Configuring Intrusion Detection Properties

This section describes the following tasks for configuring intrusion detection service:

- Configuring the IDS Rule Set on page 171
- Configuring IDS Rule Content on page 171

Configuring the IDS Rule Set

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:

```
rule-set rule-set-name {
    rule rule-name1;
    rule rule-name2;
    rule rule-name3;
    ...
}
```

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 IDS Rule Content

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 101.

To configure an IDS rule, include the `rule rule-name` statement at the `[edit services ids]` hierarchy level:

```
rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            applications [ application-names ];
            application-sets [ set-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 on page 173
- Configuring IDS Match Conditions on page 173
- Configuring IDS Actions on page 174

Configuring Match Direction

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:

```
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 Services Interfaces” on page 434.

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 IDS Match Conditions

To configure IDS match conditions, include the `from` statement at the `[edit services ids rule rule-name term term-name]` hierarchy level:

```
from {
  applications [ application-names ];
  application-sets [ set-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 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 IDS rule. For an example, see “Examples: Configuring Stateful Firewall Properties” on page 106.

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 53.

- 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 IDS Actions

To configure IDS actions, include the `then` statement at the `[edit services ids rule rule-name term term-name]` hierarchy level:

```
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**:

```

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:

```

(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:

```
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:

```
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.

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 configure more than one source address 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, not 20 connections total. The same logic applies to **destination-address** and **applications** settings.

```
[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:

```
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.

Examples: Configuring Intrusion Detection Properties

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 IDS Actions” on page 174.

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 [<i>set-names</i>];
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more target application sets.
Options	<i>set-name</i> —Name of the target application set.
Usage Guidelines	See “Configuring IDS Match Conditions” on page 173.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

applications

Syntax	applications [<i>application-names</i>];
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more applications to which IDS applies.
Options	<i>application-name</i> —Name of the target application.
Usage Guidelines	See “Configuring IDS Match Conditions” on page 173.
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 IDS Actions” on page 174.

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 IDS Actions” on page 174.

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 IDS Actions” on page 174.

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 (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4. <i>address</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the destination address for rule matching.
Options	<i>address</i> —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 IDS Match Conditions” on page 173.
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 <i>minimum-value</i> high <i>maximum-value</i> <except>;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the destination address range for rule matching.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. except—(Optional) Exempt the specified address range from rule matching.
Usage Guidelines	See “Configuring IDS Match Conditions” on page 173.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-prefix

Syntax	<code>destination-prefix <i>prefix-value</i>;</code>
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the prefix value for destination IPv4 address aggregation.
Options	<i>prefix-value</i> —Integer value. Range: 1 through 32
Usage Guidelines	See “Configuring IDS Actions” on page 174.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-prefix-ipv6

Syntax	<code>destination-prefix-ipv6 <i>prefix</i>;</code>
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Specify the prefix value for destination IPv6 address aggregation.
Options	<i>prefix-value</i> —Integer value. Range: 1 through 128
Usage Guidelines	See “Configuring IDS Actions” on page 174.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-prefix-list

Syntax	<code>destination-prefix-list <i>list-name</i> <except>;</code>
Hierarchy Level	<code>[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]</code>
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the destination prefix list for rule matching. You configure the prefix list using the <code>prefix-list</code> statement at the <code>[edit policy-options]</code> hierarchy level.
Options	<p><i>list-name</i>—Destination prefix list.</p> <p><i>except</i>—(Optional) Exclude the specified prefix list from rule matching.</p>
Usage Guidelines	See “Configuring IDS Match Conditions” on page 173.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

force-entry

Syntax	<code>(force-entry ignore-entry);</code>
Hierarchy Level	<code>[edit services ids rule <i>rule-name</i> term <i>term-name</i> then]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Specify handling of entries in the IDS events cache:</p> <ul style="list-style-type: none"> ■ 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 IDS Actions” on page 174.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

from

Syntax from {
 applications [*application-names*];
 application-sets [*set-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 IDS Match Conditions” on page 173.

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 <i>rate</i> ; }
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> 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 IDS Actions” on page 174.
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 <i>rule-name</i>]
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 IDS Match Conditions” on page 173.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mss

Syntax	<code>mss value;</code>
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> 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	<i>value</i> —MSS value. Default: 1500 Range: 128 through 8192
Usage Guidelines	See “Configuring IDS Actions” on page 174.
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 {
      applications [ application-names ];
      application-sets [ set-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	<i>rule-name</i> —Identifier for the collection of terms that constitute this rule.
Usage Guidelines	See “Configuring IDS Rule Content” on page 171.
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 <i>rule-set-name</i> { [rule <i>rule-names</i>]; }
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	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
Usage Guidelines	See “Configuring the IDS Rule Set” on page 171.
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 “Configuring Intrusion Detection Properties” on page 171.
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 IDS Actions” on page 174.

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 (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4. <i>address</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the source address for rule matching.
Options	<i>address</i> —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 IDS Match Conditions” on page 173.
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 <i>minimum-value</i> high <i>maximum-value</i> <except>;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the source address range for rule matching.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. except—(Optional) Exempt the specified address range from rule matching.
Usage Guidelines	See “Configuring IDS Match Conditions” on page 173.
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 <i>prefix-value</i> ;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the prefix value for source IPv4 address aggregation.
Options	<i>prefix-value</i> —Integer value. Range: 1 through 32
Usage Guidelines	See “Configuring IDS Actions” on page 174.
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 <i>prefix-value</i> ;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> then aggregation]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Specify the prefix value for source IPv6 address aggregation.
Options	<i>prefix-value</i> —Integer value. Range: 1 through 128
Usage Guidelines	See “Configuring IDS Actions” on page 174.
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 <i>list-name</i> <except>;
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the source prefix list for rule matching. You configure the prefix list using the <code>prefix-list</code> statement at the [edit <code>policy-options</code>] hierarchy level.
Options	<i>list-name</i> —Destination prefix list. except—(Optional) Exclude the specified prefix list from rule matching.
Usage Guidelines	See “Configuring IDS Match Conditions” on page 173.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

syn-cookie

Syntax	<pre>syn-cookie { mss <i>value</i>; threshold <i>rate</i>; }</pre>
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> 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 IDS Actions” on page 174.
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 <i>rule-name</i> term <i>term-name</i> 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 IDS Actions” on page 174.
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 {
                applications [ application-names ];
                application-sets [ set-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 Rule Content” on page 171.

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 Rule Content” on page 171.

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

threshold

Syntax	<code>threshold rate;</code>
Hierarchy Level	[edit services ids rule <i>rule-name</i> term <i>term-name</i> then logging], [edit services ids rule <i>rule-name</i> term <i>term-name</i> then syn-cookie]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the threshold for logging or applying SYN-cookie defenses.
Options	<i>rate</i> —Logging threshold number of events per second. <i>rate</i> —SYN-cookie defense number of SYN attacks per second.
Usage Guidelines	See “Configuring IDS Actions” on page 174.
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 Internet Protocol Security (IPSec) services, include the following statements at the [edit services ipsec-vpn] hierarchy level:

```
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 [ values ];
      ipv6_addr [ values ];
      key_id [ values ];
    }
    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 {
            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 ];
}
traceoptions {
    file {
        files number;
        size bytes;
    }
    flag flag;
}

```


This chapter includes the following sections:

- Minimum Security Association Configurations on page 205
- Configuring Security Associations on page 206
- Configuring an IKE Proposal on page 212
- Configuring an IKE Policy on page 215
- Configuring an IPSec Proposal on page 220
- Configuring an IPSec Policy on page 222
- Configuring IPSec Service Rules on page 224
- Configuring Dynamic Endpoints on page 230
- Tracing IPSec Operations on page 236
- Examples: Configuring IPSec Services on page 236

Minimum Security Association Configurations

The following sections show the minimum configurations necessary to set up security associations (SAs) for IPSec services:

- Manual SA Configuration on page 205
- Dynamic SA Configuration on page 205

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:

```
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;
}
```

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:

```
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 {
        proposal [ 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 207.
- **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 211.

This section includes the following topics:

- Configuring Manual Security Associations on page 207
- Configuring Dynamic Security Associations on page 211
- Clearing Security Associations on page 211



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:

```
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 Direction on page 207
- Configuring the Protocol on page 208
- Configuring the Security Parameter Index on page 209
- Configuring the Auxiliary Security Parameter Index on page 209
- Configuring Authentication on page 209
- Configuring Encryption on page 210

Configuring Direction

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:

```

direction (inbound | outbound | bidirectional) {
    ...
}

```

Example: Configuring Inbound and Outbound Direction Statements

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 12345678901234567890abcd;
    }
}

```

Example: Configuring Bidirectional Statement

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

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:

```

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:

```
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:

```
auxiliary-spi auxiliary-spi-value;
```

Configuring Authentication

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:

```
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

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:

```
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 an IKE Policy” on page 215 and “Configuring an IKE Proposal” on page 212. For more information about IPSec policies and proposals, see “Configuring an IPSec Policy” on page 222.

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.

```
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:

```
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 an IKE Proposal

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:

```
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 an IKE Authentication Algorithm on page 212
- Configuring an IKE Authentication Method on page 213
- Configuring an IKE Diffie-Hellman Group on page 213
- Configuring an IKE Encryption Algorithm on page 213
- Configuring the Lifetime for an IKE SA on page 214
- Example: Configuring an IKE Proposal on page 214

Configuring an IKE Authentication Algorithm

To configure an IKE authentication algorithm, include the **authentication-algorithm** statement at the [edit services ipsec-vpn ike proposal *proposal-name*] hierarchy level:

```
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 an IKE Authentication Method

To configure an IKE authentication method, include the **authentication-method** statement at the [edit services ipsec-vpn ike proposal *proposal-name*] hierarchy level:

```
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 an IKE Diffie-Hellman Group

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 an IKE Diffie-Hellman group, include the **dh-group** statement at the [edit services ipsec-vpn ike proposal *proposal-name*] hierarchy level:

```
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 an IKE Encryption Algorithm

To configure an IKE encryption algorithm, include the **encryption-algorithm** statement at the [edit services ipsec-vpn ike proposal *proposal-name*] hierarchy level:

```
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 IKE SA lifetime, include the **lifetime-seconds** statement at the **[edit services ipsec-vpn ike proposal *proposal-name*]** hierarchy level:

```
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 221.

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 an IKE Policy

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:

```

policy policy-name {
  description description;
  local-certificate identifier;
  local-id {
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
  }
  mode (aggressive | main);
  pre-shared-key (ascii-text key | hexadecimal key);
  proposals [ proposal-names ];
  remote-id {
    any-remote-id;
    ipv4_addr [ values ];
    ipv6_addr [ values ];
    key_id [ values ];
  }
}

```

This section includes the following topics:

- Configuring the IKE Policy Mode on page 216
- Configuring IKE Policy Proposals on page 216

- Configuring a Preshared Key on page 216
- Configuring a Local Certificate on page 217
- Configuring an IKE Policy Description on page 218
- Configuring Local and Remote IDs on page 218

For an example of an IKE policy configuration, see “Example: Configuring an IKE Policy” on page 219.

Configuring the IKE Policy Mode

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 IKE policy mode, include the **mode** statement and specify **aggressive** or **main** at the [edit services ipsec-vpn ike policy *policy-name*] hierarchy level:

```
mode (aggressive | main);
```

Configuring IKE Policy Proposals

The IKE policy includes a list of one or more proposals associated with an IKE policy.

To configure an IKE policy proposal, 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 a Preshared Key

When you configure **authentication-method pre-shared-keys** at the [edit services ipsec-vpn ike proposal *proposal-name*] hierarchy level, IKE policy preshared keys authenticate peers; for more information, see “Configuring an IKE Authentication Method” on page 213. 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 an IKE policy preshared key, include the **pre-shared-key** statement and a key at the [edit services ipsec-vpn ike policy *policy-name*] hierarchy level:

```
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 a Local Certificate

When you configure **authentication-method rsa-signatures** 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 an IKE Authentication Method” on page 213. You must identify a local certificate that is sent to the peer during the IKE authentication phase.

To configure a IKE policy local certificate, include the **local-certificate** statement at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level:

```
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:

```
trusted-ca ca-profile;
```

For more information, see “Configuring IPSec Options” on page 438.

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 an IKE Policy Description

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:

```
description description;
```

Configuring Local and Remote IDs

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:

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

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:

```
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” on page 230.

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 an IPSec Proposal

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:

```
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 an Authentication Algorithm on page 220
- Configuring an IPSec Proposal Description on page 220
- Configuring an Encryption Algorithm on page 220
- Configuring the Lifetime for an IPSec SA on page 221
- Configuring the Protocol for the Dynamic SA on page 222

Configuring an Authentication Algorithm

To configure an IPSec authentication algorithm, include the **authentication-algorithm** statement at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level:

```
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 an IPSec Proposal Description

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:

```
description description;
```

Configuring an Encryption Algorithm

To configure an IPSec encryption algorithm, include the **encryption-algorithm** statement at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level:

```
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:

```
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 the 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:

```
protocol (ah | esp | bundle);
```

Configuring an IPsec Policy

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:

```
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 an IPsec Policy Description on page 223
- Configuring Perfect Forward Secrecy on page 223
- Configuring IPsec Policy Proposals on page 223
- Example: IPsec Policy Configuration on page 223

Configuring an IPSec Policy Description

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:

```
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:

```
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 IPSec Policy Proposals

The IPSec policy includes a list of one or more proposals associated with an IPSec policy.

To configure an IPSec policy proposal, include the `proposals` statement and specify one or more proposal names at the `[edit services ipsec-vpn ipsec policy policy-name]` hierarchy level:

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

Example: IPSec Policy Configuration

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 Service Rules

This section describes the following tasks for configuring IPSec services:

- Configuring the IPSec Rule Set on page 224
- Configuring IPSec Rule Content on page 225

Configuring the IPSec Rule Set

The `rule-set` statement defines a collection of Network Address Translation (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 ipsec-vpn]` hierarchy level:

```

rule-set rule-set-name {
  rule rule-name1;
  rule rule-name2;
  rule rule-name3;
  ...
}

```

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 IPSec Rule Content

To configure an IPSec rule, include the **rule** statement and specify a rule name at the [edit services ipsec-vpn] hierarchy level:

```
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      destination-address address;
      ipsec-inside-interface interface-name;
      source-address address;
    }
    then {
      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.

In addition, each rule includes 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 (input | output)` statement at the `[edit services ipsec-vpn rule rule-name]` hierarchy level:

```
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 Services Interfaces” on page 434.

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.

For more information about configuring IPSec rules, see the following sections:

- Configuring IPSec Match Conditions on page 226
- Configuring IPSec Actions on page 227

Configuring IPSec Match Conditions

To configure IPSec match conditions, include the `from` statement at the `[edit services ipsec-vpn rule rule-name term term-name]` hierarchy level:

```
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 Services Interfaces” on page 434 and “Configuring Interface Properties” on page 469.

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.

Configuring IPSec Actions

To configure IPSec actions, include the `then` statement at the `[edit services ipsec-vpn rule rule-name term term-name]` hierarchy level:

```
then {
  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 211.

- 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 207.

You can configure the following additional properties:

- Enabling IPSec Packet Fragmentation on page 228
- Configuring the Remote Address and Backup Remote Address on page 228
- Disabling the Anti-Replay Window on page 229
- Enabling System Log Messages on page 230
- Configuring the Tunnel MTU Value on page 230

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:

```
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 the Remote Address and Backup Remote Address

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:

```
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:

```
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.



NOTE: Configuration of the `backup-remote-gateway` statement is not supported on J-series Services Routers. These routers cannot send DPD Hello messages but can respond to Hello messages sent by the peer.

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:

```
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*.

Disabling the Anti-Replay Window

To disable the IPSec anti-replay 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, anti-replay 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:

```
syslog;
```

Configuring the Tunnel MTU Value

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:

```
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 Dynamic Endpoints

IPsec tunnels can also be established using *dynamic peer* security gateways, in which the remote end of the 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 IKE Policy Mode” on page 216. 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 231
- Dynamic Implicit Rules on page 231
- Reverse Route Insertion on page 232
- Configuring an IKE Access Profile on page 232
- Configuring the Service Set on page 234

- Configuring the Interface Identifier on page 234
- Default IKE and IPSec Proposals on page 235

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.

Dynamic Implicit 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 the Remote Address and Backup Remote Address” on page 228.

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 an IKE Policy” on page 215.

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*.

```
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 *` (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 an IKE Policy” on page 215.

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.

Configuring the 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:

```
service-set name {
  next-hop-service {
    inside-service-interface interface-name;
    outside-service-interface interface-name;
  }
  ipsec-vpn-options {
    local-gateway address;
    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 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 at the **[edit interfaces interface-name]** hierarchy level:

```
unit logical-unit-number {
  dial-options {
    ipsec-interface-id identifier;
    (shared | dedicated);
  }
}
```

Specifying the interface identifier in the `dial-options` statement makes this logical interface part of the pool identified by the `ipsec-interface-id` statement.



NOTE: Only one interface identifier can be specified at a time. You can include the `ipsec-interface-id` statement or the `l2tp-interface-id` statement, but not both.

If you configure **shared** mode, it enables one logical interface to be shared across multiple tunnels. The **dedicated** statement specifies that the logical interface is used in a dedicated mode, which is necessary when you are configuring an IPSec link-type tunnel. You must include the **dedicated** statement when you specify an `ipsec-interface-id` value.

Default IKE and IPSec Proposals

The software includes implicit default IKE and IPSec proposals to match the proposals sent by the dynamic peers. The values are shown in Table 15 on page 235; if more than one value is shown, the first value is the default. For more information on IKE proposals, see “Configuring an IKE Proposal” on page 212; for more information on IPSec proposals, see “Configuring an IPSec Proposal” on page 220.



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

Table 15: Default IKE and IPSec Proposals for Dynamic Negotiations

Statement Name	Values
Implicit IKE Proposal	
authentication-method	pre-shared keys
dh-group	group1, group2
authentication-algorithm	sha1, md5, sha-256
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256
lifetime-seconds	3600 seconds
Implicit IPSec Proposal	
protocol	esp, ah, bundle
authentication-algorithm	hmac-sha1-96, hmac-md5-96
encryption-algorithm	3des-cbc, des-cbc, aes-128, aes-192, aes-256
lifetime-seconds	28,800 seconds (8 hours)

Tracing IPSec Operations

Trace operations track IPSec events and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/kmd`.

To trace IPSec operations, include the `traceoptions` statement at the `[edit services ipsec-vpn]` hierarchy level:

```
traceoptions {
  file {
    files number;
    size bytes;
  }
  flag flag;
}
```

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.

Examples: Configuring IPSec Services

See the following sections:

- [Statically Assigned Tunnels on page 236](#)
- [Dynamically Assigned Tunnels on page 239](#)

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;
  }
}

```

Dynamically Assigned Tunnels

The following examples are based on this network configuration (see Figure 3 on page 240):

- 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 240
- Configuring a Next-Hop Style Service-Set with Policy-Based Tunnels on page 242



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.

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” on page 209.

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 246
- authentication-algorithm (IPSec) on page 246

authentication-algorithm (IKE)

Syntax	authentication-algorithm (md5 sha1 sha-256);
Hierarchy Level	[edit services ipsec-vpn ike proposal <i>proposal-name</i>]
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 an IKE Authentication Algorithm” on page 212.
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 <i>ipsec-proposal-name</i>]
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 an Authentication Algorithm” on page 220.
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 <i>proposal-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an IKE authentication method.
Options	<p>dsa-signatures—Digital signature algorithm (DSA).</p> <p>rsa-signatures—Public key algorithm (supports encryption and digital signatures).</p> <p>pre-shared-keys—A key derived from an out-of-band mechanism; the key authenticates the exchange.</p>
Usage Guidelines	See “Configuring an IKE Authentication Method” on page 213.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

auxiliary-spi

Syntax	auxiliary-spi <i>spi-value</i> ;
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>]
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	<p><i>spi-value</i>—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).</p> <p>Range: 256 through 16,639</p>
Usage Guidelines	See “Configuring the Auxiliary Security Parameter Index” on page 209. For information about SPI, see “Configuring the Security Parameter Index” on page 209 and spi .
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

backup-remote-gateway

Syntax	backup-remote-gateway <i>address</i> ;
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> 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.
<hr/>	
	NOTE: Configuration of the backup-remote-gateway statement is not supported on J-series Services Routers. These routers cannot send DPD Hello messages but can respond to Hello messages sent by the peer.
<hr/>	
Options	<i>address</i> —Backup remote IPv4 or IPv6 address.
Usage Guidelines	See “Configuring the Remote Address and Backup Remote Address” on page 228.
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 <i>rule-name</i> term <i>term-name</i> 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 IPSec Actions” on page 227.
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 211.
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 211.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

description

Syntax	<code>description <i>description</i>;</code>
Hierarchy Level	[edit services ipsec-vpn ike policy <i>policy-name</i>], [edit services ipsec-vpn ike proposal <i>proposal-name</i>], [edit services ipsec-vpn ipsec policy <i>policy-name</i>], [edit services ipsec-vpn ipsec proposal <i>proposal-name</i>]
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 an IKE Policy Description” on page 218.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

destination-address

Syntax	<code>destination-address <i>address</i>;</code>
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the destination address for rule matching.
Options	<i>address</i> —Destination IP address.
Usage Guidelines	See “Configuring IPSec Match Conditions” on page 226.
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 <i>proposal-name</i>]
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 an IKE Diffie-Hellman Group” on page 213.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

direction

Syntax	<pre> direction (inbound outbound bidirectional) { protocol (ah bundle esp); spi <i>spi-value</i>; auxiliary-spi <i>spi-value</i>; authentication { algorithm (hmac-md5-96 hmac-sha1-96); key (ascii-text <i>key</i> hexadecimal <i>key</i>); } encryption { algorithm <i>algorithm</i>; key (ascii-text <i>key</i> hexadecimal <i>key</i>); } }</pre>
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the direction in which manual SAs are applied.
Options	<p>bidirectional—Apply the SA in both directions.</p> <p>inbound—Apply the SA on inbound traffic.</p> <p>outbound—Apply the SA on outbound traffic.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring IPSec Rule Content” on page 225.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

dynamic

Syntax	dynamic { ike-policy <i>policy-name</i> ; ipsec-policy <i>policy-name</i> ; }
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a dynamic IPSec SA.
Options	<p>ike-policy <i>policy-name</i>—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.</p> <p>ipsec-policy <i>policy-name</i>—Name of the IPSec policy. This statement is optional and the default policy is used if none is supplied.</p>
Usage Guidelines	See “Configuring Dynamic Security Associations” on page 211.
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” on page 210.
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 <i>algorithm</i> ;
Hierarchy Level	[edit services ipsec-vpn ike proposal <i>proposal-name</i>], [edit services ipsec-vpn ipsec proposal <i>proposal-name</i>]
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 an IKE Encryption Algorithm” on page 213 and “Configuring an Encryption Algorithm” on page 220.
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 IPSec Rule Content” on page 225.

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 [ values ];
                    ipv6_addr [ values ];
                    key_id [ values ];
                }
                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 an IKE Proposal” on page 212 and “Configuring an IKE Policy” on page 215.

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 <i>rule-name</i> term <i>term-name</i> 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 the Remote Address and Backup Remote Address” on page 228.
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	<pre> ipsec { proposal <i>proposal-name</i> { authentication-algorithm (hmac-md5-96 hmac-sha1-96); description <i>description</i>; encryption-algorithm <i>algorithm</i>; lifetime-seconds <i>seconds</i>; protocol (ah esp bundle); } policy <i>policy-name</i> { description <i>description</i>; perfect-forward-secrecy { keys (group1 group2); } proposal [<i>proposal-names</i>]; } } </pre>
Hierarchy Level	[edit services ipsec-vpn]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure IPSec.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring Security Associations” on page 206.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

ipsec-inside-interface

Syntax	<code>ipsec-inside-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> 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	<i>interface-name</i> —Service interface for internal network.
Usage Guidelines	See “Configuring IPSec Match Conditions” on page 226 or “Configuring Dynamic Endpoints” on page 230.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

lifetime-seconds

Syntax	<code>lifetime-seconds <i>seconds</i>;</code>
Hierarchy Level	[edit services ipsec-vpn ike proposal <i>proposal-name</i>], [edit services ipsec-vpn ipsec proposal <i>proposal-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the lifetime of an IKE or IPSec SA. This statement is optional.
Options	<i>seconds</i> —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 214 and “Configuring the Lifetime for an IPSec SA” on page 221.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

local-certificate

Syntax	<code>local-certificate <i>identifier</i>;</code>
Hierarchy Level	<code>[edit services ipsec-vpn ike policy <i>policy-name</i>]</code>
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	<i>identifier</i> —Name of certificate.
Usage Guidelines	See “Configuring a Local Certificate” on page 217.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

local-id

Syntax	<pre>local-id { ipv4_addr [<i>values</i>]; ipv6_addr [<i>values</i>]; key_id [<i>values</i>]; }</pre>
Hierarchy Level	<code>[edit services ike policy <i>policy-name</i>]</code>
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	<p><code>ipv4_addr [<i>values</i>]</code>—Define one or more IPv4 address identification values.</p> <p><code>ipv6_addr [<i>values</i>]</code>—Define one or more IPv6 address identification values.</p> <p><code>key_id [<i>values</i>]</code>—Define one or more key identification values.</p>
Usage Guidelines	See “Configuring Local and Remote IDs” on page 218.
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 207.

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 IPSec Rule Content” on page 225.

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 <i>policy-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define an IKE policy mode.
Default	main
Options	<p>aggressive—Takes half the number of messages of main mode, has less negotiation power, and does not provide identity protection.</p> <p>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.</p>
Usage Guidelines	See “Configuring the IKE Policy Mode” on page 216.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

no-anti-replay

Syntax	no-anti-replay;
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> 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 “Disabling the Anti-Replay Window” on page 229.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

perfect-forward-secrecy

Syntax	perfect-forward-secrecy { keys (group1 group2); }
Hierarchy Level	[edit services ipsec-vpn ipsec policy <i>policy-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define Perfect Forward Secrecy (PFS). Creates single-use keys. This statement is optional.
Options	<p>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:</p> <ul style="list-style-type: none"> ■ group1—768-bit. ■ group2—1024-bit.
Usage Guidelines	See “Configuring Perfect Forward Secrecy” on page 223.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

policy

See the following sections:

- policy (IKE) on page 264
- policy (IPSec) on page 265

policy (IKE)

Syntax `policy policy-name {
 description description;
 local-certificate identifier;
 local-id {
 ipv4_addr [values];
 ipv6_addr [values];
 key_id [values];
 }
 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 an IKE Policy” on page 215.

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

policy (IPSec)

Syntax	<pre> policy <i>policy-name</i> { description <i>description</i>; perfect-forward-secrecy { keys (group1 group2); } proposals [<i>proposal-names</i>]; } </pre>
Hierarchy Level	[edit services ipsec-vpn ipsec]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define an IPSec policy.
Options	<p><i>policy-name</i>—IPSec policy name.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring an IPSec Policy” on page 222.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

pre-shared-key

Syntax	pre-shared-key (ascii-text <i>key</i> hexadecimal <i>key</i>);
Hierarchy Level	[edit services ike policy <i>policy-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a preshared key for an IKE policy.
Options	<p><i>key</i>—Value of preshared key. The key can be one of the following:</p> <ul style="list-style-type: none"> ■ <i>ascii-text</i>—ASCII text key. ■ <i>hexadecimal</i>—Hexadecimal key.
Usage Guidelines	See “Configuring an IKE Policy” on page 215.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

proposal

See the following sections:

- [proposal \(IKE\)](#) on page 266
- [proposal \(IPSec\)](#) on page 267

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 an IKE Proposal” on page 212.

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

proposal (IPSec)

Syntax	<pre>proposal <i>proposal-name</i> { authentication-algorithm (hmac-md5-96 hmac-sha1-96); description <i>description</i>; encryption-algorithm <i>algorithm</i>; lifetime-seconds <i>seconds</i>; protocol (ah esp bundle); }</pre>
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	<p><i>proposal-name</i>—IPSec proposal name.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring an IPSec Proposal” on page 220.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

proposals

Syntax	proposals [<i>proposal-names</i>];
Hierarchy Level	[edit services ipsec-vpn ike policy <i>policy-name</i>], [edit services ipsec-vpn ipsec policy <i>policy-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a list of proposals to include in the IKE or IPSec policy.
Options	<i>proposal-names</i> —List of IKE or IPSec proposal names.
Usage Guidelines	See “Configuring an IKE Proposal” on page 212.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

protocol

Syntax	protocol (ah esp bundle);
Hierarchy Level	[edit services ipsec-vpn ipsec proposal <i>proposal-name</i>], [edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>]
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” on page 208.
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 <i>address</i> ;
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the remote address to which the IPSec traffic is directed.
Options	<i>address</i> —Remote IPv4 or IPv6 address.
Usage Guidelines	See “Configuring IPSec Actions” on page 227.
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” on page 218.

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

rule

Syntax	<pre> rule <i>rule-name</i> { match-direction (input output); term <i>term-name</i> { from { destination-address <i>address</i>; ipsec-inside-interface <i>interface-name</i>; source-address <i>address</i>; } then { backup-remote-gateway <i>address</i>; clear-dont-fragment-bit; dynamic { ike-policy <i>policy-name</i>; ipsec-policy <i>policy-name</i>; } initiate-dead-peer-detection; manual { direction (inbound outbound bidirectional) { authentication { algorithm (hmac-md5-96 hmac-sha1-96); key (ascii-text <i>key</i> hexadecimal <i>key</i>); } auxiliary-spi <i>spi-value</i>; encryption { algorithm <i>algorithm</i>; key (ascii-text <i>key</i> hexadecimal <i>key</i>); } protocol (ah bundle esp); spi <i>spi-value</i>; } } no-anti-replay; remote-gateway <i>address</i>; syslog; tunnel-mtu <i>bytes</i>; } } } </pre>
Hierarchy Level	<pre> [edit services ipsec-vpn], [edit services ipsec-vpn rule-set <i>rule-set-name</i>] </pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the rule the router uses when applying this service.
Options	<p><i>rule-name</i>—Identifier for the collection of terms that comprise this rule.</p> <p>The remaining statements are explained separately.</p>

Usage Guidelines See “Configuring IPSec Rule Content” on page 225.

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 the IPSec Rule Set” on page 224.

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 203.

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

source-address

Syntax	<code>source-address <i>address</i>;</code>
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the source address for rule matching.
Options	<i>address</i> —Source IP address.
Usage Guidelines	See “Configuring IPSec Match Conditions” on page 226.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

spi

Syntax	<code>spi <i>spi-value</i>;</code>
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then manual direction <i>direction</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the SPI for an SA.
Options	<i>spi-value</i> —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 <code>protocol</code> statement to use the <code>bundle</code> option.
Usage Guidelines	See “Configuring the Security Parameter Index” on page 209.
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 <i>rule-name</i> term <i>term-name</i> 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 IPSec Actions” on page 227.
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 {`
 `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 IPSec Rule Content” on page 225.

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

then

```

Syntax  then {
            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 IPSec Rule Content” on page 225.

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

traceoptions

Syntax traceoptions {
 file {
 files *number*;
 size *bytes*;
 }
 flag *flag*;
 }

Hierarchy Level [edit services ipsec-vpn]

Release Information Statement introduced in JUNOS Release 7.5.

Description Configure IPsec tracing operations. The messages are output 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.

size *bytes*—Maximum trace file size.

Usage Guidelines See “Tracing IPsec Operations” on page 236.

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 <i>bytes</i> ;
Hierarchy Level	[edit services ipsec-vpn rule <i>rule-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Maximum transmission unit (MTU) size for IPSec tunnels. The default MTU size depends on the device type.
Options	<i>bytes</i> —MTU size. Range: 0 through 5012 bytes
Usage Guidelines	See “Configuring the Tunnel MTU Value” on page 230.
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:

```
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 routing platforms.

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 Components on page 280
- L2TP Minimum Configuration on page 281
- Configuring L2TP Group Properties on page 283
- Configuring the Logical Interface Identifier on page 288
- L2TP Redundancy on page 289
- Tracing L2TP Operations on page 290
- Examples: Configuring L2TP Services on page 292

L2TP Services Components

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 "$9$n8HX6A01RhivL1R"; # 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 Group Properties

This section describes the tasks for configuring L2TP tunnel group properties:

- Configuring a Tunnel Group on page 284
- Configuring Access Profiles on page 284
- Configuring Addressing on page 285
- Configuring Window Size on page 285
- Configuring Timers on page 286
- Hiding Attribute-Value Pairs on page 286
- Configuring System Log Properties on page 286

Configuring a Tunnel Group

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.

Configuring Access Profiles

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 292.

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 Addressing

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

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

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

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 Log Properties

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 16 on page 287 lists the severity levels that you can specify in configuration statements at the [edit services l2tp tunnel-group *group-name* syslog host *hostname*] hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 16: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the routing platform 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 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 Logical Interface Identifier

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.

Configuring Multilink on Shared Interfaces

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 "$9$n8HX6A01RhIvL1R"; # SECRET-DATA
            }
        }
    }
    profile u {
        authentication-order radius;
    }
    radius-server {
        192.168.65.63 {
            port 1812;
            secret "$9$Vyb4ZHKPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
        }
    }
}
services {
    l2tp {
        tunnel-group 1 {
            l2tp-access-profile t;
            ppp-access-profile u;
            local-gateway {
                address 10.70.1.1;
            }
            service-interface sp-1/3/0;
        }
        traceoptions {
            flag all;
            debug-level packet-dump;
            filter {
                protocol l2tp;
                protocol ppp;
                protocol radius;
            }
        }
    }
}
}

```

L2TP Redundancy

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 475. For an example configuration, see “Examples: Configuring L2TP Services” on page 292. 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

The following is a complete example of an L2TP configuration 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$3s2690leK8X7VKM7VwgaJn/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 "$9$eq1KWxbwgZUHNdqmTF3uO1Rhr-dsoJDNd"; #
      SECRET-DATA
    group-profile sunnyvale_users;
  }
  authentication-order password;
}
profile sunnyvale_bldg_1_tunnel {
  client test {
    l2tp {
      shared-secret "$9$r3HKvLg4ZUDkX7JGjif5p0BIRS8LN"; # SECRET-DATA
      maximum-sessions-per-tunnel 75;
      interface-id west_shared;
      ppp-authentication chap;
    }
    group-profile sunnyvale_tunnel;
  }
  client production {
    l2tp {
      shared-secret
        "$9$R2QErV8X-goGylVwg4jiTz36/t0BEleWFnRhrlXxbs2aJDHqf3nCP5";
      ppp-authentication chap;
    }
    group-profile sunnyvale_tunnel;
  }
}
[edit services]
l2tp {
  tunnel-group finance-lns-server {
    l2tp-access-profile sunnyvale_bldg_1_tunnel;
    ppp-access-profile sunnyvale_bldg_1;
    local-gateway {
      address 10.1.117.3;
    }
    service-interface sp-1/3/0;
    receive-window 1500;
    maximum-send-window 1200;
    retransmit-interval 5;
    hello-interval 15;
    tunnel-timeout 55;
  }
  traceoptions {
    flag all;
  }
}
[edit interfaces sp-1/3/0]
unit0 {
  family inet;
}
unit 10 {
  dial-options {
    l2tp-interface-id foo-user;
    dedicated;
  }
}

```

```

    }
    family inet;
}
unit 11 {
    dial-options {
        l2tp-interface-id east;
        dedicated;
    }
    family inet;
}
unit 12 {
    dial-options {
        l2tp-interface-id east;
        dedicated;
    }
    family inet;
}
unit 21 {
    dial-options {
        l2tp-interface-id west;
        dedicated;
    }
    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;
}

```

The following example shows an L2TP redundancy configuration:

```

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	<code>facility-override <i>facility-name</i>;</code>
Hierarchy Level	<code>[edit services tunnel-group <i>group-name</i> syslog host <i>hostname</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Override the default facility for system log reporting.
Options	<i>facility-name</i> —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 Log Properties” on page 286.
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 <i>seconds</i> ;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the keepalive timer for L2TP tunnels.
Options	<i>seconds</i> —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” on page 286.
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 <i>name</i>]
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” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

host

Syntax	host <i>hostname</i> { services <i>severity-level</i> ; facility-override <i>facility-name</i> ; log-prefix <i>prefix-value</i> ; }
Hierarchy Level	[edit services tunnel-group <i>group-name</i> syslog]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the hostname for the system logging utility.
Options	<i>hostname</i> —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 Log Properties” on page 286.
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 <i>profile-name</i> ;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
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	<i>profile-name</i> —Identifier for the L2TP connection profile.
Usage Guidelines	See “Configuring Access Profiles” on page 284.
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 <i>address</i> ;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the local IP address for L2TP tunnel.
Options	<i>address</i> —Local IP address.
Usage Guidelines	See “Configuring Addressing” on page 285.
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 <i>prefix-value</i> ;
Hierarchy Level	[edit services tunnel-group <i>group-name</i> syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the system logging prefix value.
Options	<i>prefix-value</i> —System logging prefix value.
Usage Guidelines	See “Configuring System Log Properties” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-send-window

Syntax	<code>maximum-send-window <i>packets</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group <i>name</i>]</code>
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	<i>packets</i> —Maximum number of packets the send window can hold at one time. Default: 32
Usage Guidelines	See “Configuring Window Size” on page 285.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ppp-access-profile

Syntax	<code>ppp-access-profile <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group <i>name</i>]</code>
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	<i>profile-name</i> —Identifier for the PPP profile.
Usage Guidelines	See “Configuring Access Profiles” on page 284.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

receive-window

Syntax	<code>receive-window <i>packets</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group <i>name</i>]</code>
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	<i>packets</i> —Maximum number of packets the receive window can hold at one time. Default: 16
Usage Guidelines	See “Configuring Window Size” on page 285.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

retransmit-interval

Syntax	<code>retransmit-interval <i>seconds</i>;</code>
Hierarchy Level	<code>[edit services l2tp tunnel-group <i>name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the maximum retransmit interval for L2TP tunnels.
Options	<i>seconds</i> —Interval, in seconds, after which the server retransmits data if no acknowledgment is received. Default: 30 seconds
Usage Guidelines	See “Configuring Timers” on page 286.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-interface

Syntax	<code>service-interface sp-fpc/pic/port;</code>
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the adaptive services interface responsible for handling L2TP processing.
Options	<code>sp-fpc/pic/port</code> —AS or MultiServices PIC interface used for L2TP processing.
Usage Guidelines	See “Configuring Addressing” on page 285.
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 304
- services (L2TP System Logging) on page 305

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 279.

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	<code>services severity-level;</code>
Hierarchy Level	<code>[edit services tunnel-group <i>group-name</i> syslog host <i>hostname</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the system logging severity level.
Options	<p><i>severity-level</i>—Assigns a severity level to the facility. Valid entries include:</p> <ul style="list-style-type: none"> ■ <code>alert</code>—Conditions that should be corrected immediately. ■ <code>any</code>—Matches any level. ■ <code>critical</code>—Critical conditions. ■ <code>emergency</code>—Panic conditions. ■ <code>error</code>—Error conditions. ■ <code>info</code>—Informational messages. ■ <code>notice</code>—Conditions that require special handling. ■ <code>warning</code>—Warning messages.
Usage Guidelines	See “Configuring System Log Properties” on page 286.
Required Privilege Level	<p><code>interface</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>

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 Log Properties” on page 286.

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 290.

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 a Tunnel Group” on page 284.

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 <i>seconds</i> ;
Hierarchy Level	[edit services l2tp tunnel-group <i>name</i>]
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	<i>seconds</i> —Interval after which the tunnel is terminated if no data can be sent. Default: 120 seconds
Usage Guidelines	See “Configuring Timers” on page 286.
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) interfaces (lsq) on the Adaptive Services (AS) PIC, the internal Adaptive Services Module in the M7i platform, the MultiServices PIC, and the Link Services II PIC. LSQ interfaces are similar to link services interfaces, which are described in “Configuring Multilink and Link Services Logical Interface Properties” on page 846. 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 312
- LSQ Failure Recovery on page 313
- Link Services IQ CoS Components on page 324
- Configuring Fragmentation by Forwarding Class on page 328
- Configuring Link-Layer Overhead on page 329
- Configuring Multiclass MLPPP on page 330
- Oversubscribing Interface Bandwidth on page 331
- Providing a Guaranteed Minimum Rate on page 335
- Configuring Link Services and CoS on page 338
- Link Services CoS on J-series Services Routers on page 341
- Common Uses for the Link Services IQ Interface on page 342

Layer 2 Service Package Capabilities and Interfaces

As described in “Enabling Service Packages” on page 29, 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 “Link Services IQ CoS Components” on page 324 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 387 and “Configuring Fragmentation by Forwarding Class” on page 328.

- 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 959.



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 “Link Services IQ CoS Components” on page 324.



NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the routing platform 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*.

LSQ Failure Recovery

Link services IQ (**lsq**-) interfaces support several methods for configuring redundancy with failure recovery, depending on whether the LSQ PICs are located in different routers (interchassis) or in the same router (intrachassis), as described in the following sections:

- Interchassis LSQ Failover on page 314
- Stateless Intrachassis LSQ Failover on page 316
- Stateful Intrachassis LSQ Failover on page 316

Interchassis LSQ Failover

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 Failover Options on page 314
- Configuring Interoperability on page 315
- Configuration Restrictions on page 315

Configuring Failover Options

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. Configure the following 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:

```
no-termination-request;
```

This functionality has been extended to the 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.

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. The supported PIC types are as follows:

- Channelized OC3 IQ PICs
- Channelized OC12 IQ PICs
- Channelized STM1 IQ PICs
- Channelized STM4 IQ PICs

Configuring Interoperability

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:

```
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.

Configuration Restrictions

The following restrictions apply to LSQ failure recovery:

- It applies only to Link Services IQ PICs installed in M-series routing platforms, 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.

Stateless Intrachassis LSQ Failover

Stateless switchover from one Link Services IQ PIC to another within the same router can be configured by the same method as in the previous section, using the SONET APS mechanism. 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*.

Stateful Intrachassis LSQ Failover

You can configure LSQ redundancy on M-series and T-series routing platforms that have multiple AS or MultiServices PICs with `lsq` interfaces. To configure LSQ redundancy, you specify 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 Intrachassis LSQ Redundancy on page 317
- LSQ Failure Conditions and Constraints on page 317
- Link PIC Redundancy on page 319
- Examples: Configuring Intrachassis LSQ Redundancy on page 320

Configuring Intrachassis LSQ Redundancy

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 rlsqnumber]` hierarchy level:

```
redundancy-options {
  primary lsq-fpc/pic/port;
  secondary lsq-fpc/pic/port;
  (hot-standby | warm-standby);
}
```

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) rlsqnumber` 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.

LSQ Failure Conditions and Constraints

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 **rlsqnumber** interface configuration, change it, and reactivate it.
- The **rlsqnumber** 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 “Interchassis LSQ Failover” on page 314. 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 “FRF-16 Configuration” on page 324.



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.

Link PIC Redundancy

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:

```
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 Intrachassis LSQ Redundancy

MLPPP Configuration

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;
}
}
}

```

FRF-15 Configuration 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;
        }
    }
}

```

FRF-16 Configuration 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;
        }
    }
}

```

Link Services IQ CoS Components

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 service 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 “Scheduler Buffer Size” on page 325.
- **priority**—The transmit priority (low, high, strict-high); for more information, see “Scheduler Priority” on page 326.
- **shaping-rate**—The subscribed transmit rate; for more information, see “Scheduler Shaping Rate” on page 326.
- **drop-profile-map**—The random early detection (RED) drop profile; for more information, see “Drop Profiles” on page 326.

When you configure MLPPP and FRF.12 on M-series and T-series routing platforms, 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 routing platforms, 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 “Examples: Configuring an NxT1 Bundle Using FRF.16” on page 351. 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 platforms, 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 platforms, **lsq** interfaces do not support DiffServ code point (DSCP) and DSCP-IPv6 rewrite markers.

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 queueing algorithm guarantees enough space in the transmit buffer for two MTU-sized packets.

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.

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.

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:

```
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 Fragmentation by Forwarding Class

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 the MRRU and MTU Values” on page 855.

To configure fragmentation properties 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;
    }
  }
}
```

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 page 330.

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:

```
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 “Common Uses for the Link Services IQ Interface” on page 342.

For Link Services PIC link services interfaces (ls), fragmentation maps are not supported. Instead, you can 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 Link Services Delay-Sensitive Packet Interleaving” on page 851.

Configuring Link-Layer Overhead

Link-layer overhead can cause packet drops on constituent links because of bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information.

By default, 4 percent of the total bundle bandwidth is set aside for link-layer overhead. In most network environments, the average link-layer overhead is 1.6 percent. Therefore, we recommend 4 percent as a safeguard. For more information, see RFC 4814, *Hash and Stuffing: Overlooked Factors in Network Device Benchmarking* (March, 2007).

For link services IQ interfaces (lsq), 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

For link services IQ 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 routing platforms. For more information about the Link Services PIC support of LFI, see “Configuring Link Services Delay-Sensitive Packet Interleaving” on page 851.

For link services IQ interfaces (*lsq*) 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 Voice Services Properties” on page 385.

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 forwarding-class *class-name*] hierarchy level:

```
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

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 routing platform 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 the 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 “Providing a Guaranteed Minimum Rate” on page 335.

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 334.

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:

$$(\text{remaining delay-buffer rate} * \text{shaping rate}) / (\text{sum of shaping rates})$$

The remaining delay-buffer rate is equal to:

$$(\text{interface speed}) - (\text{sum of configured delay-buffer rates})$$

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;
    }
}

```

Providing a Guaranteed Minimum Rate

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: Providing a Guaranteed Minimum Rate” on page 337.

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: Providing a 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:

$$\text{delay-buffer-rate} < [8 \times 64 \text{ Kbps}]: 2 \text{ seconds of delay-buffer-rate}$$

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

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`:

```
service-package layer-2;
```

For more information about AS or MultiServices PIC service packages, see “Enabling Service Packages” on page 29 and “Layer 2 Service Package Capabilities and Interfaces” on page 312.

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 845.

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 845.

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;
    }
  }
}

```

Link Services CoS on J-series Services Routers

Unlike M-series and T-series platforms, J-series Services Routers support per-bundle queuing on link services (ls) interfaces. Link services interfaces for J-series Services Routers behave the same way as link services IQ interfaces (lsq). There are some exceptions, as follows:

- Queue 2 is reserved for voice traffic (LFI) on J-series Services Routers, while all other queues perform fragmentation.
- On J-series Services Routers link services interfaces (ls), you can configure compressed Real-Time Transport Protocol (CRTP) without needing to configure MLPPP. For more information, see “Voice Services Configuration Guidelines” on page 383.
- Traffic from the ls- interface on queue 0, queue 1, and queues 4 through 7 is mapped onto queue 0 of the constituent link.
- The queue 0 buffer size percentage on the constituent links should be configured as a sum of the buffer size percentages for the ls- interface queues (queue 0, queue 1, and queues 4 through 7).

For FRF.16 and MLPPP, link services interfaces on the J-series Services Router work the same way as link services IQ interfaces (lsq) on M-series and T-series platforms.

The following is a sample configuration for a link services interface on a J-series Services Router:

```

interfaces {
  ls-0/0/0 {
    per-unit-scheduler;
    unit 0 {
      encapsulation multilink-ppp;
      fragment-threshold 128;
      interleave-fragments;
      family inet {
        address 10.0.0.10/24;
      }
    }
  }
}

```

```

}
fe-0/0/1 {
  unit 0 {
    family inet {
      address 192.1.1.1/24;
    }
  }
}
se-1/0/0 {
  per-unit-scheduler;
  serial-options {
    clocking-mode dce;
    clocking-rate 2.0mhz;
  }
  unit 0 {
    family mlppp {
      bundle ls-0/0/0.0;
    }
  }
}
se-1/0/1 {
  per-unit-scheduler;
  serial-options {
    clocking-mode dce;
    clocking-rate 2.0mhz;
  }
  unit 0 {
    family mlppp {
      bundle ls-0/0/0.0;
    }
  }
}
}

```

For more information about M-series and T-series link services (ls) interfaces, see “Link and Multilink Services Configuration Guidelines” on page 845. For more information about J-series Services Router configuration and examples, see the *J-series Services Router Advanced WAN Access Configuration Guide*.

Common Uses for the Link Services IQ Interface

The following sections detail some common uses for the link services IQ interface, with configuration guidelines and a configuration example for each case.

- Configuring an NxT1 Bundle Using MLPPP on page 343
- Configuring an NxT1 Bundle Using FRF.16 on page 348
- Configuring a Single Fractional T1 Interface Using MLPPP and LFI on page 353
- Configuring a Single Fractional T1 Interface Using FRF.12 on page 357
- Configuring an NxT1 Bundle Using FRF.15 on page 364
- Configuring a T3 Link for Compressed RTP over MLPPP on page 365

- Configuring a T3 or OC3 Interface Using FRF.12 on page 366
- Configuring an ATM2 IQ Interface Using MLPPP on page 368



NOTE: Link services IQ interfaces support both T1 and E1 physical interfaces. The following examples show T1, but E1 configurations are similar.

Configuring an NxT1 Bundle Using MLPPP

To configure an NxT1 bundle using MLPPP, you aggregate *N* different T1 links into a bundle. The NxT1 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:

```
bundle lsq-fpc/pic/port.logical-unit-number;
```

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:

```
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 platforms and eight associated queues on M320 and T-series platforms. 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 platforms, 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 NxT1 Bundle Using MLPPP” on page 346.



NOTE: For M320 and T-series platforms, 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:

```
per-unit-scheduler;
```

To configure and apply the scheduling policy, include the following statements at the [edit class-of-service] hierarchy level:

```
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 NxT1 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 page 330. For more information about fragmentation maps, see “Configuring Fragmentation by Forwarding Class” on page 328.

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 the MRRU and MTU Values” on page 855.

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 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 an NxT1 Bundle Using FRF.16

To configure an NxT1 bundle using FRF.16, you aggregate *N* different T1 links into a bundle. The NxT1 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;

```

To configure the link services IQ interface properties, include 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 {
  dlcid dlcid-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:

```
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 an NxT1 Bundle Using FRF.16” on page 348.

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 platforms, 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 platforms, 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:

```
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:

```

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 the MRRU and MTU Values” on page 855.

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.

Examples: Configuring an NxT1 Bundle Using FRF.16

Configure an NxT1 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 a Single Fractional T1 Interface 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:

```
bundle lsq-fpc/pic/port.logical-unit-number;
```

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:

```

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 platforms, 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 a Single Fractional T1 Using MLPPP and LFI” on page 356.



NOTE: For M320 and T-series platforms, 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:

```
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:

```
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 Fragmentation by Forwarding Class” on page 328.

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 the MRRU and MTU Values” on page 855.

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 a Single Fractional T1 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 a Single Fractional T1 Interface 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 mfr-end-to-end]` hierarchy level:

```
bundle lsq-fpc/pic/port.logical-unit-number;
```

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:

```

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 platforms, 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 a Single Fractional T1 Interface Using FRF.12” on page 360.



NOTE: For M320 and T-series platforms, 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:

```

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:

```
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 Fragmentation by Forwarding Class” on page 328.

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 the MRRU and MTU Values” on page 855.

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 a Single 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 interface (lsq) 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 {
    dlc1 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 mlfrr-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 an NxT1 Bundle Using FRF.15

Configure an NxT1 bundle using FRF.15 on a link services IQ interface.

FRF.15 is similar to FRF.12, as described in “Configuring a Single Fractional T1 Interface Using FRF.12” on page 357. 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.

```

[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 a T3 Link for Compressed RTP over MLPPP

Bundle 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 383.

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 a T3 or OC3 Interface Using FRF.12

Configure 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 ($N \times DS0$). 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 page 330.

```
[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 a Single Fractional T1 Interface Using FRF.12” on page 357.

Configuring an ATM2 IQ Interface Using MLPPP

Configure an ATM2 IQ interface with MLPPP bundled with link services IQ interfaces. This allows you to configure LFI on ATM virtual circuits.

For this scenario, 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 <i>interface-name</i> 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 “LSQ Failure Recovery” on page 313.
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*;
 no-fragmentation;
 }

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 “Link Services IQ CoS Components” on page 324.

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 <i>bytes</i> ;
Hierarchy Level	[edit class-of-service fragmentation-maps forwarding-class <i>class-name</i>]
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 <i>interface-name</i> unit <i>logical-unit-number</i>] or [edit interfaces <i>interface-name</i> 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	<i>bytes</i> —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 “Link Services IQ CoS Components” on page 324.
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 <i>map-name</i> ;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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	<i>map-name</i> —Name of the fragmentation map.
Usage Guidelines	See “Configuring Fragmentation by Forwarding Class” on page 328.
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;
      multilink-class number;
      no-fragmentation;
    }
  }
}
```

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 “Link Services IQ CoS Components” on page 324.

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 rlsqnumber 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 “Stateful Intrachassis LSQ Failover” on page 316.

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 <i>percent</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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	<i>percent</i> —Percentage of total bundle bandwidth to be set aside for link-layer overhead. Range: 0 through 50 percent Default: 0 percent
Usage Guidelines	See “Link Services IQ CoS Components” on page 324.
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 <i>interface-name</i>]; }
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 “LSQ Failure Recovery” on page 313.
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 <i>number</i> ;
Hierarchy Level	[edit class-of-service fragmentation-maps forwarding-class <i>class-name</i>]
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	<i>number</i> —The multilink class assigned to this forwarding class. Range: 0 through 7 Default: None
Usage Guidelines	See “Configuring Multiclass MLPPP” on page 330.
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 <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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	<i>number</i> —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 page 330.
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 <i>class-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>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.</p> <p>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 Multiprotocol Label Switching (MPLS) traffic, the hash includes all MPLS labels and fields in the payload, whether the MPLS payload is IPv4 or IPv6.</p>
Default	If you do not include this statement, the traffic in forwarding class <i>class-name</i> is fragmented.
Usage Guidelines	See “Link Services IQ CoS Components” on page 324.
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 <i>interface-name</i> ppp-options], [edit interfaces lsq-fpc/pic/port lsq-failure-options]
Release Information	Statement introduced in JUNOS Release 7.4. Supported at the [edit interfaces <i>interface-name</i> ppp-options] hierarchy level in JUNOS Release 8.3.
Description	Inhibit PPP termination-request messages to the remote host if the primary circuit fails.
Usage Guidelines	See “LSQ Failure Recovery” on page 313.
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 <i>interface-name</i>]
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 page 338.
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 <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	<p>Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:</p> <ul style="list-style-type: none"> ■ Channelized OC3 IQ PIC ■ Channelized OC12 IQ PIC ■ Channelized STM1 IQ PIC <p>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.</p>
Usage Guidelines	See “Link PIC Redundancy” on page 319.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

primary

Syntax	<code>primary interface-name;</code>
Hierarchy Level	[edit interfaces <i>rlsnumber</i> redundancy-options]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	Specify the primary Link Services IQ PIC interface.
Options	<i>interface-name</i> —The identifier for the Link Services IQ PIC interface, which must be of the form <i>lsq-fpc/pic/port</i> .
Usage Guidelines	See “Stateful Intrachassis LSQ Failover” on page 316.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

redundancy-options

Syntax	<pre> redundancy-options { hot-standby; primary interface-name; secondary interface-name; warm-standby; } </pre>
Hierarchy Level	[edit interfaces <i>rlsnumber</i>]
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 “Stateful Intrachassis LSQ Failover” on page 316.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

secondary

Syntax	<code>secondary interface-name;</code>
Hierarchy Level	[edit interfaces rlsqnumber redundancy-options]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	Specify the secondary (backup) Link Services IQ PIC interface.
Options	<i>interface-name</i> —The identifier for the Link Services IQ PIC interface, which must be of the form <i>lsq-fpc/pic/port</i> .
Usage Guidelines	See “Stateful Intrachassis LSQ Failover” on page 316.
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 <i>interface-name</i>];
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	<i>interface-name</i> —Name of SONET interface.
Usage Guidelines	See “LSQ Failure Recovery” on page 313.
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 <i>rlsnumber</i> 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 “Stateful Intrachassis LSQ Failover” on page 316.
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 Services Routers use the link services (ls) interface to configure CRTP with MLPPP or PPP logical interface encapsulation. For more information, see “Configuring Compressed RTP with MLPPP Encapsulation” on page 856 and “Configuring Compressed RTP with PPP Encapsulation” on page 857.

J-series Services 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 routing platforms. For more information, see the *J-series Services Router Advanced WAN Access 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 page 338.

Link services IQ interfaces use a bundle configuration. For more information, see “Link Services IQ Interfaces Configuration Guidelines” on page 311 and “Configuring Multilink and Link Services Logical Interface Properties” on page 846.



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 387 and “Configuring Fragmentation by Forwarding Class” on page 328.

To configure voice services interface properties, include the following statements on the network interface:

```
[edit interfaces]
interface-name {
  encapsulation (frame-relay | ppp);
  unit logical-unit-number {
    encapsulation (atm-mlppp-llc | frame-relay-ppp);
    compression-device lsq-fpc/pic/port;
    dynamic-call-admission-control {
      activation-priority value;
      bearer-bandwidth-limit value;
    }
    family mlppp {
      bundle lsq-fpc/pic/port.logical-unit;
    }
  }
}
```

To configure voice services interface properties, include the following statements on the services interface:

```
(lsq | ls)-fpc/pic/port {
  unit logical-unit-number {
    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;
  }
}
```

Voice services do not require a separate service rules configuration.

This chapter contains the following sections:

- Configuring Voice Services Properties on page 385
- Configuring Encapsulation on page 388
- Configuring the Network Interface on page 389
- Configuring VoIP Routing on J-series Services Routers on page 390
- Examples: Configuring Voice Services on page 393

Configuring Voice Services Properties

You define voice service properties such as compression by configuring statements and values for a voice services interface, specified by the interface type **lsq** (**ls** on J-series Services Routers). You can configure the following statements:

```
(lsq | ls)-fpc/pic/port {
  unit logical-unit-number {
    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;
  }
}
```

This section describes the following tasks for configuring voice services properties:

- Configuring the Interface Address on page 385
- Configuring Compression on page 386
- Configuring Delay-Sensitive Packet Interleaving on page 387
- Example: Configuring Compression on page 387

Configuring the Interface Address

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 *interface-name* unit *logical-unit-number* family inet]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Assign an IP address to the interface by configuring the **address** value. The AS or MultiServices PIC supports only IP version 4 (IPv4) addresses configured using 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

You can configure several properties that specify how the 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 *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can configure the following properties at the [edit interfaces *interface-name* unit *logical-unit-number* compression rtp] hierarchy level:

- By default, the maximum number of compressed packets inserted between the transmission of full headers is 255 packets. To configure a different maximum value, include the **f-max-period** statement at the [edit interfaces *interface-name* unit *logical-unit-number* compression rtp] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number compression rtp]
f-max-period number;
```

- To specify the maximum number of RTP contexts to accept during negotiation, include the **maximum-contexts** statement at the [edit interfaces *interface-name* unit *logical-unit-number* compression rtp] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number compression rtp]
maximum-contexts number <force>;
```

The optional **force** statement 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.

- To specify the lower and upper boundaries for a range of UDP destination port values on which RTP compression takes effect, include the **port** statement at the [edit interfaces *interface-name* unit *logical-unit-number* compression rtp] hierarchy level:

```
port {
    minimum port-number;
    maximum port-number;
}
```

Values for *port-number* can range from 0 through 65,535. Within the specified range, the router software applies RTP compression to the traffic.

- To set the queues on which RTP compression takes effect, include the **queues** statement at the [edit interfaces *interface-name* unit *logical-unit-number* compression *rtp*] hierarchy level:

```
queues [ queue-numbers ];
```

You can specify q0, q1, q2, and q3 as queue numbers.

The router applies RTP compression on 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

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

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;
```

For voice services interfaces, the valid options are **atm-mlppp-llc**, **frame-relay-ppp** or **multilink-ppp**.

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*]

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 393.



NOTE: The only protocol type supported with **frame-relay-ppp** encapsulation is family mlppp.

Configuring the Network Interface

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 on page 389
- Configuring the Compression Interface on page 389

Configuring Voice Services Bundles

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 Bundles” on page 864.

Configuring the Compression Interface

To configure the physical interface for PPP encapsulation, you also need to specify the services interface to be used for voice compression; on M-series and T-series routers, the Link Services IQ (Isq) interface type is used, and on J-series Services Routers, the Link Services (Is) interface type.

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 *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring VoIP Routing on J-series Services Routers

J4350 and J6350 Services Routers support voice over IP (VoIP) routing with the Avaya IG550 Integrated Gateway. The following sections provide an overview of the product and the JUNOS configuration tasks for setting it up:

- Functional Components on page 390
- Configuring the VoIP Interface on page 390
- Configuring the Media Gateway Controller List on page 391
- Configuring Dynamic Call Admission Control on page 392

For more information, see the *J-series Services Router Advanced WAN Access Configuration Guide*.

Functional Components

The Avaya IG550 Integrated Gateway consists of the following VoIP modules that can be installed in the PCI Express slots on the J4350 and J6350 Services Routers for providing VoIP connectivity:

- TGM550—Avaya Telephony Gateway Module (TGM) with two analog telephone ports, two analog trunk ports, and one serial port for console access.
- TIM510—Avaya DS1 Telephony Interface Module (TIM) with one T1 or E1 trunk port.
- TIM514—Avaya Analog TIM with four analog telephone ports and four analog trunk ports.
- TIM521—Avaya BRI TIM with four RJ-45 ports for ISDN BRI trunk connections.

The TGM550 works with the TIM510, TIM514, and TIM521 media modules to connect IP and legacy analog telephones and trunks over IP networks.

The telephony services on the TGM550 are controlled by a Media Gateway Controller (MGC)—an Avaya media server running Avaya Communication Manager (CM) call processing software. The TGM550 is managed by an MGC located at headquarters or in a branch office. When the primary MGC is located at a remote location, the TGM550 uses standard local survivability (SLS) for partial MGC backup if the connection to the primary MGC is lost. This enables the J4350 and J6350 Services Routers to provide reliable telephony services to branch offices.

Configuring the VoIP Interface

The TGM PIC represents a single-port interface named `vp-slot/0/0`, where `slot` is the slot number. For example, a TGM PIC installed in slot 1 is named `vp-1/0/0`. To configure the interface, include the following statements:

```
interfaces vp-slot/0/0 {
  unit 0 {
    family inet {
      address address {
        destination address;
```

```

    }
  }
}

```

You can configure only one logical unit (unit 0) on the VoIP interface. You must configure a point-to-point connection between the VoIP interface and the TGM550. To configure the point-to-point connection, specify /32 as the subnet mask in the IPv4 address of the VoIP interface.

The following is a sample configuration:

```

interfaces vp-3/0/0 {
  unit 0 {
    family inet {
      address 10.10.10.1/32 {
        destination 10.10.10.2;
      }
    }
  }
}

```

Configuring the Media Gateway Controller List

To provide telephony services, the TGM550 must be registered with a Media Gateway Controller (MGC) running the Avaya Communication Manager software. You can configure the IP addresses for up to four MGCs to which the TGM550 can connect if a link fails.

The MGC list consists of the IP addresses and the order in which to reestablish the H.248 link. The first MGC on the list is the primary MGC. The TGM550 searches for the primary MGC first. If it cannot connect to the primary MGC or loses its connection to the primary MGC, it attempts to connect to the next MGC on the list, and so on.



NOTE: The MGC list is stored in the TGM 550. It is not written to the JUNOS configuration file.

You configure the MGC list by issuing operational mode commands. To configure the MGC list, perform the following steps:

1. Enter operational mode in the CLI.
2. To configure the IP addresses of the Media Gateway Controllers, enter the following command:

```

user@host> request tgm fpc slot slot-number tgm-configuration
"mgclist=ipaddress1,ipaddress2,ipaddress3,ipaddress4"

```

For example, to configure the primary MGC with the IP address 172.16.0.0 and the secondary and tertiary MGCs with the IP addresses 10.10.10.30 and 10.10.10.40 for the TGM550 installed on slot 2 of the chassis, enter:

```

user@host> request chassis fpc slot 2 tgm-configuration

```

```
"mgclist=172.16.0.0,10.10.10.30,10.10.10.40"
```

3. Log in to the TGM550. For information about logging in to the TGM550, see the *J-series Services Router Advanced WAN Access Configuration Guide*.
4. Ping the IP addresses of the configured MGCs to ensure that the IP addresses are correct. For example, to ping the IP address of the primary MGC, enter:

```
user@host># ping 172.16.0.0
```

5. To verify that the MGC list is configured correctly, see the *J-series Services Router Advanced WAN Access Configuration Guide*.

Configuring Dynamic Call Admission Control

For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation on J4350 and J6350 Services Routers supporting VoIP through the TGM550 media gateway module, you can configure dynamic call admission control (CAC). Dynamic CAC provides enhanced control over WAN bandwidth. When dynamic CAC is configured on an interface responsible for providing call bandwidth, the TGM550 informs the Media Gateway Controller (MGC) of the bandwidth limit available for voice packets on the interface and requests the MGC to block new calls when the bandwidth is exhausted.

Dynamic CAC is useful in situations where a primary link becomes unavailable and a backup link with less bandwidth takes its place. Without dynamic CAC, the MGC cannot detect the switchover to the backup link or the resulting changes in network topology and available bandwidth. The MGC would continue to admit calls at the bandwidth of the primary link, causing network congestion and possible jitter, delay, and loss of calls.

To configure dynamic CAC, you define the bearer bandwidth limit (BBL) and activation priority on each WAN interface responsible for providing call bandwidth.

The Gigabit Ethernet interface is used as the primary link for providing call bandwidth because it has the highest activation priority value. When the Gigabit Ethernet interface is active, the TGM550 reports its BBL value of 3000 Kbps to the MGC. If the Gigabit Ethernet interface fails, the TGM550 automatically switches over to the T1 interface because it has the next highest activation priority. The TGM550 now reports the BBL value of the T1 interface to the MGC. If the T1 interface also fails, the TGM550 switches over to the ISDN BRI interface and reports the BBL value of the ISDN BRI interface to the MGC. Configuring dynamic CAC on multiple WAN interfaces allows the MGC to automatically control the call bandwidth when interfaces responsible for providing call bandwidth are unavailable.

To configure dynamic CAC, include the following configuration statements:

```
interfaces interface-name {
  unit 0 {
    dynamic-call-admission-control {
      activation-priority value;
      bearer-bandwidth-limit value;
    }
  }
}
```

```
}
```

You must configure dynamic CAC on logical unit 0. 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*]

kilobits-per-second is the dynamic CAC BBL—the maximum bandwidth available for voice traffic on the interface. The TGM550 reports the BBL to the MGC. When the call bandwidth exceeds the BBL, the MGC blocks new calls and alerts the user with a busy tone. The default BBL is –1, which indicates that the complete interface bandwidth is available for voice traffic. Configure *bbl 0* to report zero bandwidth for bearer traffic to the MGC or to use the interface for signaling purposes only. The range is from 0 through 9999 Kbps.

priority is the dynamic CAC activation priority value that specifies the order in which interfaces are used for providing call bandwidth. The interface with the highest activation priority value is used as the primary link for providing call bandwidth. If the primary link becomes unavailable, the TGM550 switches over to the next active interface with the highest activation priority value, and so on. The activation priority value range is 0 through 255. The default is 50.



NOTE: Dynamic CAC works in conjunction with the Avaya Communication Manager (CM) Call Admission Control: Bandwidth Limitation (CAC-BL) feature. If you configure dynamic CAC on WAN interfaces, you must also configure CAC-BL on the Avaya CM. For more information about configuring CAC-BL, see the *Administrator Guide for Avaya Communication Manager*.

The following is a configuration example:

```
interfaces ge-0/0/3 {
  unit 0 {
    dynamic-call-admission-control {
      activation-priority 150;
      bearer-bandwidth-limit 1000;
    }
  }
}
```

Examples: Configuring Voice Services

The following is an example of a voice services configuration 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;
}
}

```

The following is an example of a voice services configuration 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;
    }
}
}

```

The following is an example of a voice services configuration using an ATM2 physical interface; it includes the corresponding class-of-service configuration 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.

activation-priority

Syntax	<code>activation-priority <i>priority</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	This statement is used only in conjunction with J4350 and J6350 Services Routers that support voice over IP with the TGM550 media gateway module. For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the dynamic call admission control (dynamic CAC) activation priority value.
Options	<i>priority</i> —The activation priority at which the interface is used for providing call bandwidth. The interface with the highest activation priority value is the primary link for providing call bandwidth. If the primary link becomes unavailable, the TGM550 switches over to the next active interface with the highest activation priority value, and so on. Range: 0 through 255 Default: 50
Usage Guidelines	See “Configuring Dynamic Call Admission Control” on page 392.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

address

Syntax	address <i>address</i> { ... }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interface address.
Options	<i>address</i> —Address of the interface.
Usage Guidelines	See “Configuring the Interface Address” on page 385; for a general discussion of <i>address</i> statement options, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

bearer-bandwidth-limit

Syntax	<code>bearer-bandwidth-limit <i>kilobits-per-second</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	This statement is used only in conjunction with J4350 and J6350 Services Routers that support voice over IP with the TGM550. For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the bearer bandwidth limit (BBL). The BBL is used for dynamic call admission control (dynamic CAC) to provide enhanced control over WAN bandwidth.
Options	<i>kilobits-per-second</i> —The bearer bandwidth limit to be reported to a TGM550, in kilobits per second (Kbps). Range: 0 through 9999 Kbps Default: –1 (dynamic CAC is not enabled on the interface)
Usage Guidelines	See “Configuring Dynamic Call Admission Control” on page 392.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

bundle

Syntax	<code>bundle (<i>lsq-fpc/pic/port</i> ...);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<i>lsq-fpc/pic/port</i> —Name of the voice services interface you are linking.
Usage Guidelines	See “Configuring the Network Interface” on page 389.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

compression

Syntax	<pre> compression { rtp { f-max-period <i>number</i>; maximum-contexts <i>number</i> <force>; queues [<i>queue-numbers</i>]; port { minimum <i>port-number</i>; maximum <i>port-number</i>; } } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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” on page 386.
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 <i>interface-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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” on page 389.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dynamic-call-admission-control

Syntax	dynamic-call-admission-control { activation-priority <i>priority</i> ; bearer-bandwidth-limit <i>kilobits-per-second</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	<p>This statement is used only in conjunction with J4350 and J6350 Services Routers that support voice over IP with the TGM550. For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure dynamic call admission control (CAC). Dynamic CAC provides enhanced control over WAN bandwidth. When dynamic CAC is configured on an interface responsible for providing call bandwidth, the TGM550 informs the Media Gateway Controller (MGC) of the bandwidth limit available for voice packets on the interface and requests that the MGC block new calls when the bandwidth is exhausted.</p> <p>Dynamic CAC must be configured on each J-series Services Router interface responsible for providing call bandwidth.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Dynamic Call Admission Control” on page 392.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

encapsulation

Syntax	encapsulation <i>type</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the logical link-layer encapsulation type.
Options	<p>atm-mlppp-llc—For ATM2 IQ physical interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC encapsulation.</p> <p>frame-relay-ppp—For Frame Relay circuits, use Frame Relay PPP encapsulation.</p> <p>multilink-ppp—By default, voice services logical interfaces use MLPPP encapsulation.</p>
Usage Guidelines	See “Configuring Encapsulation” on page 388; for information about encapsulation statement options used with other interface types, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

f-max-period

Syntax	f-max-period <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<p><i>number</i>—Maximum number of packets.</p> <p>Default: 256</p>
Usage Guidelines	See “Configuring Compression” on page 386.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

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 the Network Interface” on page 389; 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 <i>bytes</i> ;
Hierarchy Level	[edit interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For voice services interfaces, set the fragmentation threshold, in bytes.
Options	<i>bytes</i> —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 387.
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 <i>JUNOS Network Interfaces Configuration Guide</i> .
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 <i>number</i> <force>;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Specify the maximum number of RTP contexts to accept during negotiation.
Options	<i>number</i> —Maximum number of contexts. <i>force</i> —(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” on page 386.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

port

Syntax	port { minimum <i>port-number</i> ; maximum <i>port-number</i> ; }
Hierarchy Level	[edit interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i> compression rtp], [edit logical-systems <i>logical-system-name</i> interfaces <i>lsq-fpc/pic/port</i> unit <i>logical-unit-number</i> 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	<i>minimum port-number</i> —Specify the minimum port number. Range: 0 through 65,535 <i>maximum port-number</i> —Specify the maximum port number. Range: 0 through 65,535
Usage Guidelines	See “Configuring Compression” on page 386.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

queues

Syntax	<code>queues [<i>queue-numbers</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<code>queues <i>queue-numbers</i></code> —Assign one or more of the following queues: q0, q1, q2, and q3.
Usage Guidelines	See “Configuring Compression” on page 386.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rtp

Syntax	<pre> rtp { f-max-period <i>number</i>; maximum-contexts <i>number</i> <force>; queues [<i>queue-numbers</i>]; port { minimum <i>port-number</i>; maximum <i>port-number</i>; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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” on page 386.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

unit

Syntax	<pre> unit <i>logical-unit-number</i> { compression { rtp { f-max-period <i>number</i>; maximum-contexts <i>number</i> <force>; queues [<i>queue-numbers</i>]; port { minimum <i>port-number</i>; maximum <i>port-number</i>; } } } compression-device <i>interface-name</i>; dynamic-call-admission-control { activation-priority <i>priority</i>; bearer-bandwidth-limit <i>kilobits-per-second</i>; } encapsulation type; family family { address <i>address</i> { ... } bundle (lsq-fpc/pic/port ...); } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
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	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Voice Services Properties” on page 385; for a general discussion of logical interface properties, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

Chapter 20

Class-of-Service Configuration Guidelines

To configure class of service (CoS) features on the Adaptive Services (AS) or MultiServices PIC, include the `cos` statements 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 {
        applications [ application-names ];
        application-sets [ set-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 BA classification is not supported on services interfaces.

This chapter contains the following sections:

- Configuring Class-of-Service Properties on page 410
- CoS Hierarchy Configuration Examples on page 415
- Examples: Configuring Class-of-Service Properties on page 417

Configuring Class-of-Service Properties

This section describes the following tasks for configuring CoS on the AS PIC:

- Configuring the Class-of-Service Rule Set on page 410
- Configuring Class-of-Service Rule Content on page 410

Configuring the Class-of-Service Rule Set

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:

```

rule-set rule-set-name {
    rule rule-name1;
    rule rule-name2;
    rule rule-name3;
    ...
}

```

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 Class-of-Service Rule Content

To configure a CoS rule, include the `rule rule-name` statement at the `[edit services cos]` hierarchy level:

```

rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
        from {
            applications [ application-names ];

```

```

        application-sets [ set-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 describe CoS rule content in more detail:

- Configuring Class-of-Service Match Direction on page 411
- Configuring Class-of-Service Match Conditions on page 412
- Configuring Class-of-Service Actions on page 413
- CoS Configuration Restrictions on page 415
- Output Packet Rewriting on page 415

Configuring Class-of-Service Match Direction

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 Services Interfaces” on page 434.

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 Class-of-Service Match Conditions

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

```
from {
  applications [ application-names ];
  application-sets [ set-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 Properties” on page 106.

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 53.

- 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 Class-of-Service Actions

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

```
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 “CoS Hierarchy Configuration Examples” on page 415 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 on page 414
- Configuring Reflexive and Reverse CoS Actions on page 414

Configuring Application Profiles

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:

```
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 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:

```
(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.

CoS Configuration Restrictions

The following restrictions apply to adaptive services CoS configuration:

- 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.

Output Packet Rewriting

On M-series routers, you can configure rewrite rules to change packet header information and attach it to an output interface. These rules can possibly overwrite the DSCP marking configured on the AS or MultiServices PIC. It is important to bear in mind that this adverse effect can take place and to create system-wide configurations carefully.

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 rewrite rules in the output interface.
- Configure the output interface to include the most important mappings.

CoS Hierarchy Configuration Examples

To make settings consistent across Juniper Networks routers, you configure many CoS settings used on the adaptive services interface at the `[edit class-of-service]` hierarchy level. 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 the adaptive services interface. For more information, see the *JUNOS Class of Service Configuration Guide*.



NOTE: The first two configurations, forwarding-class name to forwarding-class ID mapping and forwarding-class name to queue number mapping, are mutually exclusive.

**Forwarding-Class Name
to Forwarding-Class ID
Mapping**

This configuration example illustrates the mapping between the forwarding-class name and the forwarding-class ID:

```
[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;
}
```

**Forwarding-Class Name
to Queue Number
Mapping**

This configuration specifies the mapping between the forwarding-class name and a queue number:

```
[edit class-of-service]
forwarding-classes {
  queue queue-number class-name <priority (high | low)>;
}
```

The following is a configuration example:

```
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;
}
```

Diffserv Code Point Alias to DSCP Bits Mapping

This configuration provides a mapping from an alias name to a DSCP bit value. The alias 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;
  }
}
```

Examples: Configuring Class-of-Service Properties

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;
      }
    }
  }
}
```


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” on page 414.

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

application-sets

Syntax	<code>applications-sets [<i>set-name</i>];</code>
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Define one or more target application sets.
Options	<i>set-name</i> —Name of the target application set.
Usage Guidelines	See “Configuring Class-of-Service Match Conditions” on page 412.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

applications

Syntax	<code>applications [<i>application-name</i>];</code>
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Define one or more applications to which the CoS services apply.
Options	<i>application-name</i> —Name of the target application.
Usage Guidelines	See “Configuring Class-of-Service Match Conditions” on page 412.
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 (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.1. <i>address</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.
Description	Specify the destination address for rule matching.
Options	<i>address</i> —Destination IPv4 or IPv6 address or prefix value.
Usage Guidelines	See “Configuring Class-of-Service Match Conditions” on page 412.
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 <i>list-name</i> <except>;
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the destination prefix list for rule matching. You configure the prefix list using the <i>prefix-list</i> statement at the [edit <i>policy-options</i>] hierarchy level.
Options	<i>list-name</i> —Destination prefix list. except—(Optional) Exclude the specified prefix list from rule matching.
Usage Guidelines	See “Configuring Class-of-Service Match Conditions” on page 412.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

dscp

Syntax	<code>dscp (<i>alias</i> <i>bits</i>);</code>
Hierarchy Level	[edit services cos application-profile <i>profile-name</i> (sip-text sip-video sip-voice)], [edit services cos rule <i>rule-name</i> term <i>term-name</i> then], [edit services cos rule <i>rule-name</i> term <i>term-name</i> 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	<i>alias</i> —Name assigned to a set of CoS markers. <i>bits</i> —Mapping value in the packet header.
Usage Guidelines	See “Configuring Class-of-Service Actions” on page 413.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

forwarding-class

Syntax	<code>forwarding-class <i>class-name</i>;</code>
Hierarchy Level	[edit services cos application-profile <i>profile-name</i> (sip-text sip-video sip-voice)], [edit services cos rule <i>rule-name</i> term <i>term-name</i> then], [edit services cos rule <i>rule-name</i> term <i>term-name</i> then (reflexive reverse)]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Define the forwarding class to which packets are assigned.
Options	<i>class-name</i> —Name of the target application.
Usage Guidelines	See “Configuring Class-of-Service Actions” on page 413.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

from

Syntax	from { applications [<i>application-name</i>]; application-sets [<i>set-name</i>]; destination-address <i>address</i> ; source-address <i>address</i> ; }
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i>]
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 <i>JUNOS Policy Framework Configuration Guide</i> . The remaining statements are explained separately.
Usage Guidelines	See “Configuring Class-of-Service Rule Content” on page 410.
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 <i>rule-name</i>]
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 Class-of-Service Rule Content” on page 410.
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 Actions” on page 414.

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 {
            applications [ application-names ];
            application-sets [ set-names ];
            destination-address address;
            source-address address;
        }
        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 Class-of-Service Rule Content” on page 410.

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

rule-set

Syntax	<code>rule-set <i>rule-set-name</i> { [rule <i>rule-name</i>]; }</code>
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	<i>rule-set-name</i> —Identifier for the collection of rules that constitute this rule set.
Usage Guidelines	See “Configuring the Class-of-Service Rule Set” on page 410.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

services

Syntax	<code>services cos { ... }</code>
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 “Configuring Class-of-Service Properties” on page 410.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

sip-text

Syntax	<pre>sip-text { dscp (<i>alias</i> <i>bits</i>); forwarding-class <i>class-name</i>; }</pre>
Hierarchy Level	[edit services cos application-profile <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	<p>Enable a predefined application profile for handling text data packets.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Application Profiles” on page 414.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

sip-video

Syntax	<pre>sip-video { dscp (<i>alias</i> <i>bits</i>); forwarding-class <i>class-name</i>; }</pre>
Hierarchy Level	[edit services cos application-profile <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	<p>Enable a predefined application profile for handling video data packets.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Application Profiles” on page 414.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

sip-voice

Syntax	<pre>sip-voice { dscp (<i>alias</i> <i>bits</i>); forwarding-class <i>class-name</i>; }</pre>
Hierarchy Level	[edit services cos application-profile <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	<p>Enable a predefined application profile for handling voice data packets.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Application Profiles” on page 414.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

source-address

Syntax	source-address <i>address</i> ;
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	<p>Statement introduced in JUNOS Release 8.1.</p> <p><i>address</i> option enhanced to support IPv4 and IPv6 addresses in JUNOS Release 8.5.</p>
Description	Source address for rule matching.
Options	<i>address</i> —Source IPv4 or IPv6 address or prefix value.
Usage Guidelines	See “Configuring Class-of-Service Match Conditions” on page 412.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

source-prefix-list

Syntax	source-prefix-list <i>list-name</i> <except>;
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the source prefix list for rule matching. You configure the prefix list using the <code>prefix-list</code> statement at the [edit <code>policy-options</code>] hierarchy level.
Options	<i>list-name</i> —Destination prefix list. except—(Optional) Exclude the specified prefix list from rule matching.
Usage Guidelines	See “Configuring Class-of-Service Match Conditions” on page 412.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

syslog

Syntax	syslog;
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i> then], [edit services cos rule <i>rule-name</i> term <i>term-name</i> 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 <code>/var/log</code> directory. This setting overrides any <code>syslog</code> statement setting included in the service set or interface default configuration.
Usage Guidelines	See “Configuring Class-of-Service Actions” on page 413.
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 {
    applications [ application-names ];
    application-sets [ set-names ];
    destination-address address;
    source-address address;
  }
  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 Class-of-Service Rule Content” on page 410.

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 Class-of-Service Actions” on page 413.

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 on the Adaptive Services (AS) or MultiServices PIC. To configure service sets, include the following statements at the [edit services] hierarchy level:

```
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 {
    ike-access-profile profile-name;
    local-gateway address;
    trusted-ca [ ca-profile-name ];
  }
  max-flows number;
  next-hop-service {
    inside-service-interface name.number;
    outside-service-interface name.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 on page 434
- Tracing Adaptive Services or MultiServices PIC Operations on page 446
- Example: Configuring Service Sets on page 449

Configuring Service Sets

This section describes the following tasks for configuring service sets:

- Configuring Services Interfaces on page 434
- Configuring Service Rules on page 438
- Configuring IPSec Options on page 438
- Configuring the Maximum Number of Flows on page 441
- Configuring System Log Properties on page 441
- Allowing Multicast Traffic on page 442
- Service Set Extensions for JUNOS SDK Applications on page 443

Configuring 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 an Interface Service Set on page 434
- Configuring a Next-Hop Service Set on page 435
- Determining Traffic Direction on page 436

Configuring an Interface Service Set

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:

```

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 the **service-set** statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output)] hierarchy level:

```
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 449.

Configuring a Next-Hop Service Set

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 name.number;
  outside-service-interface name.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:

```
[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:

```
[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 169.
- You configure IP Security (IPSec) rules at the `[edit services ipsec-vpn]` hierarchy level; for more information, see “IPSec Services Configuration Guidelines” on page 203.
- 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 123.
- You configure Packet Gateway Control Protocol (PGCP) rules at the `[edit services pgcp]` hierarchy level; for more information, see “Packet Gateway Configuration Guidelines” on page 499.
- You configure stateful firewall rules at the `[edit services stateful-firewall]` hierarchy level; for more information, see “Stateful Firewall Services Configuration Guidelines” on page 101.

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.

Configuring IPSec Options

IPSec service sets require additional specifications that you configure at the `[edit services service-set service-set-name ipsec-vpn-options]` hierarchy level:

```
ike-access-profile profile-name;  
local-gateway address;  
trusted-ca [ ca-profile-name ];
```

Configuration of these statements is described in the following sections:

- Configuring the Local Gateway Address on page 439
- Configuring an IKE Access Profile on page 440
- Configuring Certification Authorities on page 440

Configuring the Local Gateway Address

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 Service Rules” on page 224.

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 an IKE Access Profile

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:

```
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

You can specify one or more trusted certification authorities by including the **trusted-ca** statement:

```
trusted-ca [ ca-profile-name ];
```

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 Service Rules” on page 224.

Configuring the Maximum Number of Flows

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 169.

Configuring System Log Properties

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 17 on page 441 lists the severity levels that you can specify in configuration statements at the **[edit services service-set service-set-name syslog host hostname]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Table 17: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the routing platform to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database

Table 17: System Log Message Severity Levels (*continued*)

Severity Level	Description
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;
```

Allowing 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 “Configuring Network Address Translation for Multicast Traffic” on page 139.

Service Set Extensions for JUNOS SDK Applications

The `extension-service` statement provides a section of the configuration hierarchy at the `[edit services service-set service-set-name]` hierarchy level in which providers of JUNOS SDK applications can add their own services configuration statements using the Database Definition Language (DDL). The `extension-service` statement applies to SDK applications only.

```
[edit]
services {
  service-set service-set-name {
    extension-service service-name {
      provider-specific rules;
    }
  }
}
```

The JUNOS SDK supports transit applications similar to the way in which the JUNOS software supports packet transit. In the JUNOS software, you can configure two types of service sets:

- Interface
- Next-hop, or route-based

For next-hop service sets, routes have to be added to direct traffic to the PIC. Interface service sets have to be attached to media interfaces to direct traffic to service PICs. The following sections provide more detail:

- Interface Service Sets for JUNOS SDK Applications on page 443
- Next-Hop Service Sets for JUNOS SDK Applications on page 444
- Limitations and Constraints for SDK Service Sets on page 445
- Enabling JUNOS SDK Services on page 445

Interface Service Sets for JUNOS SDK Applications

Interface service sets apply to all packets entering and leaving a media interface (in the second example following, the media interface is `ge-0/0/0.0`). After the service set is applied, packets can be injected back to the Packet Forwarding Engine for regular forwarding.

The following example shows an interface service set configured for an SDK application:

```
[edit]
services {
  service-set sset1 {
    extension-service jnx-flow {
      provider-specific rules;
    }
    interface-service { # Existing hierarchy
      service-interface { # Indicates an interface service set
        ms-x/y/0.0; # Specifies which PIC to load this policy on
      }
    }
  }
}
```

```

    }
  }
}

```

This service set can be attached to any media interface of type `family inet`. To associate a defined service set with an interface, include the `service-set` statement at the `[edit interfaces interface-name unit logical-unit-number family inet service (input | output)]` hierarchy level:

```

[edit]
interfaces { # Existing hierarchy
  ge-0/0/0 {
    unit 0 {
      family inet {
        service {
          input {
            service-set sset1; # Attach service set
          }
          output {
            service-set sset1; # Attach service set to output as well
          }
        }
      }
    }
  }
}

```

Next-Hop Service Sets for JUNOS SDK Applications

A next-hop service set is used when a service has to be applied for traffic being forwarded from one VPN routing and forwarding (VRF) table to another.

In the following example, the service set cannot be attached to an interface; instead, two routes are added to direct traffic to the services PIC.

```

[edit]
services {
  service-set sset1 {
    extension-service jnx-flow {
      provider-specific rules;
    }
  }
  # Existing hierarchy
  next-hop-service { # Indicates service set is next-hop service
    inside-service-interface ms-1/2/0.1;
    outside-service-interface ms-1/2/0.2;
  }
}

```


Limitations and Constraints for SDK Service Sets

The following limitations currently apply to service set functionality for the JUNOS SDK:

- You cannot mix JUNOS services and SDK application services you create in the same service set.
- You cannot run JUNOS services and SDK application services on the same MultiServices PIC.
- You cannot mix services from applications developed by different SDK providers on the same PIC.
- You cannot use the same logical interface units to specify `inside-service-interface` and `outside-service-interface` values in different service sets.

Enabling JUNOS SDK Services

Two related statements enable you to configure JUNOS SDK services on a system level:

- By default, the SDK service process (`ssd`) does not run. You turn on `ssd` and enable installation of an SDK application on the router by including the `extensions` statement at the `[edit system]` hierarchy level:

```
[edit]
system {
  extensions {
    providers provider-id;
  }
}
```

You enable installation of a JUNOS SDK application by configuring its provider ID, a short string supplied by the provider. The *provider-id*, as it is called in the command-line interface (CLI), is part of the certificate name. The application's certificate name is used to sign the package during the build process, as explained in the JUNOS SDK installation instructions. When a provider ID is enabled on a router, any SDK application package built using that prefix can be installed and run on the router. Multiple prefixes can be enabled on a router.

- To configure several different types of options for JUNOS SDK applications, include the `extension-provider` statement at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package]` hierarchy level:

```
[edit]
chassis {
  fpc slot-number {
    pic pic-number {
      adaptive-services {
        service-package {
          extension-provider {
            control-cores control-number;
            data-cores data-number;
            object-cache-size value;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```

There are eight cores on a PIC. Some cores, called control cores, are dedicated to run control functionality for the application. Cores dedicated to processing data for the application are called data cores. You must set at least one control core using the **control-cores** option. Setting the **data-cores** option is optional. The total number of cores, control and data cores, that can be set using the **extension-provider** statement ranges from 1 through 8. Any cores not configured as control or data cores are treated as user cores.

The **object-cache-size** statement is used to tune SDK application performance.

In the following example, all but one of the available cores are configured. Three cores are control cores, four are data cores, and one is a user core.

```

[edit]
chassis {
  fpc 0 {
    pic 1 {
      adaptive-services {
        service-package {
          extension-provider {
            control-cores 3;
            data-cores 4;
          }
        }
      }
    }
  }
}

```

Tracing Adaptive Services or MultiServices 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> <size size> <(world-readable | no-world-readable)> <match
  regex>;
flag {
  all;
  command-queued;
  config;
  handshake;
  init;
  interfaces;
  mib;
  removed-client;
  show;
}
```

You configure 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 447
- Configuring the Number and Size of Adaptive Services Log Files on page 447
- Configuring Access to the Log File on page 448
- Configuring a Regular Expression for Lines to Be Logged on page 448
- Configuring the Trace Operations on page 448

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 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 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 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 regex;
```

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 18 on page 449 describes the meaning of the adaptive services tracing flags.

Table 18: Adaptive Services Tracing Flags

Flag	Description	Default Setting
all	Trace all operations.	Off
command-queued	Trace command enqueue events.	Off
config	Log reading of the configuration at the [edit services] hierarchy level.	Off
handshake	Trace handshake events.	Off
init	Trace initialization events.	Off
interfaces	Trace interface events.	Off
mib	Trace GGSN SNMP MIB events.	Off
removed-client	Trace client cleanup events.	Off
show	Trace CLI command servicing.	Off

To display the end of the log, issue the **show log serviced | last** operational mode command:

```
[edit]
user@host# run show log serviced | last
```

Example: Configuring Service Sets

The following example applies **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*> <size *size*> <(world-readable | no-world-readable)>
 <match *regex*>;
 flag *flag*;
 }
}

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 Adaptive Services or MultiServices PIC Operations” on page 446.

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 <i>service-set-name</i>]
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 “Allowing Multicast Traffic” on page 442.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

extension-service

Syntax	extension-service <i>service-name</i> { <i>provider-specific rules</i> ; }
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Define JUNOS SDK service set.
Options	<i>provider-specific rules</i> —Provider-specific hierarchy for services and service sets. <i>service-name</i> —Identifier for a provider-specific service. For more information, see the provider’s documentation.
Usage Guidelines	See “Service Set Extensions for JUNOS SDK Applications” on page 443.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

facility-override

Syntax	<code>facility-override <i>facility-name</i>;</code>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i> syslog host <i>hostname</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Override the default facility for system log reporting.
Options	<p><i>facility-name</i>—Name of the facility that overrides the default assignment. Valid entries are:</p> <ul style="list-style-type: none"> authorization daemon ftp kernel local0 through local7 user
Usage Guidelines	See “Configuring System Log Properties” on page 441.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

host

Syntax	<pre>host <i>hostname</i> { services <i>severity-level</i>; facility-override <i>facility-name</i>; interface-service <i>prefix-value</i>; }</pre>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i> syslog]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the hostname for the system logging utility.
Options	<p><i>hostname</i>—Name of the system logging utility host machine.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring System Log Properties” on page 441.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ids-rules

Syntax	<code>([ids-rules <i>rule-name</i>] ids-rule-sets <i>rule-set-name</i>);</code>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i>]</code>
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	<i>rule-name</i> —Identifier for the collection of terms that constitute this rule. <i>rule-set-name</i> —Identifier for the set of rules to be included.
Usage Guidelines	See “Configuring Service Rules” on page 438.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ike-access-profile

Syntax	<code>ike-access-profile <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i> ipsec-vpn-options]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Define the access profile for the IPSec traffic on dynamic tunnels.
Options	<i>profile-name</i> —Identifier for access profile, which must match the name configured at the <code>[edit access profile <i>name</i> client * ike]</code> hierarchy level.
Usage Guidelines	See “Configuring Dynamic Endpoints” on page 230 or “Configuring IPSec Options” on page 438.
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 <i>name</i> ; }
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
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 <i>name</i> —Name of the service device associated with the interface-wide service set.
Usage Guidelines	See “Configuring Services Interfaces” on page 434.
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 { ike-access-profile <i>profile-name</i> ; local-gateway <i>address</i> ; trusted-ca [<i>ca-profile-name</i>]; }
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
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 438.
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 <i>rule-name</i>] ipsec-vpn-rule-sets <i>rule-set-name</i>);
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
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	<i>rule-name</i> —Identifier for the collection of terms that constitute this rule. <i>rule-set-name</i> —Identifier for the set of rules to be included.
Usage Guidelines	See “Configuring Service Rules” on page 438.
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 <i>address</i> ;
Hierarchy Level	[edit services service-set <i>service-set-name</i> 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	<i>address</i> —Local address.
Usage Guidelines	See “Configuring Service Rules” on page 438.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

log-prefix

Syntax	<code>log-prefix <i>prefix-value</i>;</code>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i> syslog host <i>hostname</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the system logging prefix value.
Options	<i>prefix-value</i> —System logging prefix value.
Usage Guidelines	See “Configuring System Log Properties” on page 441.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

logging

Syntax	<pre>logging { traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable> <match <i>regex</i>>; flag <i>flag</i>; } }</pre>
Hierarchy Level	<code>[edit services]</code>
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 Adaptive Services or MultiServices PIC Operations” on page 446.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

max-flows

Syntax	<code>max-flows <i>number</i>;</code>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Maximum number of flows allowed for the service set.
Options	<i>number</i> —Maximum number of flows.
Usage Guidelines	See “Configuring the Maximum Number of Flows” on page 441.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

nat-rules

Syntax	<code>([nat-rules <i>rule-name</i>] nat-rule-sets <i>rule-set-name</i>);</code>
Hierarchy Level	<code>[edit services service-set <i>service-set-name</i>]</code>
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	<i>rule-name</i> —Identifier for the collection of terms that constitute this rule. <i>rule-set-name</i> —Identifier for the set of rules to be included.
Usage Guidelines	See “Configuring Service Rules” on page 438.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

next-hop-service

Syntax	<pre> next-hop-service { inside-service-interface <i>name.number</i>; outside-service-interface <i>name.number</i>; service-interface-pool <i>name</i>; } </pre>
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
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	<p>inside-service-interface <i>name.number</i>—Name and logical unit number of the service interface associated with the service set applied inside the network.</p> <p>outside-service-interface <i>name.number</i>—Name and logical unit number of the service interface associated with the service set applied outside the network.</p> <p>service-interface-pool <i>name</i>—Name of the pool of logical interfaces configured at the [edit services service-interface-pools pool <i>pool-name</i>] 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.</p>
Usage Guidelines	See “Configuring Services Interfaces” on page 434.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

pgcp-rules

Syntax	([pgcp-rules <i>rule-name</i>] pgcp-rules-sets <i>rule-set-name</i>);
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
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	<i>rule-name</i> —Identifier for the collection of terms that constitute this rule. <i>rule-set-name</i> —Identifier for the set of rules to be included.
Usage Guidelines	See “Configuring Services Interfaces” on page 434.
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 <i>interface-name</i> ;
Hierarchy Level	[edit services service-set <i>service-set-name</i> 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	<i>interface-name</i> —Identifier of the service interface.
Usage Guidelines	See “Configuring Services Interfaces” on page 434.
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 {`
 `([ids-rules rule-names] | ids-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);`
 `interface-service {`
 `service-interface name;`
 `}`
 `ipsec-vpn-options {`
 `ike-access-profile profile-name;`
 `local-gateway address;`
 `trusted-ca [ca-profile-name];`
 `}`
 `max-flows number;`
 `next-hop-service {`
 `inside-service-interface name.number;`
 `outside-service-interface name.number;`
 `}`
 `syslog {`
 `host hostname {`
 `services severity-level;`
 `facility-override facility-name;`
 `interface-service prefix-value;`
 `}`
 `}`
`}`

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 433.

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 462
- services (System Logging) on page 463

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 433.

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

services (System Logging)

Syntax	<code>services severity-level;</code>
Hierarchy Level	<code>[edit services service-set service-set-name syslog host hostname]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the severity level for system logging messages.
Options	<p><i>severity-level</i>—Assigns a severity level to the facility. Valid entries are:</p> <ul style="list-style-type: none"> ■ <i>alert</i>—Conditions that should be corrected immediately. ■ <i>any</i>—Matches any level. ■ <i>critical</i>—Critical conditions. ■ <i>emergency</i>—Panic conditions. ■ <i>error</i>—Error conditions. ■ <i>info</i>—Informational messages. ■ <i>notice</i>—Conditions that require special handling. ■ <i>warning</i>—Warning messages.
Usage Guidelines	See “Configuring System Log Properties” on page 441.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.

stateful-firewall-rules

Syntax	<code>([stateful-firewall-rules <i>rule-names</i>] stateful-firewall-rule-sets <i>rule-set-name</i>);</code>
Hierarchy Level	<code>[edit services service-set service-set-name]</code>
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	<p><i>rule-name</i>—Identifier for the collection of terms that make up this rule.</p> <p><i>rule-set-name</i>—Identifier for the set of rules to be included.</p>
Usage Guidelines	See “Configuring Service Rules” on page 438.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —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 Default System Log Properties” on page 471.

Options The remaining statements are described separately.

Usage Guidelines See “Configuring System Log Properties” on page 441.

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*> <world-readable | no-world-readable> <match
 regex>;
 flag *flag*;
 }

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: *xk* to specify KB, *xm* to specify MB, or *xg* 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 Adaptive Services or MultiServices PIC Operations” on page 446.

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 Options” on page 438.

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

Chapter 24

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 `rspnumber` 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:

- Naming Services Interfaces on page 468
- Configuring Interface Properties on page 469
- Applying Filters and Services to an Interface on page 473
- Configuring AS or MultiServices PIC Redundancy on page 475
- Examples: Configuring a Services Interface on page 478

Naming Services Interfaces

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 routing platform 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 routing platform 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 Interface Properties

This section describes the following tasks for configuring properties specific to adaptive services interfaces:

- Configuring the Interface Address and Domain on page 470
- Configuring Default Timeout Settings on page 470
- Configuring Default System Log Properties on page 471
- Enabling Fragmentation on GRE Tunnels on page 472

Configuring the Interface Address and Domain

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 Services Interfaces” on page 434.

Configuring Default Timeout Settings

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:

```
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 54.

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:

```
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 “Configuring Intrusion Detection Properties” on page 171.

Configuring Default System Log Properties

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 Log Properties” on page 441.

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 19 on page 471 lists the severity levels that you can specify in configuration statements at the [edit interfaces *interface-name* services-options syslog host *hostname*] hierarchy level. The levels from `emergency` through `info` are in order from highest severity (greatest effect on functioning) to lowest.

Table 19: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
emergency	System panic or other condition that causes the routing platform to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors

Table 19: System Log Message Severity Levels (*continued*)

Severity Level	Description
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 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:

```
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 AS or MultiServices PIC GRE tunnels. 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 an Interface

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. To associate a defined service set with an interface, include the **service-set** statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output)] 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 Services Interfaces” on page 434.



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 a Services Interface” on page 478.

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 routing platforms, except TX Matrix platforms, 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 rspnumber]` hierarchy level:

```
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 rlsqnumber]` hierarchy level. For more information, see “LSQ Failure Recovery” on page 313.

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 “Flow-Monitoring Redundancy” on page 716.



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 211.

- 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 rspnumber]` 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 a Services Interface” on page 478.

Examples: Configuring a Services Interface

The following example applies `my-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 `my_post_service` filters.

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

The following example shows configuration of two redundancy interfaces, `rsp0` and `rsp1`, along with their 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 Interface Configuration Statements

The following sections explain each of the 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 Interface Address and Domain” on page 470; 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 <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>]
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 the Adaptive Services (AS) or MultiServices Physical Interface Card (PIC). 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 472.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dial-options

Syntax	<pre>dial-options { ipsec-interface-id <i>name</i>; l2tp-interface-id <i>name</i>; (shared dedicated); }</pre>
Hierarchy Level	<pre>[edit interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i>]</pre>
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	<p>ipsec-interface-id <i>name</i>—Interface identifier for group of dynamic peers. This identifier must be replicated at the <code>[edit access profile <i>name</i> client * ike]</code> hierarchy level.</p> <p>l2tp-interface-id <i>name</i>—Interface identifier that must be replicated at the <code>[edit access profile <i>name</i>]</code> hierarchy level.</p> <p>(shared dedicated)—Specifies whether a logical interface can host one (dedicated) or multiple (shared) sessions at one time.</p>
Usage Guidelines	See “Configuring the Logical Interface Identifier” on page 288 (for L2TP configurations) or “Configuring Dynamic Endpoints” on page 230 (for IPSec configurations).
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

facility-override

Syntax	<code>facility-override <i>facility-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Override the default facility for system log reporting.
Options	<i>facility-name</i> —Name of the facility that overrides the default assignment. Valid entries include: <ul style="list-style-type: none">authorizationdaemonftpkernellocal0 through local7user
Usage Guidelines	See “Configuring Default System Log Properties” on page 471.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

family

Syntax	<pre>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>]; } } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure protocol family information for the logical interface.
Options	<p><i>family</i>—Protocol family. Valid settings for service interfaces include inet (IPv4) and mpls.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring the Interface Address and Domain” on page 470 or “Applying Filters and Services to an Interface” on page 473; for a general discussion of family statement options, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

host

Syntax	host <i>hostname</i> { services <i>severity-level</i> ; facility-override <i>facility-name</i> ; log-prefix <i>prefix-value</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the hostname for the system logging utility.
Options	<i>hostname</i> —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 an Interface” on page 473.
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 <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> 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	<i>seconds</i> —Timeout period. Default: 30 seconds Range: 4 through 86,400 seconds
Usage Guidelines	See “Configuring Default Timeout Settings” on page 470.
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 <i>service-set-name</i> <service-filter <i>filter-name</i> >]; post-service-filter <i>filter-name</i> ; }
Hierarchy Level	[edit interface <i>interface-name</i> unit <i>logical-unit-number</i> family inet service], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 an Interface” on page 473.
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 <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

log-prefix

Syntax	<code>log-prefix <i>prefix-value</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the system logging prefix value.
Options	<i>prefix-value</i> —System logging prefix value.
Usage Guidelines	See “Configuring Default System Log Properties” on page 471.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

open-timeout

Syntax	<code>open-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a timeout period for Transmission Control Protocol (TCP) session establishment.
Options	<i>seconds</i> —Timeout period. Default: 30 seconds Range: 4 through 86,400 seconds
Usage Guidelines	See “Configuring Default Timeout Settings” on page 470.
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 <i>service-set-name</i> <service-filter <i>filter-name</i> >]; }
Hierarchy Level	[edit interface <i>interface-name</i> unit <i>logical-unit-number</i> family inet service], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 an Interface” on page 473.
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 <i>filter-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<i>filter-name</i> —Identifier for the post-service filter.
Usage Guidelines	See “Applying Filters and Services to an Interface” on page 473.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

primary

Syntax	<code>primary interface-name;</code>
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	<i>interface-name</i> —The identifier for the AS or MultiServices PIC interface, which must be of the form <i>sp-fpc/pic/port</i> .
Usage Guidelines	See “Configuring AS or MultiServices PIC Redundancy” on page 475.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

redundancy-options

Syntax	<pre> redundancy-options { primary interface-name; secondary interface-name; } </pre>
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 475.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

secondary

Syntax	<code>secondary interface-name;</code>
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	<i>interface-name</i> —The identifier for the adaptive services interface, which must be of the form <i>sp-fpc/pic/port</i> .
Usage Guidelines	See “Configuring AS or MultiServices PIC Redundancy” on page 475.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service

Syntax	<pre> 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>]; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 an Interface” on page 473.
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 <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 <i>service-set-name</i>] 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 Interface Address and Domain” on page 470.
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 <i>filter-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<i>filter-name</i> —Identifies the filter to be applied in service processing.
Usage Guidelines	See “Applying Filters and Services to an Interface” on page 473.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-set

Syntax	<code>service-set <i>service-set-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<i>service-set-name</i> —Identifies the service set.
Usage Guidelines	See “Applying Filters and Services to an Interface” on page 473.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

services

Syntax	<code>services severity-level;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the system logging severity level.
Options	<p><i>severity-level</i>—Assigns a severity level to the facility. Valid entries include:</p> <ul style="list-style-type: none"> ■ <i>alert</i>—Conditions that should be corrected immediately. ■ <i>any</i>—Matches any level. ■ <i>critical</i>—Critical conditions. ■ <i>emergency</i>—Panic conditions. ■ <i>error</i>—Error conditions. ■ <i>info</i>—Informational messages. ■ <i>notice</i>—Conditions that require special handling. ■ <i>warning</i>—Warning messages.
Usage Guidelines	See “Configuring Default System Log Properties” on page 471.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

services-options

Syntax

```
services-options {
  inactivity-timeout seconds;
  open-timeout seconds;
  syslog {
    host hostname {
      services [ severity-number ];
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
}
```

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 “Configuring Interface Properties” on page 469.

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 Default System Log Properties” on page 471.

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 “Configuring Interface Properties” on page 469; 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

Packet Gateway Configuration Guidelines

To configure packet gateway services, 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;
    maximum-synchronization-time seconds;
  }
  data-inactivity-detection {
    inactivity-delay;
    latch-deadlock-delay seconds;
    send-notification-on-delay;
    inactivity-duration seconds;
    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 {
  normal-mg-execution-time default milliseconds;
  normal-mgc-execution-time default milliseconds;
  mg-provisional-response-timer-value default milliseconds;
  mgc-provisional-response-timer-value default milliseconds;
```

```

        mg-originated-pending-limit default number-of-pendings-received;
        mgc-originated-pending-limit default number-of-pendings-received;
    }
    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;
            mgc-segmentation-timer default milliseconds;
            mg-maximum-pdu-size default bytes;
            mgc-maximum-pdu-size default bytes;
        }
        traffic-management {
            peak-data-rate {
                default bytes-per-second;
                rtcp {
                    fixed-value bytes-per-second;
                    percentage percentage;
                }
            }
            sustained-data-rate {
                default bytes-per-second;
                rtcp {
                    fixed-value bytes-per-second;
                    percentage percentage;
                }
            }
            max-burst-size {
                default bytes;
                rtcp {
                    fixed-value bytes;
                    percentage percentage;
                }
            }
        }
        inactivity-timer {
            inactivity-timeout {
                detect;
                maximum-inactivity-time {
                    default 10-millisecond-units;
                }
            }
        }
    }
}
h248-options {

```



```

audit-observed-events-returns-history;
encoding {
    no-dscp-bit-mirroring;
}
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);
            failure (forced-904 | forced-906 | none);
            graceful (graceful-905 | none);
            link-loss (forced-906 | none);
        }
        virtual-interface-up {
            cancel-graceful (none | restart-918);
            warm (none | restart-900);
        }
    }
    context-indications {
        state-loss (forced-910 | forced-915 | none);
    }
}
wildcard-response-service-change;
}
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 <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
    flag flag;
}
virtual-interface interface-number {
    media-service media-service-name;
    interface interface-identifier;
    routing-instance instance-name {
        service-interface name.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;
}
}

```

This chapter discusses the following topics that provide information about configuring the packet gateway feature:

- Configuring Virtual Packet Gateways on page 503
- Configuring a Connection to the Packet Gateway Controller on page 506
- Configuring NAT Pools for the Packet Gateway on page 507
- Configuring a Media Service on page 507
- Configuring a Virtual Interface on page 507

- Configuring H.248 Timers on page 508
- Configuring Default Values for H.248 Properties on page 508
- Limiting FUF Terms on VPGs on page 510
- Limiting the Rate of Messages Sent from the PIC on page 511
- Enabling Application Layer Gateways for RTP and RTCP for Media Flows on page 511
- Configuring Graceful Restart on page 511
- Detecting Latch Deadlocks and Inactivity Delays on page 511
- Configuring the Service State of a VPG or a Virtual Interface on page 512
- Configuring the Packet Gateway Rule Content on page 513
- Configuring the Packet Gateway Rule Set on page 513
- Configuring Session Mirroring on page 514
- Configuring Overload Control on page 514
- Enabling Wildcards for Service Changes on page 514
- Enabling a History of Media Inactivity Notifications on page 514
- Disabling Bit Mirroring of DSCP Marking on page 515
- Configuring ServiceChange Method and Reason on page 515
- Tracing PGCP Operations on page 516
- Example: Configuring a Packet Gateway on page 516

Configuring Virtual Packet Gateways

To configure a VPG, include the following statements at the [edit services pgcp] hierarchy level:

```
[edit services pgcp]
  gateway gateway-name {
    cleanup-timeout seconds;
    gateway-address gateway-address;
    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;
      maximum-synchronization-time seconds;
    }
    data-inactivity-detection {
      inactivity-delay;
      latch-deadlock-delay seconds;
      send-notification-on-delay;
      inactivity-duration seconds;
    }
  }
```

```

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 {
        normal-mg-execution-time default milliseconds;
        normal-mgc-execution-time default milliseconds;
        mg-provisional-response-timer-value default milliseconds;
        mgc-provisional-response-timer-value default milliseconds;
        mg-originated-pending-limit default number-of-pendings-received;
        mgc-originated-pending-limit default number-of-pendings-received;
    }
    diffserv {
        dscp (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;
        mgc-segmentation-timer default milliseconds;
        mg-maximum-pdu-size default bytes;
        mgc-maximum-pdu-size default bytes;
    }
    traffic-management {
        peak-data-rate {
            default bytes-per-second;
            rtcp {
                fixed-value bytes-per-second;
                percentage percentage;
            }
        }
        sustained-data-rate {
            default bytes-per-second;
            rtcp {
                fixed-value bytes-per-second;
                percentage percentage;
            }
        }
        max-burst-size {
            default bytes;
            rtcp {
                fixed-value bytes;
                percentage percentage;
            }
        }
    }
}

```

```

    }
  }
  inactivity-timer {
    inactivity-timeout {
      detect;
      maximum-inactivity-time {
        default 10-millisecond-units;
      }
    }
  }
}
h248-options {
  audit-observed-events-returns-history;
  encoding {
    no-dscp-bit-mirroring;
  }
  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);
      failure (forced-904 | forced-906 | none);
      graceful (graceful-905 | none);
      link-loss (forced-906 | none);
    }
    virtual-interface-up {
      cancel-graceful (none | restart-918);
      warm (none | restart-900);
    }
  }
  context-indications {
    state-loss (forced-910 | forced-915 | none);
  }
}
wildcard-response-service-change;
}
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);
}

```

You can configure an IP address as the VPG name. However, the IP address is not used in the operation of the VPG.

To configure an IP address for the VPG, include the **gateway-address** statement. This address is the local IP address on which the VPG receives Packet Gateway Control Protocol (PGCP) messages from the packet gateway controller (PGC). You should configure this address on a physical or loopback interface so that it is advertised through the routing protocols.

To configure the number of seconds before the VPG removes gates following a disconnection from the PGC, include the **cleanup-timeout** statement.

Configuring a Connection to the Packet Gateway Controller

You can use the **gateway-controller** statement to define a connection to a PGC. To configure the information, include the **gateway-controller** statement at the **[edit services pgcp gateway gateway-name]** hierarchy level:

```

[edit services pgcp gateway gateway-name]
gateway-controller gateway-controller-name {
    controller-address ip-address;
    controller-port port-number;
    interim-ah-scheme {
        algorithm algorithm;
    }
}

```

You can configure an IP address as the PGC name. However, the IP address is not used for the connection to the PGC.

To configure an IP address for the controller, include the **controller-address** statement.

To set the authentication algorithm used on connections to the PGC to hashed message authentication code (HMAC), include the **interim-ah-scheme** statement.

If the underlying network layer does not support IP Security (IPSec), you can use the interim authentication header (AH) scheme to provide security on the PGCP connection. The interim AH scheme defines an authentication header with the H.248.1 protocol header. To use the interim AH scheme, configure the security algorithm that the interim AH scheme uses. Currently, HMAC null is the only algorithm supported.

Configuring NAT Pools for the Packet Gateway

To configure a network address translation (NAT) pool that is exclusive to the packet gateway feature, include the `pgcp` statement at the `[edit services nat pool pool-name]` hierarchy level.

```
[edit services nat pool pool-name]
pgcp {
  hint hint-string;
  remotely-controlled;
  ports-per-session ports;
  transport;
}
```

By default, the packet gateway chooses the NAT pool using the value of the interface field in the virtual interface. You can configure the packet gateway to choose the NAT pool using the hint value in the Direction field of a nonstandard termination ID. To configure a hint that the packet gateway matches with the hint in the Direction field, include the `hint` statement and specify a specific hint value or a list of hint values.

You can configure separate NAT pools that can be controlled by either the PG or the PGC. By default, the PG controls the addresses and ports in a pool. To configure the addresses and ports in a NAT pool to be explicitly controlled by the PGC, include the `remotely-controlled` statement. The PGC reserves the addresses and ports when it requests specific local NAT bindings for remote addresses.

To configure the number of ports allocated to voice and video flows on the MultiServices PIC, include the `ports-per-session` statement. By default, the number of available ports is 2. This value is useful when one port is allocated for Real-Time Transport Protocol (RTP), and the accompanying Real-Time Control Protocol (RTCP) flow uses the other port. To support the extra ports required for combined voice and video flows, you can specify 4 ports.

Configuring a Media Service

A media service is applied to PGCP packets. A media service specifies the name of a NAT pool that is used on a virtual interface. You need to configure a media service for each virtual interface. To configure a media service, include the `media-service` statement at the `[edit services pgcp]` hierarchy level:

```
[edit services pgcp]
media-service media-service-name {
  nat-pool nat-pool-name;
}
```

Configuring a Virtual Interface

A virtual interface provides the mapping between interface names that appear in the H.248 termination ID in H.248 messages and the following:

- The media service (NAT pools) to be used for a gate.

- A physical interface on the router.

A voice call always travels between two virtual interfaces. The configuration of the two virtual interfaces is used to determine the media services and logical system interface to be used by the call.

To create a virtual interface and configure the media service and the logical interface (that is, a specific unit on a physical interface), include the `virtual-interface` statement at the `[edit services pgcp]` hierarchy level.

```
[edit services pgcp]
virtual-interface name {
  media-service media-service-name;
  interface logical-interface-name;
  routing-instance instance-name {
    service-interface name.number;
  }
  service-state (in-service | out-of-service-forced | out-of-service-graceful);
}
```

Configuring H.248 Timers

To configure H.248 timers for the PGCP connection between the VPG and the PGC, include the `h248-timers` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level:

```
[edit services pgcp gateway gateway-name]
h248-timers {
  initial-average-ack-delay milliseconds;
  maximum-net-propagation-delay milliseconds;
  maximum-waiting-delay milliseconds;
  tmax-retransmission-delay milliseconds;
}
```

See clause 9.2 and annex D of the *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*, for details about these timers.

Configuring Default Values for H.248 Properties

You can configure default values for various H.248 properties. The VPG uses these values unless the PGC overrides them with a PGCP command. You can set default values for properties defined in the following packages:

- The base root package is defined in annex E.2 of the *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*. The properties in this package mostly affect the timers used when the VPG and PGC send and receive provisional responses to H.248 commands.
- The Differentiated Services (Diffserv) package is defined in annex A of the *Gate control protocol, ETSI technical specification TS 102 333 v1.2.0, January 2008*. You can configure a default Differentiated Services code point (DSCP) value that the VPG uses for outgoing traffic when the DSCP value is not already defined by the

PGC. All eight bits are exposed, but only the six leading bits are used by the packet.

- The segmentation package is defined in annex E.14 of the *Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005*. The properties in this package affect the limits used when long H.248 replies are segmented into several H.248 messages.
- The hanging termination detection package is defined in *Gateway control protocol: Hanging Termination Detection package H.248.36, September 2005*.
- The inactivity timer package is defined in *Gateway control protocol: Inactivity timer package H.248.14, March 2002*.
- The Notification Behavior package is defined in annex E.15 of the *Gateway control protocol v3, ITU T Recommendation H.248.1, September 2005*.

To configure default values for H.248 properties, include the **h248-properties** statement at the [edit services pgcp gateway *gateway-name*] hierarchy level:

```
[edit services pgcp gateway gateway-name h248-properties]
application-data-inactivity-detection {
  ip-flow-stop-detection (regulated-notify | immediate-notify);
}
base-root {
  normal-mg-execution-time default milliseconds;
  normal-mgc-execution-time default milliseconds;
  mg-provisional-response-timer-value default milliseconds;
  mgc-provisional-response-timer-value default milliseconds;
  mg-originated-pending-limit default number-of-pendings-received;
  mgc-originated-pending-limit default number-of-pendings-received;
}
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;
  mgc-segmentation-timer default milliseconds;
  mg-maximum-pdu-size default bytes;
  mgc-maximum-pdu-size default bytes;
}
traffic-management {
  peak-data-rate {
    default bytes-per-second;
    rtcp {
```

```

        fixed-value bytes-per-second;
        percentage percentage;
    }
}
sustained-data-rate {
    default bytes-per-second;
    rtcp {
        fixed-value bytes-per-second;
        percentage percentage;
    }
}
max-burst-size {
    default bytes;
    rtcp {
        fixed-value bytes;
        percentage percentage;
    }
}
}
inactivity-timer {
    inactivity-timeout {
        detect;
        maximum-inactivity-time {
            default 10-millisecond-units;
        }
    }
}
}

```

Limiting FUF Terms on VPGs

When fast update filters (FUFs) are installed on the Packet Forwarding Engine and PIC, the Packet Forwarding Engine and PIC discard packets flowing through gates if they exceed the rate limits set in the FUF. To get accurate statistics for gates, the JUNOS software includes these discarded packets in statistics that it collects for gates.

Collecting statistics on packets that are dropped on a gate can impact system performance. To improve performance, the software can limit the number of FUF terms installed on the Packet Forwarding Engine for a VPG. This limit is the maximum value of the following parameters configured on the router:

- The maximum number of FUF terms installed for the VPG
- The maximum percentage of gates with FUF filters relative to all gates currently installed for the VPG

To limit the number of FUF terms installed on the Packet Forwarding Engine for a VPG, include the `fast-update-filters` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level:

```

[edit services pgcp gateway gateway-name]
fast-update-filters {
    maximum-terms number-of-terms;
    maximum-fuf-percentage percentage;
}

```

Limiting the Rate of Messages Sent from the PIC

To limit the rate of messages the PIC sends to the PGC, include the `notification-rate-limit` statement at the `[edit services pgcp]` hierarchy level. This enables the system to maintain a steady state when the PIC generates a large volume of messages.

Enabling Application Layer Gateways for RTP and RTCP for Media Flows

You can enable RTP and RTCP application layer gateways (ALGs) for twice NAT flows created when the PGC installs media gates on the VPG. The ALGs monitor packets on the gate and provide statistics.

You can only enable these ALGs for flows created by the VPG. You cannot enable them within standalone NAT rules.

You can choose to monitor either RTCP or RTP, or both. To enable RTP ALGs or RTCP ALGs, include the `monitor media rtp` statement or the `monitor media rtcp` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level. To enable both ALGs simultaneously, you can include the `monitor media` statement without specifying `rtp` or `rtcp` at the `[edit services pgcp gateway gateway-name]` hierarchy level:

```
monitor {
  media {
    rtcp;
    rtp;
  }
}
```

Configuring Graceful Restart

When the inter-process communication (IPC) connection between the `pgcpd` process and the MultiServices PIC is being established or reestablished, a synchronization procedure is performed. After the synchronization with the PIC is successfully completed, the `pgcpd` process attempts to establish the PGCP connection to the PGC.

To configure graceful restart properties for synchronization with the PIC, include the `graceful-restart` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level:

```
[edit services pgcp gateway gateway-name]
graceful-restart {
  maximum-synchronization-mismatches number-of-mismatches;
  maximum-synchronization-time seconds;
}
```

Detecting Latch Deadlocks and Inactivity Delays

The PG uses data inactivity detection to detect latch deadlocks or other data inactivity on a gate. When a gate is installed, the PG begins inactivity detection after a configurable delay. You can configure two types of delay—a delay for terminations that do not include a latch event and a delay for terminations that do include a latch

event. You can configure a very short delay for latch deadlock detection and a longer delay for other data inactivity.

To configure detection of latch deadlocks on a gate, include the `data-inactivity-detection` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level.

```
[edit services pgcp gateway gateway-name]
data-inactivity-detection {
  inactivity-delay seconds;
  latch-deadlock-delay seconds;
  send-notification-on-delay;
  inactivity-duration seconds;
  stop-detection-on-drop;
  report-service-change {
    service-change-type (forced-906 | forced-910);
  }
}
```

Configuring the Service State of a VPG or a Virtual Interface

You can set the service state of a VPG or a virtual interface. You can set the service state to in-service, or you can perform a forced or graceful shutdown of the VPG or virtual interface.

Configuring the service state of a virtual interface is useful when you do not want to shut down the entire VPG.

Configuring the Service State of a VPG

To set the service state of a VPG, enter the `set service-state` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level:

```
service-state (in-service | out-of-service-forced | out-of-service-graceful);
```

To cause the VPG to be operational and available for traffic, set the service state to in-service. When the VPG is in service, it attempts to connect to the PGC and accepts all PGCP commands from the PGC.

There are two ways to shut down the VPG—forced or graceful:

- **Forced**—The VPG immediately removes all gates and disconnects from the PGC. The packet gateway does not attempt to establish a new connection.
- **Graceful**—The VPG 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 VPG accepts only subtract commands from the PGC.

Configuring the Service State of a Virtual Interface

To set the service state of a virtual interface, enter the `set service-state` statement at the `[edit services pgcp gateway gateway-name virtual-interface]` hierarchy level:

```
service-state (in-service | out-of-service-forced | out-of-service-graceful);
```

To cause the virtual interface to be operational and available for traffic, set the service state to in-service. When the virtual interface is in service, it is connected to the physical interface and accepts all voice calls.

There are two ways to shut down the virtual interface—forced or graceful:

- **Forced**—The virtual interface immediately removes all calls and disconnects from the physical interface. The virtual interface does not attempt to establish a new connection.
- **Graceful**—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 PGC.

Configuring the Packet Gateway Rule Content

To configure a packet gateway rule, include the **rule** statement:

```
rule rule-name {
    gateway gateway-name;
    media-service media-service-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit services pgcp]
- [edit services service set service-set-name]

Each rule consists of the name of the VPG to which the rule is applied and the name of the media service that is applied to the VPG.

Configuring the Packet Gateway Rule Set

The **rule-set** statement defines a collection of rules that determine what actions the packet gateway performs on packets in the data stream. Then, you specify the order of the rules by including the **rule-set** statement:

```
rule-set rule-set-name {
    rule rule-name1;
    rule rule-name2;
    rule rule-name3;
}
```

You can include this statement at the following hierarchy levels:

- [edit services pgcp]
- [edit services service set service-set-name]

The packet gateway processes the rules in the order in which you specify them in the configuration. If a rule matches the packet, the packet gateway performs the corresponding action and the rule processing stops. If no rule matches the packet,

processing continues to the next rule in the rule set. If none of the rules match the packet, the packet is dropped by default.

Configuring Session Mirroring

To configure session mirroring, include the `session-mirroring` statement at the `[edit services pgcp]` hierarchy level or the `[edit services pgcp gateway gateway-name session-mirroring]`:

```
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;
}
```

Configuring Overload Control

To enable the PG to send overload messages to the PGC, include the `queue-limit-percentage` statement at the `[edit services pgcp gateway gateway-name overload-control]` hierarchy level:

```
[edit services pgcp gateway gateway-name overload-control]
queue-limit-percentage percentage;
```

Enabling Wildcards for Service Changes

You can enable the VPG to issue service change commands as wildcard-response commands, which trigger a short response from the PGC. If you do not enable the use of wildcard response for service change commands, the PGC will generate an individual response for every termination that matches the service change command.

To enable the VPG to issue service change commands as wildcard-response commands, enter the `wildcard-response-service-change` statement at the `[edit services pgcp gateway gateway-name h248-options]` hierarchy level:

```
[edit services pgcp gateway gateway-name h248-options]
wildcard-response-service-change;
```

Enabling a History of Media Inactivity Notifications

To enable the PGC to access a history of media inactivity notifications from the VPG, including any not sent to the PGC, include the `audit-observed-events-returns-history` statement at the `[edit services pgcp gateway gateway-name h248-options]` hierarchy level:

```
[edit services pgcp gateway gateway-name h248-options]
```

audit-observed-events-returns-history

Disabling Bit Mirroring of DSCP Marking

The VPG applies bit mirroring to DSCP marking. To disable bit mirroring, include the `no-dscp-bit-mirroring` statement at the `[edit services pgcp gateway gateway-name h248-options encoding]` hierarchy level:

```
[edit services pgcp gateway gateway-name h248-options encoding]
no-dscp-bit-mirroring
```

Configuring ServiceChange Method and Reason

To configure the method and reason that the VPG includes in ServiceChange commands that it sends to the PGC when the state of a control association, virtual interface, or context changes, include the `service-change` statement at the `[edit services pgcp gateway gateway-name h248-options]` hierarchy level.

```
[edit services pgcp gateway gateway-name h248-options]
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);
      failure (forced-904 | forced-906 | none);
      graceful (graceful-905 | none);
      link-loss (forced-906 | none);
    }
    virtual-interface-up {
      cancel-graceful (none | restart-918);
      warm (none | restart-900);
    }
  }
  context-indications {
    state-loss (forced-910 | forced-915 | none);
  }
}
```

Tracing PGCP Operations

To trace PGCP operations and record them in a log file, include the `traceoptions` statement at the `[edit services pgcp]` hierarchy level:

```
[edit services pgcp]
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regex>;
  flag flag;
}
```

You can specify the following PGCP tracing flags:

- `all`—Trace everything.
- `configuration`—Trace configuration events.
- `debug`—Trace debug messages.
- `error`—Trace error messages.
- `firewall`—Trace firewall messages.
- `gate`—Trace gate request and reply events.
- `media-function`—Trace media function events.
- `pgc-connection`—Trace PGC connection events.
- `pgcp-stack-debug`—Trace PGCP stack debug events.
- `pgcp-stack-h248`—Trace PGCP stack H.248 events.
- `pgcp-stack-trace`—Trace PGCP stack function call events.
- `routing-socket`—Trace routing-socket events.
- `trace`—Trace process function calls.

Example: Configuring a Packet Gateway

The following is an example of a packet gateway configuration:

```
[edit services]
pgcp
  gateway vpgOne {
    cleanup-timeout 3600;
    gateway-address 10.20.10.100;
    gateway-port 2944;
    fast-update-filters {
      maximum-terms 1500;
      maximum-fuf-percentage 10;
    }
  }
  gateway-controller pgc1 {
    controller-address 10.10.33.200;
    controller-port 2944;
    interim ah scheme {
      algorithm hmac-null;
    }
  }
```



```

    }
  }
  gateway-controller pgc2 {
    controller-address 10.20.33.300;
    controller-port 2944;
    interim ah scheme {
      algorithm hmac-null;
    }
  }
  graceful-restart {
    maximum-synchronization-mismatches 300;
    maximum-synchronization-time 30;
  }
  data-inactivity-detection {
    inactivity-delay 30;
    latch-deadlock-delay 10;
    send-notification-on-delay;
    inactivity-duration 120;
    stop-detection-on-drop;
    report-service-change {
      service-change-type forced-906;
    }
  }
}
h248-properties {
  application-data-inactivity-detection {
    ip-flow-stop-detection immediate-notify;
  }
  {
    base-root {
      normal-mg-execution-time default 500;
      normal-mgc-execution-time default 500;
      mg-provisional-response-timer default 2000;
      mgc-provisional-response-timer default 2000;
      mg-originated-pending-limit default 4;
      mg-originated-pending-limit default 4;
    }
    diffserv {
      dscp {
        default do-not-change;
      }
    }
  }
  traffic-management {
    sustained-data-rate {
      default 3000000;
      rtcp {
        fixed-value 200000;
      }
    }
    peak-data-rate {
      default 2000000;
      rtcp {
        fixed-value 100000;
      }
    }
  }
  max-burst-size {
    default 3000000;
    rtcp {

```

```

        percentage 1000;
    }
}
segmentation {
    mg-segmentation-timer default 4000;
    mgc-segmentation-timer default 4000;
    mg-maximum-pdu-size default 1472;
    mgc-maximum-pdu-size default 1472;
}
hanging-termination-detection {
    timerx 30;
}
}
h248-options {
    wildcard-response-service-change;
    audit-observed-events-returns-history
}
h248-timers {
    initial-average-ack-delay 4000;
    maximum-net-propagation-delay 40000;
    maximum-waiting-delay 30000;
    tmax-retransmission-delay 25000;
}
}
gateway vpgTwo {
    cleanup-timeout 3600;
    gateway-address 10.20.10.200;
    gateway-port 2944;
    gateway-controller pgc1 {
        controller-address 10.10.33.200;
        controller-port 2944;
        interim ah scheme {
            algorithm hmac-null;
        }
    }
    gateway-controller pgc2 {
        controller-address 10.10.33.300;
        controller-port 2944;
        interim ah scheme {
            algorithm hmac-null;
        }
    }
}
graceful-restart {
    maximum-synchronization-mismatches 300;
    maximum-synchronization-time 30;
}
h248-properties {
    base-root {
        normal-mg-execution-time default 500;
        normal-mgc-execution-time default 500;
        mg-provisional-response-timer default 2000;
        mgc-provisional-response-timer default 2000;
        mg-originated-pending-limit default 4;
        mg-originated-pending-limit default 4;
    }
}

```

```

diffserv {
  dscp {
    default do-not-change;
  }
}
segmentation {
  mg-segmentation-timer-value default 4000;
  mgc-segmentation-timer-value default 4000;
  mg-maximum-pdu-size default 1472;
  mgc-maximum-pdu-size default 1472;
}
traffic-management {
  sustained-data-rate {
    default 3000000;
    rtcp {
      fixed-value 200000;
    }
  }
  peak-data-rate {
    default 2000000;
    rtcp {
      fixed-value 100000;
    }
  }
  max-burst-size {
    default 3000000;
    rtcp {
      percentage 1000;
    }
  }
}
hanging-termination-detection {
  timerx 30;
}
}
h248-timers {
  initial-average-ack-delay 4000;
  maximum-net-propagation-delay 40000;
  maximum-waiting-delay 30000;
  tmax-retransmission-delay 25000;
}
monitor {
  media {
    rtp;
  }
}
}
media-service ms1 {
  nat-pool pgcp-pool;
}
notification-rate-limit 200
overload-control {
  queue-limit-percentage 50
}
rule pgcp-rule {
  gateway pg1;
}

```

```

        media-service ms1;
    }
    traceoptions {
        file pgcpd;
        flag routing-socket;
        flag gate;
        flag media-function;
        flag pgc-connection;
        flag pgcp-stack-h248;
        flag error;
        flag configuration;
    }
    virtual-interface 1 {
        media-service ms1;
        interface fe-1/0/0;
        service-state in-service;
    }
    virtual-interface 2 {
        media-service ms1;
        interface fe-1/1/0;
        service-state in-service;
    }
    session-mirroring {
        delivery-function df-1 {
            destination-address 10.1.1.63;
            destination-port 15000;
            network-operator-id ABCDE;
            source-address 10.1.1.43;
            source-port 10000;
        }
    }
}
nat {
    pool pgcp-pool {
        pgcp;
        address-range low 20.20.20.100/32;
        port automatic;
    }
}

```

Chapter 27

Summary of Packet Gateway Configuration Statements

The following sections explain each of the packet gateway statements. The statements are organized alphabetically.

administrative

See the following sections:

- administrative (Control Association Down) on page 522
- administrative (Virtual Interface Down) on page 523

administrative (Control Association Down)

Syntax	administrative (forced-905 forced-908 none)
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Unregistration Messages in ServiceChange commands that it sends to the PGC when a control association transitions to Out-of-Service because of an administrative operation.
Default	If you do not specify an option, the VPG includes FO/905 (forced-905).
Options	<p>forced-905—Termination is being taken out of service. The VPG is transitioning to Out-of-Service because of an administrative operation.</p> <p>forced-908—Termination is being taken out of service. The VPG is transitioning to Out-of-Service because of an administrative operation or error.</p> <p>none—The VPG does not send a ServiceChange command to the PGC.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

administrative (Virtual Interface Down)

Syntax	administrative (forced-905 forced-906 none)
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when a virtual interface changes to Out-of-Service because of an administrative operation.
Default	If you do not specify an option, the VPG includes FO/905 (forced-905).
Options	<p>forced-905—Termination is being taken out of service. The virtual interface is transitioning to Out-of-Service because of an administrative operation.</p> <p>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.</p> <p>none—VPG does not send a ServiceChange command.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

algorithm

Syntax	algorithm <i>algorithm</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> gateway-controller <i>gateway-controller-name</i> 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	<p>algorithm—Algorithm used for the interim AH scheme. HMAC null is currently the only algorithm supported.</p> <p>Values: hmac-null</p>
Usage Guidelines	See “Configuring a Connection to the Packet Gateway Controller” on page 506.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

application-data-inactivity-detection

Syntax	application-data-inactivity { ip-flow-stop-detection (regulated-notify immediate-notify); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

audit-observed-events-returns-history

Syntax	audit-observed-events-returns-history;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Enable a history of media inactivity events to be viewed by the PGC.
Usage Guidelines	See “Enabling a History of Media Inactivity Notifications” on page 514.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

base-root

Syntax base-root {
 normal-mg-execution-time default *milliseconds*;
 normal-mgc-execution-time default *milliseconds*;
 mg-provisional-response-timer-value default *milliseconds*;
 mgc-provisional-response-timer-value default *milliseconds*;
 mg-originated-pending-limit default *number-of-pendings-received*;
 mgc-originated-pending-limit default *number-of-pendings-received*;
 }

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.

Usage Guidelines See “Configuring Default Values for H.248 Properties” on page 508.

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

cancel-graceful

- cancel-graceful (Control Association Up) on page 526
- cancel-graceful (Virtual Interface Up) on page 527

cancel-graceful (Control Association Up)

Syntax	cancel-graceful (none restart-918);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Notification ServiceChange commands that it sends to the PGC when the control association transitions from the Draining state to the Forwarding state.
Default	If you do not specify an option, the VPG does not send a ServiceChange command.
Options	<p>none—The VPG does not send a ServiceChange command to the PGC.</p> <p>restart-918—The control association has returned to the Forwarding state.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide.</i>

cancel-graceful (Virtual Interface Up)

Syntax	cancel-graceful (none restart-918);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Notification ServiceChange commands that it sends to the PGC when the virtual interface transitions from In-Service to Out-of-Service-Graceful.
Default	If you do not specify an option, the VPG does not send a ServiceChange command.
Options	none—VPG does not send a ServiceChange command. restart-918—Cancel graceful. The virtual interface has entered the Draining state.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

cleanup-timeout

Syntax	cleanup-timeout <i>seconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure the number of seconds before the virtual packet gateway (VPG) automatically deletes all gates following a disconnection from the packet gateway controller (PGC).
Options	<i>seconds</i> —Interval before inactivity detection starts. Range: 0 through 65,535 seconds Default: 3600 seconds
Usage Guidelines	See “Configuring Virtual Packet Gateways” on page 503.
Required Privilege Level	interface-level—To view this statement in the configuration. interface-control—To add this statement to the configuration.

context-indications

Syntax	context-indications { state-loss (forced-910 forced-915 none); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when the gates of a context no longer provide their configured services. When the VPG sends a Service-Interruption message, both terminations in the context become Out-of-Service.
Options	The options are explained separately.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

control-association-indications

Syntax	<pre>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); } }</pre>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in ServiceChange commands that it sends to the PGC when the state of the control association changes.
Options	The options are explained separately.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

controller-address

Syntax	controller-address <i>ip-address</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> gateway-controller <i>gateway-controller-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure an IP address for the PGC.
Options	<i>ip-address</i> —IP address of the PGC.
Usage Guidelines	See “Configuring a Connection to the Packet Gateway Controller” on page 506.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

controller-failure

Syntax	controller-failure (failover-909 restart-902);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Registration Request ServiceChange commands when it attempts to reregister with the PGC or register with a new PGC after the control association is disconnected.
Default	If you do not specify an option, the VPG includes RS/902 (restart-902).
Options	failover-909—PGC impending failure. The VPG is reregistering with a new PGC following a disconnection of the VPG and PGC. restart-902—Warm boot. The VPG is attempting to reregister with existing states after a PGC failure.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

controller-port

Syntax	<code>controller-port <i>port-number</i>;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> gateway-controller <i>gateway-controller-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure the port number of the PGC listening port. The VPG sends H.248 messages to this port.
Options	<i>port-number</i> —Port number of the PGC.
Usage Guidelines	See “Configuring a Connection to the Packet Gateway Controller” on page 506.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

data-inactivity-detection

Syntax	<pre> data-inactivity-detection { inactivity-delay; latch-deadlock-delay <i>seconds</i>; send-notification-on-delay; inactivity-duration <i>seconds</i>; stop-detection-on-drop; report-service-change { service-change-type (forced-906) forced-910); } } </pre>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i>]</code>
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.
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

delivery-function

Syntax	delivery-function <i>delivery-function-name</i> { destination-address <i>destination-address</i> ; destination-port <i>destination-port</i> ; network-operator-id <i>network-operator-id</i> ; source-address <i>source-address</i> ; source-port <i>source-port</i> ; }
Hierarchy Level	[edit services pgcp session-mirroring] [edit services pgcp gateway <i>gateway-name</i> session-mirroring]
Release Information	See “Configuring Session Mirroring” on page 514.
Description	Configure the delivery function that receives the session mirroring information. You can configure only one delivery function.
Options	<i>delivery-function-name</i> —Name of the delivery function that receives the session mirroring information.
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
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	<i>JUNOS Multiplay Solutions Guide</i> .

destination-address

Syntax	destination-address <i>destination-address</i> ;
Hierarchy Level	[edit services pgcp session-mirroring delivery-function <i>delivery-function-name</i>] [edit services pgcp gateway <i>gateway-name</i> session-mirroring delivery-function <i>delivery-function-name</i>]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the address of the delivery function server to which the PG sends session-mirroring information.
Options	<i>destination-address</i> —Address of the server to which the PG sends session-mirroring information.
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
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	<i>JUNOS Multiplay Solutions Guide</i> .

destination-port

Syntax	<code>destination-port destination-port;</code>
Hierarchy Level	[edit services pgcp session-mirroring delivery-function <i>delivery-function-name</i>] [edit services pgcp gateway <i>gateway-name</i> session-mirroring delivery-function <i>delivery-function-name</i>]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the port on the delivery function server that receives session-mirroring information.
Options	<i>destination-port</i> —Port on the delivery function server that receives session-mirroring information. Range: 1 through 65535
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
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	<i>JUNOS Multiplay Solutions Guide</i> .

detect

Syntax	<code>detect;</code>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties inactivity-timer inactivity-timeout]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specifies whether the PG detects inactivity timeout events received from the PG by default.
Default	The PG does not detect inactivity timeout events by default.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

diffserv

Syntax	diffserv { dscp { default (<i>dscp-value</i> <i>alias</i> do-not-change); } }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 <i>Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005</i> .
Options	Statements are explained separately.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable-session-mirroring

Syntax	disable-session-mirroring;
Hierarchy Level	[edit services pgcp session-mirroring] [edit services pgcp gateway <i>gateway-name</i> session-mirroring]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Disable or enable session mirroring on the packet gateway. To disable session mirroring, enter set disable-session-mirroring . To enable session mirroring, enter delete disable-session-mirroring .
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
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	<i>JUNOS Multiplay Solutions Guide</i> .

disconnect

Syntax	disconnect { controller-failure (failover-909 restart-902) reconnect (disconnected-900 restart-902) }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change control-association-indications]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in Registration Request ServiceChange commands when it attempts to reregister with the PGC or register with a new PGC after the control association is disconnected.
Options	The options are explained separately.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

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 VPG includes in Unregistration Messages in ServiceChange commands that it sends to the PGC when a control association transitions to Out-of-Service because of a failure. The failure can be the result of a MultiServices PIC or Flexible PIC Concentrator (FPC) failure, or because the MultiServices PIC was powered off or removed.

Options The options are explained separately.

Usage Guidelines See “Configuring ServiceChange Method and Reason” on page 515.

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	<pre>dscp { default (<i>dscp-value</i> <i>alias</i> do-not-change); }</pre>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties diffserv]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Configure default values for DSCP marking that the VPG uses for outgoing traffic when the DSCP value is not already defined by the PGC.
Default	The default DSCP value that the VPG uses is zero (0x00).
Options	<p><i>dscp-value</i>—Specify a string of eight bits or a 1-byte hexadecimal value using the format: 0xXX. Currently, only six bits are used by the packet.</p> <p><i>alias</i>—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.</p> <p>do-not-change—Specify that no DSCP action should be performed on the PIC. The egress value on the gate is the same as the ingress DSCP value.</p>
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

encoding

Syntax	<pre>encoding { no-dscp-bit-mirroring; }</pre>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Change encoding defaults.
Options	The statement is explained separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

event-timestamp-notification

Syntax	event-timestamp-notification { request-timestamp (requested suppressed autonomous); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Enable or disable the PGC to access timestamp information for media inactivity event notifications.
Options	The statement is explained separately.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

failover-cold

Syntax	failover-cold (failover-920 restart-901);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Registration ServiceChange commands when it attempts to register with a new PGC following a cold failover.
Default	If you do not specify an option, the VPG includes RS/901 (<i>restart-901</i>).
Options	<i>failover-920</i> —Cold failover. The VPG is registering following a graceful Routing Engine switchover. The installed state is reset. <i>restart-901</i> —Cold boot. The VPG is transitioning to In-Service. The previously installed state is not retained.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

failover-warm

Syntax	failover-warm (failover-919 restart-902);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Registration ServiceChange commands when it attempts to register with a new PGC following a warm failover.
Default	If you do not specify an option, the VPG includes RS/902 (restart-902).
Options	<p>failover-919—PGC impending failure. The VPG is registering with a new PGC after the VPG and the PGC were disconnected.</p> <p>restart-902—Warm boot. The VPG is transitioning to In-Service. The previously installed state is retained.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide.</i>

failure

See the following sections:

- failure (Control Association Down) on page 540
- failure (Virtual Interface Down) on page 541

failure (Control Association Down)

Syntax	failure (forced-904 forced-908 none);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG 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 VPG includes FO/904 (forced-904).
Options	<p>forced-904—Termination malfunctioning. The VPG is transitioning to Out-of-Service because of a failure.</p> <p>forced-908—The VPG is transitioning to Out-of-Service due to administrator action or a failure.</p> <p>none—The VPG does not send a ServiceChange command to the PGC.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

failure (Virtual Interface Down)

Syntax	failure (forced-904 forced-906 none);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when a virtual interface transitions to Out-of-Service because of a failure.
Default	If you do not specify an option, the VPG includes FO/904 (forced-904).
Options	<p>forced-904—Termination malfunctioning. The virtual interface is transitioning to Out-of-Service because of an internal failure.</p> <p>forced-906—Loss of lower-layer connectivity. The virtual interface is transitioning to Out-of-Service because of a loss of layer 2 connectivity on the logical or physical interface.</p> <p>none—VPG does not send a ServiceChange command.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

fast-update-filters

Syntax	fast-update-filters { maximum-terms <i>number-of-terms</i> ; maximum-fuf-percentage <i>percentage-of-gates</i> ; }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Limit the number of FUF terms installed on the Packet Forwarding Engine for a VPG 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.
Usage Guidelines	See “Limiting FUF Terms on VPGs” on page 510.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

gateway

```

Syntax  pgcp {
    gateway gateway-name {
        cleanup-timeout seconds;
        gateway-address gateway-address;
        gateway-controller gateway-controller-name {
            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;
        maximum-synchronization-time seconds;
    }
    data-inactivity-detection {
        inactivity-delay seconds;
        latch-deadlock-delay seconds;
        send-notification-on-delay;
        inactivity-duration seconds;
        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 {
            normal-mg-execution-time default milliseconds;
            normal-mgc-execution-time default milliseconds;
            mg-provisional-response-timer-value default milliseconds;
            mgc-provisional-response-timer-value default milliseconds;
            mg-originated-pending-limit default number-of-pendings-received;
            mgc-originated-pending-limit default number-of-pendings-received;
        }
        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;
    mgc-segmentation-timer default milliseconds;
    mg-maximum-pdu-size default bytes;
    mgc-maximum-pdu-size default bytes;
}
traffic-management {
    peak-data-rate {
        default bytes-per-second;
        rtcp {
            fixed-value bytes-per-second;
            percentage percentage;
        }
    }
    sustained-data-rate {
        default bytes-per-second;
        rtcp {
            fixed-value bytes-per-second;
            percentage percentage;
        }
    }
    max-burst-size {
        default bytes;
        rtcp {
            fixed-value bytes;
            percentage percentage;
        }
    }
}
inactivity-timer {
    inactivity-timeout {
        detect;
        maximum-inactivity-time {
            default 10-millisecond-units;
        }
    }
}
}
h248-options {
    audit-observed-events-returns-history;
    encoding {
        no-dscp-bit-mirroring;
    }
    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);
            failure (forced-904 | forced-906 | none);
            graceful (graceful-905 | none);
            link-loss (forced-906 | none);
        }
        virtual-interface-up {
            cancel-graceful (none | restart-918);
            warm (none | restart-900);
        }
    }
    context-indications {
        state-loss (forced-910 | forced-915 | none);
    }
}
wildcard-response-service-change;
}
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. traffic-management option introduced in JUNOS Release 9.2. overload-control option introduced in JUNOS Release 9.3. data-inactivity-detection option introduced in JUNOS Release 9.3. audit-observed-events-returns-history option introduced in JUNOS Release 9.3. encoding option introduced in JUNOS 9.3. event-timestamp-notification option introduced in JUNOS 9.3. service-change option introduced in JUNOS 9.3.
Description	Configure a VPG on the router.
Options	<i>gateway-name</i> —Identifier of the VPG. You can configure an IP address as the gateway name. However, the IP address is not used in the operation of the VPG. The remaining statements are explained separately.
Usage Guidelines	See “Configuring Virtual Packet Gateways” on page 503.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

gateway-address

Syntax	<i>gateway-address gateway-address;</i>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure the IP address of the VPG.
Options	<i>gateway-address</i> —IP address of the VPG that you are configuring on the router.
Usage Guidelines	See “Configuring Virtual Packet Gateways” on page 503.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

gateway-controller

Syntax	gateway-controller <i>gateway-controller-name</i> { controller-address <i>ip-address</i> ; controller-port <i>port-number</i> ; interim-ah-scheme { algorithm <i>algorithm</i> ; } }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure a packet gateway controller.
Options	<i>gateway-controller-name</i> —Name of the PGC. You can configure an IP address as the PGC name. However, the IP address is not used for the connection to the PGC. The remaining statements are explained separately.
Usage Guidelines	See “Configuring a Connection to the Packet Gateway Controller” on page 506.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

gateway-port

Syntax	gateway-port <i>gateway-port</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure a port number for the VPG.
Options	<i>gateway-port</i> —Port number of the VPG that you are configuring on the router. Range: 0 through 65,535 Default: 2944
Usage Guidelines	See “Configuring Virtual Packet Gateways” on page 503.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

graceful

See the following sections:

- graceful (Control Association Down) on page 547
- graceful (Virtual Interface Down) on page 548

graceful (Control Association Down)

Syntax	graceful (graceful-905 none);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Notification ServiceChange commands that it sends to the PGC when the control association transitions from In-Service to Out-of-Service-Graceful.
Default	If you do not specify an option, the VPG does not send a ServiceChange command.
Options	<p>graceful-905—Termination is being taken out of service. The control association has entered the Draining state.</p> <p>none—The VPG does not send a ServiceChange command to the PGC.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide.</i>

graceful (Virtual Interface Down)

Syntax	<code>graceful (graceful-905 none);</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> h248-options service-change virtual-interface-indications virtual-interface-down]</code>
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in Notification ServiceChange commands that it sends to the PGC when the virtual interface transitions from In-Service to Out-of-Service-Graceful.
Default	If you do not specify an option, the VPG does not send a ServiceChange command.
Options	<p><code>graceful-905</code>—Termination is being taken out of service. The interface has entered the Draining state.</p> <p><code>none</code>—VPG does not send a ServiceChange command.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide.</i>

graceful-restart

Syntax	<pre>graceful-restart { maximum-synchronization-mismatches <i>seconds</i>; maximum-synchronization-time <i>seconds</i>; }</pre>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i>]</code>
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.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Graceful Restart” on page 511.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

h248-options

Syntax

```
h248-options {
  audit-observed-events-returns-history;
  encoding {
    no-dscp-bit-mirroring;
  }
  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);
      failure (forced-904 | forced-906 | none);
      graceful (graceful-905 | none);
      link-loss (forced-906 | none);
    }
    virtual-interface-up {
      cancel-graceful (none | restart-918);
      warm (none | restart-900);
    }
  }
  context-indications {
    state-loss (forced-910 | forced-915 | none);
  }
}
wildcard-response-service-change;
```

Hierarchy Level [edit services pgcp gateway *gateway-name*]

Release Information Statement introduced in JUNOS Release 8.5.
 audit-observed-events-returns-history option introduced in JUNOS Release 9.3.
 encoding option introduced in JUNOS Release 9.3.
 service-change options introduced in JUNOS Release 9.3.

Description Configure options that affect VPG H.248 behavior.

Options The statements are explained separately.

Usage Guidelines See “Enabling a History of Media Inactivity Notifications” on page 514, “Enabling Wildcards for Service Changes” on page 514, “Disabling Bit Mirroring of DSCP Marking” on page 515, and “Configuring ServiceChange Method and Reason” on page 515.

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

h248-properties

```

Syntax h248-properties {
    application-data-inactivity-detection {
        ip-flow-stop-detection (regulated-notify | immediate-notify)
    }
    base-root {
        normal-mg-execution-time default milliseconds;
        normal-mgc-execution-time default milliseconds;
        mg-provisional-response-timer-value default milliseconds;
        mgc-provisional-response-timer-value default milliseconds;
        mg-originated-pending-limit default number-of-pendings-received;
        mgc-originated-pending-limit default number-of-pendings-received;
    }
    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;
        mgc-segmentation-timer default milliseconds;
        mg-maximum-pdu-size default bytes;
        mgc-maximum-pdu-size default bytes;
    }
    traffic-management {
        peak-data-rate {
            default bytes-per-second;
            rtcp {
                fixed-value bytes-per-second;
                percentage percentage;
            }
        }
        sustained-data-rate {
            default bytes-per-second;
            rtcp {
                fixed-value bytes-per-second;
                percentage percentage;
            }
        }
        max-burst-size {
            default bytes;
            rtcp {
                fixed-value bytes;
                percentage percentage;
            }
        }
    }
}

```

```

    inactivity-timer {
        inactivity-timeout {
            detect;
            maximum-inactivity-time {
                default 10-millisecond-units;
            }
        }
    }
}

```

Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
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.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

h248-timers

Syntax	h248-timers { initial-average-ack-delay <i>milliseconds</i> ; maximum-net-propagation-delay <i>milliseconds</i> ; maximum-waiting-delay <i>milliseconds</i> ; tmax-retransmission-delay <i>milliseconds</i> ; }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure H.248 timers for the PGCP connection.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring H.248 Timers” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

hanging-termination-detection

Syntax	hanging-termination-detection { timerx seconds; }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Enable and configure hanging termination detection.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

inactivity-timeout

Syntax	inactivity-timeout { detect; maximum-inactivity-time { default 10-millisecond-units; } }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties inactivity-timer]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configures the inactivity timeout event. The inactivity timeout event is used to detect that the inactivity timer has expired.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

inactivity-delay

Syntax	<code>inactivity-delay seconds;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> data-inactivity-detection]</code>
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the time after which the VPG begins checking for data packets on terminations that do not include a latch event.
Options	<i>seconds</i> —Time interval before checking for media inactivity. Range: 0 through 3600 Default: 5
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

inactivity-duration

Syntax	<code>inactivity-duration seconds;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> data-inactivity-detection]</code>
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the time interval that determines inactivity. When the VPG determines that the time since the last packet was received exceeds this duration, the VPG 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	<i>seconds</i> —Time during which no packets are received. Range: 5 through 86400 Default: 30
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

inactivity-timer

Syntax inactivity-timer {
 inactivity-timeout {
 detect;
 maximum-inactivity-time {
 default 10-millisecond-units;
 }
 }
 }

Hierarchy Level [edit services pgcp gateway *gateway-name* h248-properties]

Release Information Statement introduced in JUNOS Release 9.2.

Description Configures the inactivity timer package, which allows the PG to use message inactivity to detect that its active PGC has failed.

Options The statements are explained separately.

Usage Guidelines See “Configuring Default Values for H.248 Properties” on page 508.

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 VPG 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

Usage Guidelines See “Configuring H.248 Timers” on page 508.

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

interface

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	<code>[edit services pgcp virtual-interface <i>interface-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Map the virtual interface to a service interface or a physical (media) interface or both. If you do not require ingress filtering, you do not need to specify a physical interface.
Options	<i>interface-name</i> —Interface name. Include the logical portion of the name, which corresponds to the logical unit number.
Usage Guidelines	See “Configuring a Virtual Interface” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interim-ah-scheme

Syntax	<code>interim-ah-scheme { algorithm hmac-null; }</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> gateway-controller <i>gateway-controller-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Set up the packet gateway to use the interim AH scheme.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring a Connection to the Packet Gateway Controller” on page 506.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ip-flow-stop-detection

Syntax	ip-flow-stop-detection (regulated-notify immediate-notify);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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	<p>regulated-notify—Activate regulated notification of media inactivity events.</p> <p>immediate-notify—Activate non-regulated notification of media inactivity events.</p>
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

latch-deadlock-delay

Syntax	latch-deadlock-delay <i>seconds</i>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> data-inactivity-detection]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the time after which the VPG begins checking for data packets on terminations that include a latch event.
Options	<p><i>seconds</i>—Time interval before checking for data packets.</p> <p>Range: 0 through 3600</p>
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

link-loss

Syntax	link-loss (forced-906 none);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC when the virtual interface transitions to Out-of-Service because of a link loss.
Default	If you do not specify an option, the VPG includes FO/906 (forced-906).
Options	<p>forced-906—Loss of lower-layer connectivity. The virtual interface is transitioning to Out-of-Service because of a loss of layer 2 connectivity on the logical or physical interface.</p> <p>none—VPG does not send a ServiceChange command.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

max-burst-size

See the following sections:

- max-burst-size (All Streams) on page 559
- max-burst-size (RTCP Streams) on page 560

max-burst-size (All Streams)

Syntax	max-burst-size default <i>bytes</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties traffic-management]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the maximum burst size for gate streams of any protocol, including RTP.
Default	The PG uses the default value of 1000 bytes if the Policy command in H.248 messages in ON and both of the following apply: <ul style="list-style-type: none"> ■ The maximum burst size is not set in the H.248 message. ■ There is no CLI configuration for maximum burst size.
Options	<i>bytes</i> —Maximum burst size. Range: 20 through 4,294,967,295
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

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.

Description Configure the maximum burst size 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 PG 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 sustained data 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.

bytes—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: 0 through 1000

Usage Guidelines See “Configuring Default Values for H.248 Properties” on page 508.

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 <i>number-of-calls</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	<p>Configure the maximum number of concurrent calls on the VPG. If you configure multiple VPGs for one service PIC, you can use this statement to achieve a fair distribution of resources between the VPGs. For example, the MultiServices 500 PIC is capable of 10,000 concurrent calls, and you can divide this number between its associated VPGs.</p> <p>You can overbook concurrent calls to avoid resource idleness. The configured total of all VPG maximum concurrent calls can be greater than the PIC limit. For example, vpg-1 and vpg-2 are connected to single PIC. If you configure 6000 maximum concurrent calls on vpg-1 and 8000 on vpg-2, vpg-1 can open up to 6000 concurrent calls, and vpg-2 can open up to 8000 concurrent calls. However, when the total number of calls reaches 10,000, neither of the VPGs will be able to open a new context.</p> <p>If the resources on the PIC are exhausted and no more calls are allowed, the VPG sends an H.248 error message to the PGC in response to new call requests.</p>
Options	<i>number-of-calls</i> —Maximum number of concurrent calls on the VPG. Range: 0 through 10,000
Usage Guidelines	See “Configuring Virtual Packet Gateways” on page 503.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

maximum-fuf-percentage

Syntax	maximum-fuf-percentage <i>percentage</i>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> fast-update-filters]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Along with the <code>maximum-terms</code> statement, limit the number of FUF terms installed on the Packet Forwarding Engine for a VPG. This limit is the maximum value of the <code>maximum-terms</code> and <code>maximum-fuf-percentage</code> statements.
Options	<i>percentage</i> —Maximum percentage of gates with FUF filters relative to all gates currently installed for the VPG. Range: 0 through 100 Default: 10
Usage Guidelines	See “Limiting FUF Terms on VPGs” on page 510.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-inactivity-time

Syntax	<pre>maximum-inactivity-time { default 10-millisecond-units; }</pre>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties inactivity-timer inactivity-timeout]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specifies a default value for the maximum inactivity time. The default value is used if the PGC requests that the PG detect the inactivity timeout event, but the PGC does not set a value for the maximum inactivity time. If the PG does not receive a message from the PGC before the maximum inactivity time expires, it sends a Notify message to the PGC. This timer resets each time the PG receives a message from the PGC.
Options	default—Default value for the maximum inactivity time. Range: 100 through 65,535 10 millisecond units Default: 12,000
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

maximum-net-propagation-delay

Syntax	maximum-net-propagation-delay <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 <i>Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005</i> .
Options	<i>milliseconds</i> —Duration of the maximum network propagation delay time. Range: 0 through 65,535 milliseconds Default: 40,000
Usage Guidelines	See “Configuring H.248 Timers” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-synchronization-mismatches

Syntax	maximum-synchronization-mismatches <i>number-of-mismatches</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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. If the number of mismatches exceeds this number, the pgcpd process clears the state of the PIC and the state of the pgcpd process.
Options	<i>number-of-mismatches</i> —Maximum number of mismatches allowed during the synchronization procedure with the PIC. Range: 0 through 3000
Usage Guidelines	See “Configuring Graceful Restart” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-synchronization-time

Syntax	maximum-synchronization-time <i>seconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> graceful-restart]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Configure the time within which the pgcpd process and the PIC should complete the synchronization process. If the synchronization process is not complete when this time expires, the pgcpd process clears the state of the PIC and the state of the pgcpd process.
Options	<i>seconds</i> —Maximum time allowed for the synchronization procedure with the PIC. Range: 0 through 300
Usage Guidelines	See “Configuring Graceful Restart” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-terms

Syntax	maximum-terms <i>number-of-terms</i>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> fast-update-filters]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Along with the <i>maximum-fuf-percentage</i> statement, limit the number of FUF terms installed on the Packet Forwarding Engine for a VPG. This limit is the maximum value of the <i>maximum-terms</i> and <i>maximum-fuf-percentage</i> statements.
Options	<i>number-of-terms</i> —Maximum number of FUF terms installed for the VPG. Range: 0 through 20000 Default: 20000
Usage Guidelines	See “Limiting FUF Terms on VPGs” on page 510.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-waiting-delay

Syntax	maximum-waiting-delay <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-timers]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Define a maximum waiting delay (MWD) value. When the VPG loses its connection to a PGC, it attempts to reconnect to the PGC. If the VPG cannot reconnect to the PGC, it traverses its list of PGCs and attempts to connect to one of the PGCs. If the VPG finishes traversing its list of PGCs, and has not connected to a PGC, the VPG waits for a random value between 0 and MWD milliseconds before it begins another attempt to connect to a PGC. See section 9.2 of <i>Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005</i> .
Options	<i>milliseconds</i> —Maximum time the VPG waits before contacting a new PGC when the connection to the controlling PGC is lost. Range: 1 through 36,000 milliseconds Default: 3000
Usage Guidelines	See “Configuring H.248 Timers” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

media

Syntax	media { rtcp; rtp; }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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.
Usage Guidelines	See “Enabling Application Layer Gateways for RTP and RTCP for Media Flows” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

media-service

Syntax	<code>media-service <i>media-service-name</i> { nat-pool <i>nat-pool-name</i>; }</code>
Hierarchy Level	<code>[edit services pgcp media-service <i>media-service-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure media services for the packet gateway configuration. Media services are applied to PGCP packets.
Options	<p><i>media-service-name</i>—Identifier for the media service name.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring a Media Service” on page 507.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

mg-maximum-pdu-size

Syntax	<code>mg-maximum-pdu-size default <i>bytes</i>;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> h248-properties segmentation]</code>
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MG maximum PDU size property of the segmentation package.
Options	<p><i>bytes</i>—Maximum size of messages that the PGC sends to the packet gateway.</p> <p>Range: 512 through 65,507</p> <p>Default: 1472</p>
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

mg-originated-pending-limit

Syntax	mg-originated-pending-limit default <i>number-of-messages</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties base-root]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MG originated pending limit property of the base root package.
Options	<i>number-of-messages</i> —Number of transaction pending messages that the PGC can receive from the VPG. Range: 1 through 30 Default: 4
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mg-provisional-response-timer-value

Syntax	mg-provisional-response-timer-value default <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties base-root]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MG provisional response timer property of the base root package.
Options	<i>milliseconds</i> —Time within which the PGC should expect a pending response from the VPG if a transaction cannot be completed. Range: 500 through 4,000 Default: 2000
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mg-segmentation-timer

Syntax	mg-segmentation-timer default <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties segmentation]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MG segmentation timer value property of the segmentation package.
Options	<i>milliseconds</i> —Time within which the PGC should expect to receive outstanding message segments from the VPG after it receives the SegmentationCompleteToken. Range: 500 through 10,000 Default: 4000
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mgc-maximum-pdu-size

Syntax	mgc-maximum-pdu-size default <i>bytes</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties segmentation]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MGC maximum PDU size property of the segmentation package.
Options	<i>bytes</i> —Maximum size of messages that the VPG sends to the PGC. Range: 512 through 65,507 Default: 1472
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mgc-originated-pending-limit

Syntax	mgc-originated-pending-limit default <i>number-of-messages</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties base-root]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MGC originated pending limit property of the base root package.
Options	<i>number-of-messages</i> —Number of transaction pending messages that the VPG can receive from the PGC. Range: 1 through 30 Default: 4
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mgc-provisional-response-timer-value

Syntax	mgc-provisional-response-timer-value default <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties base-root]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MGC provisional response timer value property of the base root package.
Options	<i>milliseconds</i> —Time within which the VPG should expect a pending response from the PGC if a transaction cannot be completed. Range: 500 through 4,000 Default: 2000
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mgc-segmentation-timer

Syntax	<code>mgc-segmentation-timer default <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> h248-properties segmentation]</code>
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the MGC segmentation timer value property of the segmentation package.
Options	<p><i>milliseconds</i>—Time within which the VPG should expect to receive outstanding message segments from the PGC after it receives the SegmentationCompleteToken.</p> <p>Range: 500 through 10,000</p> <p>Default: 4000</p>
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

monitor

Syntax	<pre>monitor { media { rtcp; rtp; } }</pre>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i>]</code>
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.
Usage Guidelines	See “Enabling Application Layer Gateways for RTP and RTCP for Media Flows” on page 511.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

nat-pool

Syntax	<code>nat-pool <i>nat-pool-name</i>;</code>
Hierarchy Level	<code>[edit services pgcp media-service <i>media-service-name</i>]</code>
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	<i>nat-pool-name</i> —Identifier for the NAT address pool.
Usage Guidelines	See “Configuring a Media Service” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

network-operator-id

Syntax	<code>network-operator-id <i>network-operator-id</i>;</code>
Hierarchy Level	<code>[edit services pgcp session-mirroring delivery-function <i>deliver-function-name</i>] [edit services pgcp gateway <i>gateway-name</i> session-mirroring delivery-function <i>deliver-function-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the network operator ID. The PG 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	<i>network-operator-id</i> —The network operator ID can be up to five characters.
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
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	<i>JUNOS Multiplay Solutions Guide</i>

normal-mg-execution-time

Syntax	normal-mg-execution-time default <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties base-root]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the normal MG execution time property of the base root package.
Options	<i>milliseconds</i> —Interval within which the PGC expects a response to transactions from the VPG. Range: 500 through 4,000 Default: 500
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

normal-mgc-execution-time

Syntax	normal-mgc-execution-time default <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties base-root]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set a default value for the normal MGC execution time property of the base root package.
Options	<i>milliseconds</i> —Interval within which the VPG expects a response to transactions from the PGC. Range: 500 through 4,000 Default: 500
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

notification-behavior

Syntax	notification-behavior { notification-regulation default (once 0 – 100); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the default frequency for regulated media inactivity notifications sent by the PG.
Options	The statement is explained separately.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

notification-rate-limit

Syntax	notification-rate-limit <i>rate</i> ;
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.
Options	<i>rate</i> —Maximum number of notifications per second the PIC sends to a PGC. Range: 10 through 10,000
Usage Guidelines	See “Limiting the Rate of Messages Sent from the PIC” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

notification-regulation

Syntax	notification-regulation (once 0 – 100);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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	<p>once—Send only one media inactivity notification for a regulated event to the PGC.</p> <p>0 – 100—The percentage of media inactivity notifications for regulated events to send to the PGC.</p>
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

no-dscp-bit-mirroring

Syntax	no-dscp-bit-mirroring;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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.
Usage Guidelines	See “Disabling Bit Mirroring of DSCP Marking” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

overload-control

Syntax	overload-control { queue-limit-percentage <i>percentage</i> ; }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the PG to send overload messages to the PGC based on the status of its work queue. The overload messages cause the PGC to lower the rate at which it admits packets for processing.
Options	The statement is described separately.
Usage Guidelines	See “Configuring Overload Control” on page 514.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

peak-data-rate

See the following sections:

- [peak-data-rate \(All Streams\)](#) on page 576
- [peak-data-rate \(RTCP Streams\)](#) on page 577

peak-data-rate (All Streams)

Syntax	<code>peak-data-rate default <i>bytes-per-second</i>;</code>
Hierarchy Level	[edit services services pgcp gateway <i>gateway-name</i> h248-properties traffic-management]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the peak data rate for gate streams of any protocol.
Default	The PG 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: <ul style="list-style-type: none"> ■ The peak data rate is not set in the H.248 message. ■ There is no CLI configuration for peak data rate.
Options	<i>bytes-per-second</i> —Peak data rate. Range: 125 through 4,294,967,295
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

peak-data-rate (RTCP Streams)

Syntax `peak-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 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 PG 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 bytes 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 rate.

percentage—Peak data rate as a percentage of the RTP gate's rate.

Range: 0 through 1000

Usage Guidelines See “Configuring Default Values for H.248 Properties” on page 508.

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*

queue-limit-percentage

Syntax	queue-limit-percentage <i>percentage</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 work queue utilization reaches this limit percentage, overload control is activated. An overload notification is generated for each transaction on a gate that contains an ADD. When 100 percent of the queue is in use, transactions are dropped with error 510 (insufficient resources).
Options	<p><i>percentage</i>—Percentage of the overload control work queue in use that triggers creation of an overload event.</p> <p>Range: 55 through 95</p> <p>Default: 80</p>
Usage Guidelines	See “Configuring Overload Control” on page 514.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

reconnect

Syntax	reconnect (disconnected-900 restart-902);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Registration Request ServiceChange commands when it attempts to reregister with the PGC or register with a new PGC after the control association is disconnected.
Default	If you do not specify an option, the VPG includes DC/900 (disconnected-900).
Options	<p>disconnected-900—Service restored. The VPG is registering with the last controlling PGC following a disconnection of the VPG and PGC.</p> <p>restart-902—Warm boot. The VPG is transitioning to In-Service, and the previously installed state is retained.</p>
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide.</i>

report-service-change

Syntax	report-service-change { service-change-type (forced-906 forced-910); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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.
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511,
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

request-timestamp

Syntax	request-timestamp (requested suppressed autonomous);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties event-timestamp]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specific whether timestamp information is made available to the PGC or is suppressed.
Options	requested—Enables PGC access to timestamp information for notifications. suppressed—Disables PGC access to timestamp information for notifications. autonomous—Equivalent to suppressed.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

routing-instance

Syntax	<code>routing-instance <i>instance-name</i> { service-interface <i>name.number</i>; }</code>
Hierarchy Level	<code>[edit services pgcp virtual-interface <i>interface-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4. service-interface option added 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	<i>instance-name</i> —Name of a routing instance that has been configured at the <code>[edit routing-instance]</code> hierarchy level. The remainder of the statements are explained separately.
Usage Guidelines	See “Configuring a Virtual Interface” on page 507.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rtcp

Syntax	<code>rtcp;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> monitor media]</code>
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 PGC installs media gates on the VPG.
Usage Guidelines	See “Enabling Application Layer Gateways for RTP and RTCP for Media Flows” on page 511
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rtp

Syntax	<code>rtp;</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> monitor media]</code>
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 packet gateway controller (PGC) installs media gates on the virtual packet gateway (VPG).
Usage Guidelines	See “Enabling Application Layer Gateways for RTP and RTCP for Media Flows” on page 511.
Required Privilege Level	interface-level—To view this statement in the configuration. interface-level—To add this statement to the configuration.

rule

Syntax	<pre>rule <i>rule-name</i> { gateway <i>gateway-name</i>; media-service <i>media-service-name</i>; }</pre>
Hierarchy Level	<code>[edit services pgcp],</code> <code>[edit services service-set <i>service-set-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Specify the rule that the router uses when it applies the media service.
Options	<p><i>rule-name</i>—Identifier for the rule.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring the Packet Gateway Rule Content” on page 513.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rule-set

Syntax	<code>rule-set <i>rule-set-name</i> { [rule <i>rule-name</i>] }</code>
Hierarchy Level	<code>[edit services pgcp], [edit services service-set <i>service-set-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Specify the rule set the router uses when applying this service.
Options	<i>rule-set-name</i> —Identifier for the collection of rules that make up this rule set.
Usage Guidelines	See “Configuring the Packet Gateway Rule Set” on page 513.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

segmentation

Syntax	<code>segmentation { mg-segmentation-timer default <i>milliseconds</i>; mgc-segmentation-timer default <i>milliseconds</i>; mg-maximum-pdu-size default <i>bytes</i>; mgc-maximum-pdu-size default <i>bytes</i>; }</code>
Hierarchy Level	<code>[edit services pgcp gateway <i>gateway-name</i> h248-properties]</code>
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Configure default values for properties in the segmentation package defined in Annex E of <i>Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005</i> .
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

send-notification-on-delay

Syntax	send-notification-on-delay;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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.
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
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

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);
                failure (forced-904 | forced-906 | none);
                graceful (graceful-905 | none);
                link-loss (forced-906 | 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 VPG includes in ServiceChange commands that it sends to the PGC when the state of a control association, virtual interface, or context changes.

Options The options are explained separately.

Usage Guidelines See “Configuring ServiceChange Method and Reason” on page 515.

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 <i>gateway-name</i> 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	<p>forced-906—Service is terminated using a forced termination method with reason code 906 (loss of lower layer connectivity).</p> <p>forced-910—Service is terminated using a forced termination with reason code 910 (media capability failure).</p>
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

service-interface

Syntax	service-interface <i>name.number</i> ;
Hierarchy Level	[edit services pgcp virtual-interface <i>virtual-interface-name</i> 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	<i>name.number</i> —Name and logical interface number of the service interface.
Usage Guidelines	See “Configuring a Virtual Interface” on page 507.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

service-state

See the following sections:

- [service-state \(VPG\)](#) on page 586
- [service-state \(Virtual Interface\)](#) on page 587

service-state (VPG)

Syntax	service-state (in-service out-of-service-forced out-of-service-graceful);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i>]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Set the service state of the virtual packet gateway (VPG).
Options	<p>in-service—The VPG is operational and available for traffic. When the VPG is in service, it attempts to connect to the packet gateway controller (PGC) and accepts all Packet Gateway Control Protocol (PGCP) commands from the PGC.</p> <p>out-of-service-forced—Force the VPG out of service. When the VPG is forced out of service, it immediately removes all gates and disconnects from the PGC. The VPG does not attempt to establish a new connection.</p> <p>out-of-service-graceful—Cause the VPG 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 packet gateway accepts only subtract commands from the PGC.</p>
Usage Guidelines	See “Configuring the Service State of a VPG” on page 512.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

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	<p>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.</p> <p>out-of-service-forced—Forces 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.</p> <p>out-of-service-graceful—Causes 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 PGC.</p>
Usage Guidelines	See “Configuring the Service State of a Virtual Interface” on page 512.
Required Privilege Level	<p>interface-level—To view this statement in the configuration.</p> <p>interface-level—To add this statement to the configuration.</p>

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.
Usage Guidelines	See “Configuring Virtual Packet Gateways” on page 503.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

session-mirroring

Syntax	<pre> session-mirroring { delivery-function <i>delivery-function-name</i> { destination-address <i>destination-address</i>; destination-port <i>destination-port</i>; network-operator-id <i>network-operator-id</i>; source-address <i>source-address</i>; source-port <i>source-port</i>; } disable-session-mirroring; } </pre>
Hierarchy Level	<pre> [edit services pgcp] [edit services pgcp gateway <i>gateway-name</i>] </pre>
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the session mirroring feature.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
Required Privilege Level	pgcp-session-mirroring—To view this statement in the configuration. pgcp-session-mirroring-control—To add this statement to the configuration.

source-address

Syntax	<pre> source-address <i>source-address</i>; </pre>
Hierarchy Level	<pre> [edit services pgcp session-mirroring delivery-function <i>deliver-function-name</i>] [edit services pgcp gateway <i>gateway-name</i> session-mirroring delivery-function <i>deliver-function-name</i>] </pre>
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the source address that is applied to mirrored packets.
Options	<i>destination-address</i> —Address of the interface on which the PG sends session-mirroring data to the deliver function server.
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
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	<i>JUNOS Multiplay Solutions Guide</i>

source-port

Syntax	source-port <i>source-port</i> ;
Hierarchy Level	[edit services pgcp session-mirroring delivery-function <i>deliver-function-name</i>] [edit services pgcp gateway <i>gateway-name</i> session-mirroring delivery-function <i>deliver-function-name</i>]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the source port applied to the mirrored packets.
Options	<i>destination-port</i> —Port on which the PG sends session-mirroring data to the delivery function server. Range: 1 through 65,535
Usage Guidelines	See “Configuring Session Mirroring” on page 514.
Required Privilege Level	pgcp-session-mirroring—To view this statement in the configuration. pgcp-session-mirroring-control—To add this statement to the configuration.

state-loss

Syntax	state-loss (forced-910 forced-915 none);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change context-indications]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in Service-Interruption ServiceChange commands that it sends to the PGC after a state loss on a specific context.
Default	If you do not specify an option, the VPG 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 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 states was detected on one or more of the context’s gates. none—VPG does not send a ServiceChange command to the PGC.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

stop-detection-on-drop

Syntax	stop-detection-on-drop;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> data-inactivity-detection]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure the PG to stop inactivity detection when a gate action is set to drop. When the call is resumed, the PG starts the delay time and resumes data inactivity detection.
Usage Guidelines	See “Detecting Latch Deadlocks and Inactivity Delays” on page 511.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

sustained-data-rate

See the following sections:

- `sustained-data-rate` (All Streams) on page 591
- `sustained-data-rate` (RTCP Streams) on page 592

sustained-data-rate (All Streams)

Syntax	<code>sustained-data-rate default <i>bytes-per-second</i>;</code>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties traffic-management]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure the sustained data rate for streams of any protocol, including RTP.
Default	The PG 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: <ul style="list-style-type: none"> ■ The sustained data rate is not set in the H.248 message. ■ There is no CLI configuration for sustained data rate.
Options	<i>bytes-per-second</i> —Sustained data rate. Range: 125 through 4,294,967,295
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i>

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 PG 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 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 bytes per second.

bytes-per-second—Sustained data rate.

Range: 125 through 4,294,967,295

percentage—Value entered is a percentage of the RTP gate's rate.

percentage—Sustained data rate as a percentage of the RTP gate's rate.

Range: 0 through 1000

Usage Guidelines See “Configuring Default Values for H.248 Properties” on page 508.

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 <i>seconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-properties hanging-termination-detection]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	<p>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 PG and the PGC for a termination before this time expires, the PG sends a notification to the PGC. The timer resets when the PG and the PGC 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 PGC.</p> <p>Your configuration takes effect on new and modified terminations.</p>
Options	<p><i>seconds</i>—Number of seconds between the last message exchanged for this termination and when the PG sends a notification to the PGC. Setting the timer to zero (0) deactivates hanging termination detection.</p> <p>Range: 0 through 4,294,967,295</p>
Usage Guidelines	See “Configuring Default Values for H.248 Properties” on page 508.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

traffic-management

Syntax

```

traffic-management {
  peak-data-rate {
    default bytes-per-second;
    rtcp {
      fixed-value bytes-per-second;
      percentage percentage;
    }
  }
  sustained-data-rate {
    default bytes-per-second;
    rtcp {
      fixed-value bytes-per-second;
      percentage percentage;
    }
  }
  max-burst-size {
    default bytes;
    rtcp {
      fixed-value bytes;
      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.

Usage Guidelines See “Configuring Default Values for H.248 Properties” on page 508.

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 <i>milliseconds</i> ;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 <i>Gateway control protocol v3, ITU-T Recommendation H.248.1, September 2005</i> .
Options	<i>milliseconds</i> —Duration of the delay before the packet gateway considers the PGC to be down. Range: 0 through 65,535 Default: 25000
Usage Guidelines	See “Configuring H.248 Timers” on page 508.
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*> <world-readable | no-world-readable> <match
 regex>;
 flag *flag*;
 }

Hierarchy Level [edit services pgcp]

Release Information Statement introduced in JUNOS Release 8.4.
 gate, pgcp-stack-debug, pgcp-stack-h248, pgcp-stack-trace, and trace introduced in JUNOS Release 8.5.

Description Configure PGCP tracing operations. The messages are output to `/var/log/pgcpd`.

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 1,000 files

Default: 3 files

flag *flag*—Tracing operation to perform:

- all—Trace everything.
- configuration—Trace configuration events.
- debug—Trace debug messages.
- error—Trace error messages.
- gate—Trace gate request and reply events.
- media-function—Trace media function events.
- pgc-connection—Trace PGC connection events.
- pgcp-stack-debug—Trace PGCP stack debug events.
- pgcp-stack-h248—Trace PGCP stack H.248 events.
- pgcp-stack-trace—Trace PGCP stack function call events.
- routing-socket—Trace routing-socket events.
- trace—Trace process function calls.

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: xk to specify KB, xm to specify MB, or xg 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 PGCP Operations” on page 516.

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

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 VPG includes in Notification Messages or Registration commands in ServiceChange commands when a control association transitions to In-Service.

Options The options are explained separately.

Usage Guidelines See “Configuring ServiceChange Method and Reason” on page 515.

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 name {
 media-service media-service-name;
 interface interface-identifier;
 routing-instance instance-name {
 service-interface name.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 added in JUNOS Release 9.0.
service-interface option added in JUNOS Release 9.3.

Description Configure a virtual interface for the packet gateway.

Options *interface-number*—Identifier for the interface.

The remainder of the statements are explained separately.

Usage Guidelines See “Configuring a Virtual Interface” on page 507.

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

virtual-interface-down

Syntax	virtual-interface-down { administrative (forced-905 forced-906 none); failure (forced-904 forced-906 none); graceful (graceful-905 none); link-loss (forced-906 none); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change virtual-interface-indications]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in ServiceChange commands that it sends to the PGC when the state of the virtual interface changes to Out-of-Service.
Options	The options are explained separately.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

virtual-interface-indications

Syntax	<pre>virtual-interface-indications { virtual-interface-down { administrative (forced-905 forced-906 none); failure (forced-904 forced-906 none); graceful (graceful-905 none); link-loss (forced-906 none); } virtual-interface-up { cancel-graceful (none restart-918); warm(none restart-900); } }</pre>
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the method and reason that the VPG includes in ServiceChange commands that it sends to the PGC when the state of the virtual interface changes.
Options	The options are explained separately.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

virtual-interface-up

Syntax	virtual-interface-up { cancel-graceful (none restart-918); warm (none restart-900); }
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options service-change virtual-interface-indications]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the ServiceChange command that the VPG sends to the PGC when the state of the virtual interface changes to In-Service.
Options	The options are explained separately.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

warm

Syntax	warm (none restart-900);
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> 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 VPG includes in Service-Restoration ServiceChange commands that it sends to the PGC when a virtual interface transitions to In-Service.
Default	If you do not specify an option, the VPG includes RS/900 (restart-900).
Options	none—VPG does not send a ServiceChange command. restart-900—Service restored. The virtual interface has become In-Service and is in the Forwarding state.
Usage Guidelines	See “Configuring ServiceChange Method and Reason” on page 515.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

wildcard-response-service-change

Syntax	wildcard-response-service-change;
Hierarchy Level	[edit services pgcp gateway <i>gateway-name</i> h248-options]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Enable the VPG to issue service change commands as wildcard-response commands, which trigger a short response from the PGC. If you do not enable the use of wildcard responses for service change commands, the PGC will generate an individual response for every termination that matches the service change command.
Usage Guidelines	See “Enabling Wildcards for Service Changes” on page 514.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

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 name.number;
  }
}
```

This chapter discusses the following topics that provide information about configuring service interface pools:

- [Configuring Service Interface Pools on page 603](#)

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 name.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	<code>interface <i>name.number</i>;</code>
Hierarchy Level	<code>[edit services service-interface-pools pool <i>pool-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Add logical service interfaces to the pool of service interfaces.
Options	<p><i>name.number</i>—Name and logical unit number of the service interface.</p> <ul style="list-style-type: none">■ All interfaces in a pool must belong to the same service PIC.■ 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.
Usage Guidelines	See “Configuring Service Interface Pools” on page 603.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

pool

Syntax	<code>pool <i>pool-name</i> { interface <i>name.number</i>; }</code>
Hierarchy Level	[edit services service-interface-pools]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Configure a service interface pool for PGCP VPN aggregation.
Options	<i>pool-name</i> —Name of the service interface pool. The remaining options are explained separately.
Usage Guidelines	See “Configuring Service Interface Pools” on page 603.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

service-interface-pools

Syntax	<code>service-interface-pools { pool <i>pool-name</i> { interface <i>name.number</i>; } }</code>
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.
Usage Guidelines	See “Configuring Service Interface Pools” on page 603.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Multiplay Solutions Guide</i> .

Part 3

Data Link Switching

- Data Link Switching Overview on page 609
- Data Link Switching Configuration Guidelines on page 611
- Summary of Data Link Switching Configuration Statements on page 623

Chapter 30

Data Link Switching Overview

The data link switching (DLSw) protocol allows you to tunnel System Network Architecture (SNA) Logical Link Control (LLC2) traffic over an Internet Protocol (IP) network.



NOTE: DLSw configuration is supported for the J-series Services Router only.

This chapter includes the following:

- Overview on page 609
- DLSw Standards on page 609

Overview

DLSw enables you to allow SNA clients to communicate with SNA applications on a mainframe through an IP network. The IP network between an SNA client and an SNA application becomes transparent with DLSw. DLSw is configured on peer IP routers.

A DLSw connection is a Transport Control Protocol (TCP) connection that carries control and information frames between two DLSw peers within the IP network. After a connection is established, a DLSw circuit can be created for transporting SNA traffic.

DLSw Standards

DLSw is defined in the following documents:

- RFC 1795, *Data Link Switching: Switch-to-Switch Protocol AIW DLSw RIG: DLSw Closed Pages, DLSw Standard Version 1*
- RFC 2166, *APPN Implementer's Workshop Closed Pages Document DLSw v2.0 Enhancements*

Chapter 31

Data Link Switching Configuration Guidelines

The data link switching (DLSw) protocol allows you to tunnel System Network Architecture (SNA) Logical Link Control (LLC2) traffic over an Internet Protocol (IP) network.



NOTE: DLSw configuration is supported on J-series Services Routers only.

To configure DLSw properties on a J-series Services Router, you include the following statements at the [edit protocols] hierarchy level:

```
dls {
  connection-idle-timeout time;
  explorer-wait-time seconds;
  load-balance circuit-weight;
  local-peer peer-address;
  multicast-address address;
  promiscuous;
  reachability-cache-timeout seconds;
  receive-initial-pacing number;
  remote-peer peer-address {
    circuit-weight weight;
    cost cost;
  }
  dls-cos {
    destination-interface interface-name;
    type-of-service number;
  }
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp> <(world-readable |
      no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
}
```

This chapter discusses the following topics that provide information about configuring DLSw:

- Configuring DLSw Properties on page 612
- Configuring Logical Link Control on an Interface on page 617
- Configuring DLSw Ethernet Redundancy Using LLC2 Properties on page 618

Configuring DLSw Properties

This section includes the following sections:

- Minimum DLSw Configuration on page 612
- Configuring the Remote Peer on page 612
- Configuring Load Balancing on page 613
- Configuring DLSw Timers on page 613
- Configuring the Local Peer on page 614
- Configuring the Initial Pacing Window on page 615
- Configuring the Idle Timeout on page 615
- Configuring the Multicast Address on page 615
- Configuring Class of Service on page 615
- Tracing DLSw Protocol Traffic on page 617

Minimum DLSw Configuration

To enable DLSw on a router, you must include at least the following statements at the [edit protocols dlsw] hierarchy level:

```
local-peer peer-address;  
promiscuous;
```

Specify the **promiscuous** statement to allow the router to accept all incoming peer connections.

Configuring the Remote Peer

You can specify one or more remote peers to initiate a DLSw connection. Configuring the remote peer initiates the connection and also accepts connection from that remote peer only, if promiscuous mode is not configured.

To configure a specific DLSw remote peer, include the **remote-peer** statement at the [edit protocols dlsw] hierarchy level:

```
remote-peer peer-address;
```

You can configure more than one remote peer.

Configuring Load Balancing

When more than one remote peer provides alternative paths to a destination, you can specify preferences among the available routers or enable load balancing for alternatives that have the same lowest cost value. The DLSw router maintains a cache of paired media access control (MAC) address and IP address entries to determine whether an SNA host can be reached by means of any of the peers it knows about. There are separate caches for local and remote peers, called the *local reachability cache* and the *remote reachability cache*.

To configure DLSw load balancing, include the **load-balance** statement at the **[edit protocols dlsw]** hierarchy level:

```
load-balance circuit-weight;
```

By default, load balancing is disabled. Load balancing uses the circuit-weight algorithm to make its determinations. This requires that, if there are multiple routes to a destination that share the same lowest cost value, the number of circuits traversing each peer is balanced according to the circuit weight you configure for each peer:

- To configure the *circuit cost*, or how much preference is attached to establishing a circuit using a peer, include the **cost** statement at the **[edit protocols dlsw remote-peer peer-address]** hierarchy level:

```
cost cost;
```

By default, the **cost** value is 100. The range is from 0 through 127. The lower the configured **cost** value, the higher the preference.

- To configure the *circuit weight*, or the extent to which a peer should participate in establishing circuits, include the **circuit-weight** statement at the **[edit protocols dlsw remote-peer peer-address]** hierarchy level:

```
circuit-weight weight;
```

By default, the **circuit-weight** value is 1. The range is from 1 through 127. The higher the configured **circuit-weight** value, the greater the percentage of total circuits established with this remote peer.

Configuring DLSw Timers

You can also configure two timer values that influence how circuits on the DLSw router are established:

- *Explorer timeout* is the maximum time the DLSw router waits for its peers to respond to its explorer requests. To configure a timeout period, include the **explorer-wait-time** statement at the **[edit protocols dlsw]** hierarchy level:

```
explorer-wait-time seconds;
```

By default, the **explorer-wait-time** value is 10 seconds. The range is from 5 through 60 seconds.

- *Reachability cache timeout* is the maximum time an entry remains in the reachability cache before it is deleted. To configure a timeout period, include the `reachability-cache-timeout` statement at the `[edit protocols dls]` hierarchy level:

```
reachability-cache-timeout seconds;
```

By default, the `reachability-cache-timeout` value is 900 seconds. The range is from 0 through 3600 seconds. A lower value helps detect new routes that might have become available during the timeout period.

Configuring the Local Peer

You can specify a local peer from which you will accept connection requests specifically. To configure a DLSw local peer, include the `local-peer` statement at the `[edit protocols dls]` hierarchy level:

```
local-peer peer-address;
```

Examples: Configuring DLSw Peers

Configure the router when it is closer to the client and expects to initiate connections to peers:

```
dls {
  local-peer peer-address;
  remote-peer peer-address;
}
```

Configure the router when it is closer to the server and does not expect to initiate connections:

```
dls {
  local-peer peer-address;
  promiscuous;
}
```

Configure specific peers and accept incoming connections only from those peers:

```
dls {
  local-peer peer-address;
  promiscuous;
}
```

Configure so that for some applications you initiate connections, and for other applications you receive connections from other peers:

```
dls {
  local-peer peer-address;
  remote-peer peer-address;
  promiscuous;
}
```

Configuring the Initial Pacing Window

To configure the DLSw initial pacing window size, include the `receive-initial-pacing` statement at the `[edit protocols dlsw]` hierarchy level:

```
receive-initial-pacing number;
```

Specify the number value to configure the initial value of the receive pacing window size for the sending peer.

Configuring the Idle Timeout

To configure an idle timeout period, include the `connection-idle-timeout` statement at the `[edit protocols dlsw]` hierarchy level:

```
connection-idle-timeout time;
```

Specify the `connection-idle-timeout` statement to configure the idle time period after which the connection is torn down.

Configuring the Multicast Address

To configure the multicast address for DLSw, include the `multicast-address` statement at the `[edit protocols dlsw]` hierarchy level:

```
multicast-address address ;
```

Configuring Class of Service

Class of service (CoS) can be used to classify DLSw packets, which are sent to a logical tunnel interface that loops the packets into the Packet Forwarding Engine. The Packet Forwarding Engine classifies and queues the packets based on the configured type-of-service (ToS) value.

To configure CoS for DLSw, include the `dlsw-cos` statement at the `[edit protocols dlsw]` hierarchy level:

```
dlsw-cos {
  destination-interface interface-name;
  type-of-service number;
}
```

To enable traffic through a specified destination interface, specify the `destination-interface` statement. To configure the ToS value, specify the `type-of-service` statement.

To configure the Differentiated Services code point (DSCP) for DLSw, include the `dscp` statement at the `[edit class-of-service interfaces interface-name unit unit-number classifiers]` hierarchy level:

```
dscp {
  default;
}
```

For information on configuring class-of-service properties, see the *JUNOS Class of Service Configuration Guide*.

Example: Configuring CoS for a DLSw Connection

Configure the lt-0/0/0 interface:

```
[edit]
interfaces lt-0/0/0 {
  unit 0 {
    encapsulation frame-relay;
    dlci 100;
    peer-unit 1;
    family inet;
  }
  unit 1 {
    encapsulation frame-relay;
    dlci 100;
    peer-unit 0;
    family inet;
  }
}
```

Enable the DLSw protocol:

```
[edit protocols]
dls {
  local-peer 1.1.1.1;
  remote-peer 1.1.1.2;
  dls-cos {
    type-of-service 192;
    destination-interface lt-0/0/0.0;
  }
}
```

Configure the CoS parameters on the lt-0/0/0 interface:

```
[edit class-of-service]
classifiers {
  inet-precedence ipPREC {
    forwarding-class best-effort {
      loss-priority low code-points 000;
      loss-priority high code-points 001;
    }
    forwarding-class expedited-forwarding {
      loss-priority low code-points 010;
      loss-priority high code-points 011;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-points 100;
      loss-priority high code-points 101;
    }
    forwarding-class network-control {
      loss-priority low code-points 110;
      loss-priority high code-points 111;
    }
  }
}
```

```

    }
  }
}
interfaces {
  lt-0/0/0 {
    unit 1 {
      classifiers {
        dscp default;
      }
    }
  }
}
}

```

Tracing DLSw Protocol Traffic

To debug the DLSw protocol or trace DLSw protocol traffic, you can specify DLSw-specific options by including the `traceoptions` statement at the `[edit protocols dlsw]` hierarchy level:

```

traceoptions{
  file name <replace> <size size> <files number> <no-stamp> <(world-readable |
    no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}

```

You can specify the following DLSw-specific trace options in the `flag` statement:

- `all`—Everything
- `memory`—Memory traces
- `packets`—Packet traces
- `parse`—Configuration parse traces
- `route-socket`—Route socket traces

Configuring Logical Link Control on an Interface

You must configure logical link control (LLC) on an interface to enable a DLSw connection. To configure LLC on an interface, include the following statements on the DLSw interface:

```

interfaces interface-name {
  unit logical-unit-number {
    family llc2 {
      ack-max count;
      ack-delay-time time;
      idle-time time;
      local-window count;
      max-retry number;
      p-bit-timeout time;
      t1-time time;
      t2-timeout time;
      trej-time time;
    }
  }
}

```

```

    }
  }
}

```

For information on configuring LLC and interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Example: Configuring LLC Options on an Interface

Configure LLC options on an unnumbered interface:

```

[edit]
interfaces {
  fe-0/0/0 {
    unit 0 {
      family inet;
      address 10.10.10.2/24;
    }
    family llc2 {
      ack-delay-time 3000;
      ack-max 10;
      idle-time 102;
      local-window 15;
      max-retry 20;
      p-bit-timeout 100;
      t1-time 101;
      t2-time 101;
      max-retry 5;
      trej-time 4000;
    }
  }
}

```

Configuring DLSw Ethernet Redundancy Using LLC2 Properties

To achieve fault tolerance and load sharing, two or more redundant DLSw routers can be deployed on the same LAN segment using Ethernet redundancy support. These redundant routers provide alternate paths to the destinations and avoid a single point of failure.

When DLSw Ethernet redundancy is configured on a LAN segment, a master router is selected from a group of DLSw neighbors. The master router establishes the circuits.

To configure DLSw Ethernet redundancy, include the **redundancy-group** statement and define redundancy group options:

```

llc2 {
  redundancy-group group-number {
    advertise-interval seconds;
    map {
      local-mac mac-address remote-mac mac-address;
      preempt hold-time seconds;
      no-preempt;
      priority priority;
    }
  }
}

```

```

track {
  dls {
    destination mac-address priority-cost priority;
    peer ip-address priority-cost priority;
  }
  interface interface-name priority-cost priority;
}
}
}
}
}

```

You can include these statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family]

You can configure the following redundancy options:

- *group-number*—The group to which this router belongs. Specify the group number, in the range from 0 through 255.
- *advertise-interval*—The advertisement interval of DLSw peers on the network. All routers in the redundancy group must use the same advertisement interval. Specify the number of seconds, from 1 through 255. The default is 1 second.
- *map*—Map a local peer MAC address to a remote peer MAC address.
 - *local-mac*—The local peer MAC address to be mapped to a remote peer MAC address.
 - *mac-address*—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn.nn.nn.nn.nn.nn*. For example, 0011.2233.4455 or 00:11:22:33:44:55.
 - *remote-mac*—The remote destination MAC address to be mapped to a local peer MAC address.
- *preempt hold-time seconds*—Configure the time to wait before a higher-priority backup router can preempt the master router. Specify the number of seconds, from 0 through 3600. The default is not to implement DLSw preemption.
- *no-preempt*—Prohibit the preemption of the master router.
- *priority priority*—The router's priority for becoming the master router. The router with the highest priority within the redundancy group becomes the master. A larger value indicates a higher priority for being elected. Specify the priority from 1 through 255. The default is 100 (for backup routers).
- *track*—Enable the following tracking options for the remote peer and the destination peer:
 - *dls*—DLSw protocol.
 - *destination mac-address priority-cost priority*—The local MAC address and the priority. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn.nn.nn.nn.nn.nn*. For


example, 0011.2233.4455 or 00:11:22:33:44:55. The priority cost is the value subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.

- **peer ip-address priority-cost priority**—IP address of the remote peer. The priority cost is the value subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
- **interface interface-name**—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.

Example: Configuring DLSw Ethernet Redundancy

In Figure 4 on page 620, the local hosts share the same destination MAC address of 00:00:5E:00:01:01 and send DLSw traffic to the remote host with a MAC address of 00:02:00:00:00:01. Router 1 and Router 2 are configured for DLSw redundancy and map the local destination MAC address to the remote MAC address. Router 1 is also the designated master. If Router 1 becomes unavailable, Router 2 is the backup router.

Figure 4: DLSw Ethernet Redundancy Topology



To configure DLSw Ethernet redundancy, do the following:

Configuration on Router 1 Configure the redundancy group, redundancy group options, and the priority cost of each redundancy group option:

```
[edit]
interfaces {
  fe-0/0/0 {
    unit 0 {
      family llc2 {
        redundancy-group 1 {
          advertise-interval 1;
          map {
            local-mac 00:00:5e:00:01:01 remote-mac 00:02:00:00:00:01;
          }
          preempt hold-time 20;
          priority 200;
          track {
            dls {
              destination 00:02:00:00:00:01 priority-cost 50;
            }
          }
        }
      }
    }
  }
}
```



```

        peer 10.10.10.10 priority-cost 25;
    }
    interface e1-0/0/2.0 priority-cost 40;
    }
}
}
}
}
}
}
}

```

Configuration on Router 2 Configure the redundancy group, redundancy group options, and the priority cost of each redundancy group option:

```

[edit]
interfaces {
  fe-0/0/1 {
    unit 0 {
      family llc2 {
        redundancy-group 1 {
          map {
            local-mac 00:00:5e:00:01:01 remote-mac 00:02:00:00:00:01;
          }
          priority-cost 190;
          track {
            dlsw {
              destination 00:02:00:00:00:01 priority-cost 50;
              peer 10.10.10.10 priority-cost 25;
            }
            interface e1-0/0/2.0 priority-cost 40;
          }
        }
      }
    }
  }
}
}
}
}
}
}

```


Chapter 32

Summary of Data Link Switching Configuration Statements

The following sections explain each of the data link switching (DLSw) configuration statements. The statements are organized alphabetically.

advertise-interval

Syntax	advertise-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>For J-series Services Routers only. For Ethernet interfaces configured for DLSw Ethernet redundancy, configure the advertisement interval of DLSw neighbors on the network.</p> <p>All routers in the redundancy group must use the same advertisement interval.</p>
Options	<p><i>seconds</i>—Time interval between advertisement packets.</p> <p>Range: 1 through 255 seconds</p> <p>Default: 1 second</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

circuit-weight

Syntax	<code>circuit-weight <i>weight</i>;</code>
Hierarchy Level	[edit protocols dlsw remote-peer <i>peer-address</i>]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Configure the extent to which this peer should participate in establishing circuits. The higher the circuit weight, the greater the percentage of total circuits established with this remote peer.
Options	<i>weight</i> —Value of this peer in establishing circuits. Range: 1 through 127 Default: 1
Usage Guidelines	See “Configuring Load Balancing” on page 613.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

connection-idle-timeout

Syntax	<code>connection-idle-timeout <i>time</i>;</code>
Hierarchy Level	[edit protocols dlsw]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure an idle timeout for the connection with the remote DLSw peer.
Options	<i>time</i> —Configure the time period for which the connection to the peer can remain idle before the connection is taken down, in seconds. Range: 0 through 60,000 seconds Default: 60 seconds
Usage Guidelines	See “Configuring the Idle Timeout” on page 615.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

cost

Syntax	<code>cost cost;</code>
Hierarchy Level	[edit protocols dlsw remote-peer <i>peer-address</i>]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Configure the preference for establishing circuits with this peer. The lower the cost, the higher the preference.
Options	<p><i>cost</i>—Preference for this peer in establishing circuits.</p> <p>Range: 0 through 127</p> <p>Default: 100</p>
Usage Guidelines	See “Configuring Load Balancing” on page 613.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

destination

Syntax	<code>destination <i>mac-address</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw]</p>
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J-series Services Routers only. For Ethernet interfaces configured for DLSw Ethernet redundancy, enable tracking options for a destination MAC address.
Options	<p><i>mac-address</i>—Local MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn.nn.nn.nn.nn.nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p><i>priority-cost cost</i>—Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

destination-interface

Syntax	<code>destination-interface <i>interface-name</i>;</code>
Hierarchy Level	<code>[edit protocols dls w dls w-cos]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure a destination interface for class-of-service (CoS) traffic classification.
Options	<i>interface-name</i> —Name of the destination interface.
Usage Guidelines	See “Configuring Class of Service” on page 615.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

dls w

Syntax	<code>dls w { destination <i>mac-address</i> priority-cost <i>priority</i>; peer <i>ip-address</i> priority-cost <i>priority</i>; }</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track]</code>
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J-series Services Routers only. For Ethernet interfaces configured for DLSw, enable tracking options for a remote peer or destination MAC address. The statements are explained separately.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dlsw-cos

Syntax	dlsw-cos { destination-interface <i>interface-name</i> ; type-of-service <i>number</i> ; }
Hierarchy Level	[edit protocols dlsw]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure properties for outgoing DLSw traffic based on the type-of-service (ToS) classification.
Options	The remaining statements are explained separately.
Usage Guidelines	See “Configuring Class of Service” on page 615.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

explorer-wait-time

Syntax	explorer-wait-time <i>seconds</i> ;
Hierarchy Level	[edit protocols dlsw]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Configure the interval during which the DLSw router waits for responses to its explorer requests.
Options	<i>seconds</i> —Time interval for explorer responses. Range: 5 through 60 seconds Default: 10 seconds
Usage Guidelines	See “Configuring DLSw Timers” on page 613.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hold-time

Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> preempt], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> preempt]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Hold time before a higher-priority backup router preempts the master router.
Default	No DLSw preemption.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (No DLSw preemption)
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface

Syntax	<code>interface <i>interface-name</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J-series Services Routers only. For Ethernet interfaces configured for DLSw Ethernet redundancy, enable tracking options for an interface.
Options	<i>interface-name</i> —Interface name. Include the logical portion of the name, which corresponds to the logical unit number. <i>peer ip-address</i> —IP address of the remote peer. <i>priority-cost cost</i> —Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

load-balance

Syntax	<code>load-balance circuit-weight;</code>
Hierarchy Level	[edit protocols dlsw]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Enable load balancing based on circuit weight, in which the router picks among the available peers that have the same lowest cost value.
Default	By default, the <code>load-balance</code> statement is disabled.
Options	<i>circuit-weight</i> —Use the circuit-weight algorithm for load-balancing. (This is the only option available.)
Usage Guidelines	See “Configuring Load Balancing” on page 613.
Required Privilege Level	<i>routing</i> —To view this statement in the configuration. <i>routing-control</i> —To add this statement to the configuration.

local-mac

Syntax	local-mac <i>mac-address</i> remote-mac <i>mac-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> map]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J-series Services Routers only. For Ethernet interfaces configured for DLSw, specify the local MAC address to be mapped to a remote destination MAC address.
Options	<i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn.nn.nn.nn.nn.nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

local-peer

Syntax	local-peer <i>peer-address</i> ;
Hierarchy Level	[edit protocols dls w]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure the local peer.
Options	<i>peer-address</i> —IP address of the local DLSw peer.
Usage Guidelines	See “Configuring the Local Peer” on page 614.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

map

Syntax	map { local-mac <i>mac-address</i> remote-mac <i>mac-address</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J-series Services Routers only. For Ethernet interfaces configured for DLSw Ethernet redundancy, map a local peer MAC address to a remote peer MAC address. The statements are explained separately.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

multicast-address

Syntax	multicast-address <i>address</i> ;
Hierarchy Level	[edit protocols dlsw]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure a multicast address for the DLSw connection and enable multicast functionality.
Options	<i>address</i> —Multicast address.
Usage Guidelines	See “Configuring the Multicast Address” on page 615.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-preempt

See preempt

peer

Syntax	<code>peer ip-address priority-cost priority;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J-series Services Routers only. For Ethernet interfaces configured for DLSw, enable tracking options for a remote peer.
Options	<i>ip-address</i> —IP address of the remote peer. <i>priority-cost cost</i> —Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

preempt

Syntax	(preempt no-preempt) { hold-time <i>seconds</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Configure whether a backup router can preempt a master router: <ul style="list-style-type: none"> ■ preempt—Allow the master router to be preempted. ■ no-preempt—Prohibit the preemption of the master router. <p>The remaining statement is explained separately.</p>
Default	If you omit this statement, the backup router cannot preempt a master router.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

priority

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet redundancy-group <i>group-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet redundancy-group <i>group-number</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	When configuring DLSw Ethernet redundancy on Fast Ethernet and Gigabit Ethernet interfaces, configure a DLSw router's priority for becoming the master router. The router with the highest priority within the group becomes the master.
Options	<i>priority</i> —Router's priority for being elected master router in the VRRP group. A larger value indicates a higher priority for being elected. Range: 1 through 255 Default: 100 (for backup routers)
Usage Guidelines	See "Configuring DLSw Ethernet Redundancy Using LLC2 Properties" on page 618.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

promiscuous

Syntax	<code>promiscuous;</code>
Hierarchy Level	[edit protocols dlsw]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Enable the router to accept all incoming peer connections.
Usage Guidelines	See "Minimum DLSw Configuration" on page 612.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

protocols

Syntax	protocols dls { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the properties associated with a particular protocol.
Options	dls—Enable DLSw properties.
Usage Guidelines	See “Minimum DLSw Configuration” on page 612.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

reachability-cache-timeout

Syntax	reachability-cache-timeout <i>seconds</i> ;
Hierarchy Level	[edit protocols dls]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Configure the interval for retaining entries in the reachability cache.
Options	<i>seconds</i> —Lifetime of reachability cache entries. Range: 0 through 3600 seconds Default: 900 seconds
Usage Guidelines	See “Configuring DLSw Timers” on page 613.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

receive-initial-pacing

Syntax	receive-initial-pacing <i>number</i> ;
Hierarchy Level	[edit protocols dlsr]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure the size of the initial receive pacing window for incoming transport connections with the DLSr peer.
Options	<i>number</i> —The pacing window size. Range: 1 through 255 Default: 32
Usage Guidelines	See “Configuring the Initial Pacing Window” on page 615.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

redundancy-group

Syntax	<pre> redundancy-group <i>group-number</i> { advertise-interval <i>seconds</i>; map { local-mac <i>mac-address</i> remote-mac <i>mac-address</i>; } preempt hold-time <i>seconds</i>; no-preempt; priority <i>priority</i>; track { dls { peer <i>ip-address</i> priority-cost <i>priority</i>; destination <i>mac-address</i> priority-cost <i>priority</i>; } interface <i>interface-name</i> priority-cost <i>priority</i>; } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>For J-series Services Routers only. On Ethernet interfaces configured for DLSw, configure the router for DLSw redundancy.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

remote-mac

See local-mac

remote-peer

Syntax	<pre>remote-peer <i>peer-address</i> { circuit-weight <i>weight</i>; cost <i>cost</i>; }</pre>
Hierarchy Level	[edit protocols dls w]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure the IP address of the remote DLSw peer.
Options	<p><i>peer-address</i>—IP address of the remote DLSw peer.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring the Remote Peer” on page 612.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

traceoptions

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp><replace> <size <i>size</i>> <(world-readable no-world-readable)>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit protocols dls w]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	<p>Configure DLSw protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	The default protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—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. We recommend that you place ES-IS tracing output in the file esis-log.</p> <p>files <i>number</i>—(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.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag—Tracing operation to perform. To specify more than one flag, include multiple flag statements.</p> <ul style="list-style-type: none"> ■ all—Everything ■ memory—Memory traces ■ packets—Packet traces ■ parse—Configuration parse traces ■ route-socket—Route socket traces

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- *detail*—Detailed trace information
- *receive*—Packets being received
- *send*—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Disallow any user to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace 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: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Usage Guidelines See “Tracing DLSw Protocol Traffic” on page 617.

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.

track

Syntax	<pre>track { dlsw { destination <i>mac-address</i> priority-cost <i>priority</i>; peer <i>ip-address</i> priority-cost <i>priority</i>; } interface <i>interface-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>]</p>
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>For J-series Services Routers only. On Ethernet interfaces configured for DLSw Ethernet redundancy, enable tracking options for an interface, remote peer, or destination MAC address.</p> <p>The statements are explained separately.</p>
Options	<p>destination <i>mac-address</i>—Local MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn.nn.nn.nn.nn.nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>dlsw—DLSw protocol.</p> <p>interface <i>interface-name</i>—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.</p> <p>peer <i>ip-address</i>—IP address of the remote peer.</p> <p>priority-cost <i>cost</i>—Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 618.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

type-of-service

Syntax	type-of-service <i>number</i> ;
Hierarchy Level	[edit protocols dlsw dlsw-cos]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure the type-of-service (ToS) value used for CoS.
Options	<i>number</i> —Type-of-service value. Range: 0 through 255 Default: 0
Usage Guidelines	See “Configuring Class of Service” on page 615.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

Part 4

Encryption Services

- Encryption Overview on page 645
- Encryption Interfaces Configuration Guidelines on page 647
- Summary of Encryption Configuration Statements on page 657

Chapter 33

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 34

Encryption Interfaces Configuration Guidelines

To enable encryption interfaces, you can configure the following properties:

- Configuring an Encryption Interface on page 647
- Configuring Traffic on page 649
- Configuring an ES Tunnel Interface for a Layer 3 VPN on page 654
- Configuring ES PIC Redundancy on page 654
- Configuring IPSec Tunnel Redundancy on page 655

Configuring an Encryption Interface

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 routing platforms.

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

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 Traffic” on page 649.

Configuring the MTU for an Encryption Interface

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 Traffic

This section contains the following topics:

- Traffic Overview on page 649
- Configuring the Security Association on page 650
- Configuring an Outbound Traffic Filter on page 651
- Applying the Outbound Traffic Filter on page 652
- Configuring an Inbound Traffic Filter on page 652
- Applying the Inbound Traffic Filter to the Encryption Interface on page 653

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 5 on page 649, 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 5: 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” on page 648.

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 5 on page 649). 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 653. 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 653.



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 routing platforms 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 routing platforms.

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 653.

```
[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 6 on page 656 shows IPSec primary and backup tunnels.

Figure 6: 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 routing platforms.

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 959.

Chapter 35

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 an Encryption Interface” on page 647.

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 <i>destination-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel interfaces, specify the remote address of the backup tunnel.
Options	<i>destination-address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring IPSec Tunnel Redundancy” on page 655.
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 <i>interface-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> 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	<i>interface-name</i> —Name of ES interface to serve as the backup.
Usage Guidelines	See “Configuring ES PIC Redundancy” on page 654.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination

Syntax	<code>destination <i>destination-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel and encryption interfaces, specify the remote address.
Options	<i>destination-address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring an Encryption Interface” on page 647, “Configuring Traffic Sampling” on page 675, and “Configuring Flow Monitoring” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

es-options

Syntax	<code>es-options { backup-interface <i>interface-name</i>; }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On ES interfaces, configure ES interface-specific interface properties. The <code>backup-interface</code> statement is explained separately.
Usage Guidelines	See “Configuring ES PIC Redundancy” on page 654.
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 <i>sa-name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure protocol family information for the logical interface.
Options	<p><i>family</i>—Protocol family:</p> <ul style="list-style-type: none"> ■ <i>ccc</i>—Circuit cross-connect protocol suite ■ <i>inet</i>—IP version 4 suite ■ <i>inet6</i>—IP version 6 suite ■ <i>iso</i>—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite ■ <i>mlfr-end-to-end</i>—Multilink Frame Relay FRF.15 ■ <i>mlfr-uni-nni</i>—Multilink Frame Relay FRF.16 ■ <i>multilink-ppp</i>—Multilink Point-to-Point Protocol ■ <i>mpls</i>—Multiprotocol Label Switching ■ <i>tcc</i>—Translational cross-connect protocol suite ■ <i>tnp</i>—Trivial Network Protocol ■ <i>vpls</i>—Virtual private LAN service <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Encryption Interfaces Configuration Guidelines” on page 647; for a general discussion of family statement options, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

filter

Syntax	filter { input <i>filter-name</i> ; output <i>filter-name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the filters to be applied on an interface.
Options	input <i>filter-name</i> —Identifier for the input filter. output <i>filter-name</i> —Identifier for the output filter.
Usage Guidelines	See “Configuring Traffic” on page 649.
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 <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ipsec-sa

Syntax	<code>ipsec-sa sa-name;</code>
Hierarchy Level	[edit interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the IP Security (IPSec) SA name associated with the interface.
Options	<i>sa-name</i> —IPSec SA name.
Usage Guidelines	See “Encryption Interfaces Configuration Guidelines” on page 647.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS System Basics Configuration Guide</i>

source

Syntax	<code>source source-address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel and encryption interfaces, specify the source address.
Options	<i>source-address</i> —Address of the source side of the connection.
Usage Guidelines	See “Configuring an Encryption Interface” on page 647, “Configuring Traffic Sampling” on page 675, and “Configuring Flow Monitoring” on page 683.
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 647 and “Tunnel Interfaces Configuration Guidelines” on page 959.

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 647; 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 667
- Flow Monitoring and Discard Accounting Configuration Guidelines on page 671
- Summary of Flow-Monitoring Configuration Statements on page 719
- Flow Collection Configuration Guidelines on page 775
- Summary of Flow Collection Configuration Statements on page 787
- Dynamic Flow Capture Configuration Guidelines on page 805
- Flow-Tap Configuration Guidelines on page 817
- Summary of Dynamic Flow Capture and Flow-Tap Configuration Statements on page 823

Chapter 36

Flow Monitoring and Discard Accounting Overview

Using a Juniper Networks M-series or T-series routing platform, a selection of PICs (including the Monitoring Services PIC, Adaptive Services [AS] PIC, or MultiServices PIC) 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).



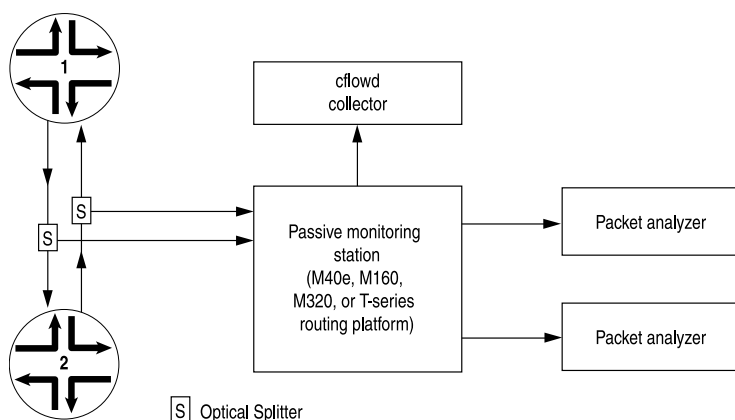
NOTE: Monitoring Services PICs, AS PICs, and MultiServices PICs must be mounted on an Enhanced Flexible PIC Concentrator (FPC) in an M-series or T-series routing platform.

This section provides general information on the following topics:

- Passive Flow Monitoring on page 667
- Active Flow Monitoring on page 668

Passive Flow Monitoring

The routing platform 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 7 on page 668 shows a typical topology for the passive flow-monitoring application.

Figure 7: 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 routing platform. 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:

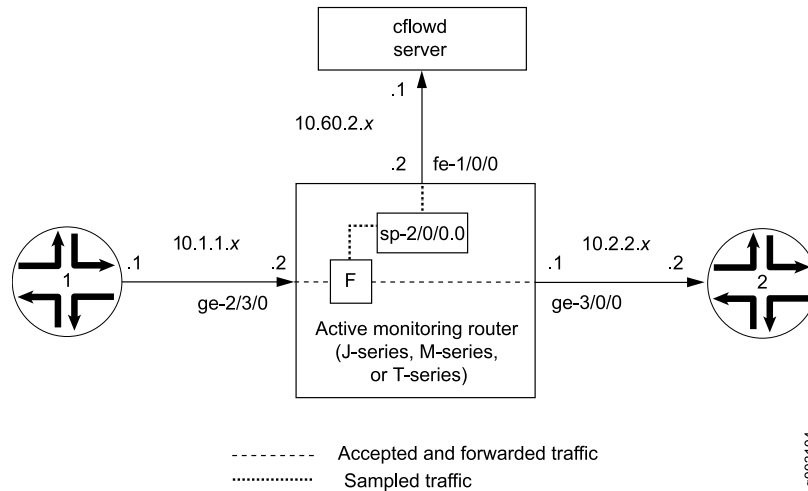
- The router can perform sampling or port mirroring at any one time.
- The router can perform forwarding or discard accounting at any one time.

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 8 on page 670 shows a sample topology.

Figure 8: Active Monitoring Configuration Topology



In Figure 8 on page 670, 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 37

Flow Monitoring and Discard Accounting Configuration Guidelines

To configure flow monitoring and accounting properties, include the following statements at the [edit interfaces] hierarchy level:

```
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:

```
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 [ group-names ] {
  interface interface-name {
    next-hop [ addresses ];
  }
}

port-mirroring {
  input {
    family (inet | inet6) {
      rate rate;
      run-length number;
    }
  }
}

```

```

output {
  interface interface-name {
    next-hop address;
  }
  no-filter-check;
}
traceoptions {
  file filename {
    files number;
    size bytes;
    (world-readable | no-world-readable);
  }
}
}
sampling {
  disable;
  input {
    family (inet | mpls) {
      max-packets-per-second number;
      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 cflowd version 9 properties, include the following statements at the [edit services] hierarchy level:

```

flow-monitoring {
    version9 {
        template template-name {
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            ipv4-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 describes the following tasks for configuring traffic sampling and flow-monitoring properties:

- Minimum Traffic Sampling or Forwarding Configuration on page 675
- Configuring Traffic Sampling on page 675
- Configuring Flow Monitoring on page 683
- Configuring Flow Aggregation on page 687
- Configuring Port Mirroring on page 698
- Load Balancing Among Multiple Monitoring Interfaces on page 707
- Configuring Discard Accounting on page 710
- Enabling Passive Flow Monitoring on page 711
- Flow-Monitoring Redundancy on page 716

Minimum Traffic Sampling or Forwarding Configuration

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

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.

The following sections provide information about traffic sampling configuration:

- Configuring Traffic Sampling Properties on page 676
- Disabling Traffic Sampling on page 677
- Configuring Traffic Sampling Output on page 677
- Tracing Traffic Sampling Operations on page 679
- Examples: Configuring Traffic Sampling on page 680

Configuring Traffic Sampling Properties

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 9 on page 677).

Figure 9: Configure Sampling Rate**Rate and Run-length**

Case #1 Rate =10, run-length =4



1746

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 $x = \text{run-length} + 1$. 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 677.

Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the **disable** statement at the [edit forwarding-options sampling] hierarchy level:

```
disable;
```

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 “Configuring Flow Aggregation” on page 687.

The **aggregate-export-interval** statement is described in “Configuring Discard Accounting” on page 710, and the **flow-active-timeout** and **flow-inactive-timeout** statements are described in “Configuring Flow Monitoring” on page 683.

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 Files

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
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
# Feb  1 20:31:21
# Dest                Src Dest  Src Proto TOS  Pkt Intf  IP    TCP
# addr                addr  port  port
# len  num  frag flags
# 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);
  }
}
```

Examples: Configuring Traffic Sampling

The following sections provide examples of configuring traffic sampling:

- Sampling a Single SONET Interface on page 680
- Sampling All Traffic from a Single IP Address on page 681
- Sampling All FTP Traffic on page 682

Sampling a Single SONET 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;
    }
}

```

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;
    }
}

```

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 the Flow-Monitoring Interface on page 683
- Configuring Flow-Monitoring Properties on page 684
- Example: Configuring Flow Monitoring on page 686

Configuring the Flow-Monitoring Interface

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. 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:

- Configuring the Flow-Monitoring Interface on page 684
- Exporting Flows on page 685
- Configuring Timers on page 685

Configuring the Flow-Monitoring Interface

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 Simple Network Management Protocol (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 platforms, 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 775.

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 “Configuring Flow Aggregation” on page 687.

Configuring Timers

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 “Configuring Flow Aggregation” on page 687.

```
[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;
}
}
}
}

```

Configuring 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):

```
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 traffic, but you cannot 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.

The following sections describe the configuration methods:

- Configuring Version 5 or Version 8 cflowd on page 688
- Configuring Version 9 Flow Templates on page 690
- Replicating Flows to Multiple Flow Servers on page 695
- Debugging cflowd Flow Aggregation on page 697

Configuring 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 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.

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 698.

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 Version 9 Flow Templates

Use of version 9 allows you to define a flow record template suitable for IPv4 traffic, MPLS traffic, or a combination of the two. 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.

The following sections contain additional information:

- Configuring the Traffic to be Sampled on page 690
- Configuring the Version 9 Template Properties on page 691
- Restrictions on page 692
- Fields Included in Each Template Type on page 692
- MPLS Sampling Behavior on page 693
- Verification on page 694
- Examples: Configuring Version 9 Flow Templates on page 694

Configuring the Traffic to be Sampled

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

```
forwarding-options {
```

```
sampling {
  input {
    family (inet | mpls) {
      max-packets-per-second value;
      rate value;
      run-length value;
    }
  }
}
```

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:

```
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 | mpls-ipv4-template | mpls-template) {
          label-position [ label ];
        }
      }
    }
  }
}
```

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`, `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.

Restrictions

The following restrictions apply to version 9 templates:

- You can export only one version 9 template at a time.
- You can sample packets from only one protocol family (**inet** or **mpls**) at a time.
- 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

The MPLS template includes the following specific fields:

- MPLS Label #1
- MPLS Label #2
- MPLS Label #3

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;
      }
    }
  }
}

```

Replicating 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:

- Replicating Routing Engine–Based Sampling to Multiple Flow Servers on page 696
- Replicating Version 9 Flow Aggregation to Multiple Flow Servers on page 696

Replicating Routing Engine–Based Sampling to Multiple Flow 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;
      }
    }
  }
}
```

Replicating Version 9 Flow Aggregation to Multiple Flow 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;
      }
    }
  }
}
```

Debugging cflowd Flow Aggregation

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 677.



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:

```
input {
  family (inet | inet6) {
    rate rate;
    run-length number;
  }
}
output {
  interface interface-name {
    next-hop address;
  }
  no-filter-check;
}
traceoptions {
  file filename {
    files number;
    size bytes;
    (world-readable | no-world-readable);
  }
}
```

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.

The interface used to send the packets to the analyzer is the output interface configured above at the [edit forwarding-options port-mirroring 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 family (inet | inet6)] 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 679.

For more information about port mirroring, see the following sections:

- Configuring Tunnels on page 700
- Filter-Based Forwarding with Multiple Monitoring Interfaces on page 701
- Restrictions on page 701
- Configuring Port Mirroring on Services Interfaces on page 702
- Examples: Configuring Port Mirroring on page 703

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 959.

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:

```
next-hop-group [ group-names ] {
  interface interface-name {
    next-hop [ addresses ];
  }
}
```

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.

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 703.

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, T320, and T640 routing platforms, simultaneous IPv4 and IPv6 port mirroring is supported. Multiple next-hop mirroring is not supported.
- On M-series platforms 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 {
  family inet {
    input {
      rate 1;
    }
    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 {
        family inet {
            rate 1; # This rate port mirrors one packet for every one received (1:1 = all
                # packets).
        }
    }
    output { # This sends traffic to a tunnel interface to prepare for multipoint 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 ];
  }
}

```

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 {
    family inet {
      rate 1;
    }
  }
  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 “Configuring Flow Aggregation” on page 687. The **flow-active-timeout** and **flow-inactive-timeout** statements are described in “Configuring Flow Monitoring” on page 683.

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 routing platform:

- 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 *interface-name* unit *logical-unit-number*] hierarchy level:

```
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:

```
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:

```
family inet;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see “Configuring the Flow-Monitoring Interface” on page 683.

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:

```
receive-options-packets;
receive-ttl-exceeded;
```

For more information, see the following sections:

- Passive Flow Monitoring for MPLS Encapsulated Packets on page 712
- Example: Enabling Passive Flow Monitoring on page 714

Passive Flow Monitoring for MPLS Encapsulated Packets

On monitoring services interfaces, you can process Multiprotocol Label Switching (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:

```
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:

```

pop-all-labels {
    required-depth number;
}

```

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:

```

required-depth number;

```

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;
    }
}
}

```

Flow-Monitoring Redundancy

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 475. 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 38

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 [ hostnames ] {
      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 710.

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

address

Syntax	address <i>address</i> { destination <i>address</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interface address.
Options	<i>address</i> —Address of the interface. The remaining statement is explained separately.
Usage Guidelines	See “Configuring Flow Monitoring” on page 683 or “Configuring Traffic Sampling” on page 675.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other options not associated with flow monitoring.

aggregate-export-interval

Syntax	aggregate-export-interval <i>seconds</i> ;
Hierarchy Level	[edit forwarding-options accounting <i>name</i> 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	<i>seconds</i> —Duration.
Usage Guidelines	See “Configuring Discard Accounting” on page 710.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

aggregation

Syntax	<pre>aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; }</pre>
Hierarchy Level	[edit forwarding-options accounting output cflowd <i>hostname</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>]
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	<p>autonomous-system—Aggregate by autonomous system (AS) number.</p> <p>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 <i>cflowd Configuration Guide</i>, dated August 30, 1999.</p> <p>destination-prefix—Aggregate by destination prefix.</p> <p>protocol-port—Aggregate by protocol and port number.</p> <p>source-destination-prefix—Aggregate by source and destination prefix.</p> <p>source-prefix—Aggregate by source prefix.</p>
Usage Guidelines	See "Configuring Flow Aggregation" on page 687.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

autonomous-system-type

Syntax	autonomous-system-type (origin peer);
Hierarchy Level	[edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the type of AS numbers that cflowd exports.
Default	origin
Options	<p>origin—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.</p> <p>peer—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.</p>
Usage Guidelines	See “Configuring Flow Aggregation” on page 687.
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 724
- cflowd (Flow Monitoring) on page 725

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 “Configuring Flow Aggregation” on page 687.

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

cflowd (Flow Monitoring)

Syntax	<code>cflowd hostname { port port-number; }</code>
Hierarchy Level	[edit forwarding-options monitoring <i>name</i> family inet output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility <code>cfcollect</code>.</p> <p>You can configure up to eight version 5 flow formats at the [edit forwarding-options monitoring <i>name</i> output] hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.</p>
Options	<p><i>hostname</i>—The IP address or identifier of the host system (the workstation running the <code>cflowd</code> utility).</p> <p>The remaining statement is explained separately.</p>
Usage Guidelines	See “Configuring Flow Aggregation” on page 687.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

core-dump

Syntax	<code>(core-dump no-core-dump);</code>
Hierarchy Level	[edit interfaces <i>mo-fpc/pic/port</i> multiservice-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory <code>/var/tmp</code> 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):</p> <ul style="list-style-type: none"> ■ <code>core-dump</code>—Enable the core dumping operation. ■ <code>no-core-dump</code>—Disable the core dumping operation.
Usage Guidelines	See “Configuring Flow Monitoring” on page 683.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

destination

Syntax	<code>destination destination-address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel interfaces, specify the remote address of the tunnel.
Options	<i>destination-address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring a Unicast Tunnel” on page 959, “Configuring Traffic Sampling” on page 675, and “Configuring Flow Monitoring” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable

Syntax	<code>disable;</code>
Hierarchy Level	[edit forwarding-options sampling], [edit forwarding-options sampling output file]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable traffic accounting or sampling.
Usage Guidelines	See “Configuring Traffic Sampling” on page 675.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

engine-id

Syntax	<code>engine-id <i>number</i>;</code>
Hierarchy Level	[edit forwarding-options accounting <i>name</i> output interface <i>interface-name</i>], [edit forwarding-options monitoring <i>name</i> output interface <i>interface-name</i>], [edit forwarding-options sampling output interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the engine ID number for flow monitoring and accounting services.
Options	<i>number</i> —Identity of accounting interface.
Usage Guidelines	See “Configuring Traffic Sampling” on page 675, “Configuring Flow Monitoring” on page 683, or “Configuring Discard Accounting” on page 710.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

engine-type

Syntax	<code>engine-type <i>number</i>;</code>
Hierarchy Level	[edit forwarding-options accounting <i>name</i> output interface <i>interface-name</i>], [edit forwarding-options monitoring <i>name</i> output interface <i>interface-name</i>], [edit forwarding-options sampling output interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the engine type number for flow monitoring and accounting services.
Options	<i>number</i> —Platform-specific accounting interface type.
Usage Guidelines	See “Configuring Traffic Sampling” on page 675, “Configuring Flow Monitoring” on page 683, or “Configuring Discard Accounting” on page 710.
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 <i>format</i> ;
Hierarchy Level	[edit forwarding-options monitoring <i>name</i> output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Flow monitoring export format.
Options	<i>format</i> —Format of the flows. Values: 5 or 8 Default: 5
Usage Guidelines	See “Exporting Flows” on page 685.
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 729
- family (Monitoring) on page 730
- family (Port Mirroring) on page 731
- family (Sampling) on page 731

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 683.

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 683.

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 { rate <i>number</i> ; run-length <i>number</i> ; }
Hierarchy Level	[edit forwarding-options port-mirroring input]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the protocol family to be sampled. Only IPv4 (<i>inet</i>) is supported. The statements are explained separately.
Usage Guidelines	See “Configuring Port Mirroring” on page 698.
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 mpls) { max-packets-per-second <i>number</i> ; rate <i>number</i> ; run-length <i>number</i> ; }
Hierarchy Level	[edit forwarding-options sampling input]
Release Information	Statement introduced before JUNOS Release 7.4. mpls option added in Release 8.3.
Description	Configure the protocol family to be sampled. IPv4 (<i>inet</i>) is supported for most purposes, but you can configure <i>family mpls</i> to collect and export MPLS label information using flow aggregation version 9. The remaining statements are explained separately.
Usage Guidelines	See “Configuring Traffic Sampling” on page 675.
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 732
- file (Trace Options) on page 733

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 675.

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 <i>filename</i> { files <i>number</i> ; size <i>bytes</i> ; (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	<i>filename</i> —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 679.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

filename

Syntax	filename <i>filename</i> ;
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	<i>filename</i> —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 675.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

files

Syntax	<code>files <i>number</i>;</code>
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	<p><i>number</i>—Maximum number of traffic sampling or trace log files. When a file named <i>sampling-file</i> reaches its maximum size, it is renamed <i>sampling-file.0</i>, then <i>sampling-file.1</i>, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.</p> <p>Range: 1 through 100 files</p> <p>Default: 5 files for sampling output; 10 files for trace log information</p>
Usage Guidelines	See “Configuring Port Mirroring” on page 698 or “Configuring Traffic Sampling” on page 675.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

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 683.

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 Timers” on page 685 or “Configuring the Version 9 Template Properties” on page 691.

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 <i>group-name</i> 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 685.
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 <i>seconds</i> ;
Hierarchy Level	[edit forwarding-options accounting <i>name</i> output], [edit forwarding-options monitoring <i>name</i> 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	<i>seconds</i> —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 Timers” on page 685 or “Configuring the Version 9 Template Properties” on page 691.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

flow-monitoring

Syntax	<pre> flow-monitoring { version9 { template <i>template-name</i> { flow-active-timeout <i>seconds</i>; flow-inactive-timeout <i>seconds</i>; ipv4-template; mpls-template { label-position [<i>positions</i>]; } mpls-ipv4-template { label-position [<i>positions</i>]; } option-refresh-rate packets <i>packets</i> seconds <i>seconds</i>; template-refresh-rate packets <i>packets</i> seconds <i>seconds</i>; } } } </pre>
Hierarchy Level	[edit services]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>Specify the active monitoring properties for flow aggregation version 9.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring Version 9 Flow Templates” on page 690.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

forwarding-options

Syntax	forwarding-options { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure traffic forwarding.</p> <p>The statements that apply to services interfaces are explained separately. For other statements, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>
Usage Guidelines	See “Flow Monitoring and Discard Accounting Configuration Guidelines” on page 671.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

input

See the following sections:

- input (Port Mirroring) on page 739
- input (Sampling) on page 739

input (Port Mirroring)

Syntax input {
 family inet {
 rate *number*;
 run-length *number*;
 }
 }

Hierarchy Level [edit forwarding-options port-mirroring]

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 698.

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 | 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.

Description Configure traffic sampling on a logical interface.

The statements are explained separately.

Usage Guidelines See “Configuring Traffic Sampling” on page 675.

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 <i>number</i> ;
Hierarchy Level	[edit forwarding-options monitoring <i>name</i> output interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify a value for the input interface index that overrides the default supplied by Simple Network Management Protocol (SNMP).
Options	<i>number</i> —Input interface index value.
Usage Guidelines	See “Configuring Flow Monitoring” on page 683.
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 741
- interface (Monitoring) on page 742
- interface (Port Mirroring) on page 742

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 710 or “Configuring Traffic Sampling” on page 675.

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 683.

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 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 698.

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 <i>JUNOS Network Interfaces Configuration Guide</i> 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 <i>template-name</i>]
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 Version 9 Flow Templates” on page 690.
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 [<i>positions</i>];
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i> mpls-ipv4-template], [edit services flow-monitoring version9 template <i>template-name</i> 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	<i>positions</i> —Numbered positions for the labels.
Usage Guidelines	See “Configuring Version 9 Flow Templates” on page 690.
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 <i>hostname</i>]
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 “Configuring Flow Aggregation” on page 687.
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 <i>number</i> ;
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	<i>number</i> —Maximum number of packets per second. Range: 0 through 65,535 Default: 1000
Usage Guidelines	See “Configuring Traffic Sampling” on page 675.
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 683.

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 [<i>positions</i>]; }
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
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 Version 9 Flow Templates” on page 690.
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 [<i>positions</i>]; }
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
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 Version 9 Flow Templates” on page 690.
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 <i>mo-fpc/pic/port</i>]
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 683.
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 <i>address</i> ;
Hierarchy Level	[edit forwarding-options port-mirroring output interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the next-hop address for sending copies of packets to an analyzer.
Options	<i>address</i> —IP address of the next-hop router.
Usage Guidelines	See “Configuring Port Mirroring” on page 698.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

next-hop-group

Syntax `next-hop-group [group-names] {
 interface interface-name {
 next-hop [addresses];
 }
 }`

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 *addresses*—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-names—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 698.

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 output]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Disable filter checking on the port-mirroring interface.</p> <p>This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it.</p>
Usage Guidelines	See “Configuring Port Mirroring” on page 698.
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 <i>packets</i> seconds <i>seconds</i> ;
Hierarchy Level	[edit services flow-monitoring version9], [edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Specify the refresh rate, in either packets or seconds.
Options	<p><i>packets</i>—Refresh rate, in number of packets. Range: 1 through 480,000 Default: 4800</p> <p><i>seconds</i>—Refresh rate, in number of seconds. Range: 10 through 600 Default: 60</p>
Usage Guidelines	See “Configuring Version 9 Flow Templates” on page 690.
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 752
- output (Monitoring) on page 753
- output (Port Mirroring) on page 753
- output (Sampling) on page 754

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 710.

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

output (Monitoring)

Syntax output {
 cfowd *hostname* port *port-number*;
 export-format *format*;
 flow-active-timeout *seconds*;
 flow-export-destination {
 (cfowd-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 683.

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]

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 698.

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]

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 675.

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 <i>number</i> ;
Hierarchy Level	[edit forwarding-options monitoring <i>name</i> output interface <i>interface-name</i>]
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.
<hr/>	
	NOTE: On J-series platforms, cflowd sampling in the input direction of an interface reports the output interface index as 0.
<hr/>	
Options	<i>number</i> —Output interface index value.
Usage Guidelines	See “Configuring Flow Monitoring” on page 683.
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 <i>interface-name</i> unit <i>logical-unit-number</i>]
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 711.
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 <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options mpls], [edit interfaces <i>interface-name</i> fastether-options mpls], [edit interfaces <i>interface-name</i> gigether-options mpls], [edit interfaces <i>interface-name</i> sonet-options mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, removes up to two Multiprotocol Label Switching (MPLS) labels from incoming IP packets.</p> <p>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.</p> <p>The remaining statement is explained separately.</p>
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 712.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

port

Syntax	<code>port <i>port-number</i>;</code>
Hierarchy Level	[edit forwarding-options accounting <i>name</i> output cflowd <i>hostname</i>], [edit forwarding-options monitoring <i>name</i> family inet output cflowd <i>hostname</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the User Datagram Protocol (UDP) port number on the cflowd host system.
Options	<i>port-number</i> —Any valid UDP port number on the host system.
Usage Guidelines	See “Configuring Flow Aggregation” on page 687.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

port-mirroring

```

Syntax  port-mirroring {
            input {
              family inet {
                rate rate;
                run-length number;
              }
            }
            output {
              interface interface-name {
                next-hop address;
              }
              no-filter-check;
            }
            traceoptions {
              file filename {
                files number;
                size bytes;
                (world-readable | no-world-readable);
              }
            }
          }

```

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 698.

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

rate

Syntax	<code>rate number;</code>
Hierarchy Level	[edit forwarding-options port-mirroring input family inet], [edit forwarding-options sampling input family (inet mpls)]
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	<i>number</i> —Denominator of the ratio. Range: 1 through 65,535
Usage Guidelines	See “Configuring Port Mirroring” on page 698 or “Configuring Traffic Sampling” on page 675.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

receive-options-packets

Syntax	<code>receive-options-packets;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 711.
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 <i>interface-name</i> unit <i>logical-unit-number</i> 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 711.
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 <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> fastether-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> gigeether-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> 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	<i>number</i> —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 712.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

run-length

Syntax	<code>run-length <i>number</i>;</code>
Hierarchy Level	[edit forwarding-options port-mirroring input family inet], [edit forwarding-options sampling input family (inet mpls)]
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	<i>number</i> —Number of samples. Range: 0 through 20 Default: 0
Usage Guidelines	See “Configuring Port Mirroring” on page 698 or “Configuring Traffic Sampling” on page 675.
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 763](#)
- [sampling \(Interfaces\) on page 764](#)

sampling (Forwarding Options)

```

Syntax  sampling {
            disable;
            input {
                family (inet | mpls) {
                    max-packets-per-second number;
                    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 675.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

sampling (Interfaces)

Syntax	sampling <i>direction</i> ;
Hierarchy Level	[edit interfaces <i>mo-fpc/pic/port</i> unit <i>logical-unit-number</i> 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 683.
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 Version 9 Flow Templates” on page 690.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

size

Syntax	<code>size bytes;</code>
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	<p>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.</p> <p>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.</p>
Options	<p>bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).</p> <p>Syntax: <code>xk</code> to specify KB, <code>xm</code> to specify MB, or <code>xg</code> to specify GB</p> <p>Range: 10 KB through the maximum file size supported on your router</p> <p>Default: 1 MB for sampling data; 128 KB for log information</p>
Usage Guidelines	See “Configuring Port Mirroring” on page 698 or “Configuring Traffic Sampling” on page 675.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

source-address

Syntax	source-address <i>address</i> ;
Hierarchy Level	[edit forwarding-options accounting <i>name</i> output interface <i>interface-name</i>], [edit forwarding-options monitoring <i>name</i> family inet output interface <i>interface-name</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>], [edit forwarding-options sampling output interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the source address for monitored packets.
Options	<i>address</i> —Interface source address.
Usage Guidelines	See “Configuring Discard Accounting” on page 710, “Configuring Flow Monitoring” on page 683, or “Configuring Traffic Sampling” on page 675.
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 675.
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 <i>mo-fpc/pic/port</i> multiservice-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the <code>/var/log</code> directory.</p> <ul style="list-style-type: none">■ <code>syslog</code>—Enable PIC system logging.■ <code>no-syslog</code>—Disable PIC system logging.
Usage Guidelines	See “Configuring Flow Monitoring” on page 683.
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 768
- [template \(Services\)](#) on page 769

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 Version 9 Flow Templates” on page 690.

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

template (Services)

Syntax	<pre> template <i>template-name</i> { flow-active-timeout <i>seconds</i>; flow-inactive-timeout <i>seconds</i>; ipv4-template; mpls-template { label-position [<i>positions</i>]; } mpls-ipv4-template { label-position [<i>positions</i>]; } option-refresh-rate packets <i>packets</i> seconds <i>seconds</i>; template-refresh-rate packets <i>packets</i> seconds <i>seconds</i>; } </pre>
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	<i>template-name</i> —Name of version 9 template.
Usage Guidelines	See “Configuring Version 9 Flow Templates” on page 690.
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 <i>packets</i> seconds <i>seconds</i> ;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Specify the refresh rate, in either packets or seconds.
Options	<i>packets</i> —Refresh rate, in number of packets. Range: 1 through 480,000 Default: 4800 <i>seconds</i> —Refresh rate, in number of seconds. Range: 10 through 600 Default: 60
Usage Guidelines	See “Configuring Version 9 Flow Templates” on page 690.
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 679.

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 not affect services interfaces.

version

Syntax	<code>version <i>format</i>;</code>
Hierarchy Level	[edit forwarding-options accounting <i>name</i> output cflowd <i>hostname</i>], [edit forwarding-options sampling output cflowd <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the version format of the aggregated flows exported to a cflowd server.
Options	<i>format</i> —Format of the flows. Values: 5 or 8 Default: 5
Usage Guidelines	See “Configuring Flow Aggregation” on page 687.
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 773
- version9 (Services) on page 774

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 Version 9 Flow Templates” on page 690.

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

version9 (Services)

Syntax	<pre> version9 { template <i>template-name</i> { flow-active-timeout <i>seconds</i>; flow-inactive-timeout <i>seconds</i>; ipv4-template; mpls-template { label-position [<i>positions</i>]; } mpls-ipv4-template { label-position [<i>positions</i>]; } option-refresh-rate packets <i>packets</i> seconds <i>seconds</i>; template-refresh-rate packets <i>packets</i> seconds <i>seconds</i>; } } </pre>
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 Version 9 Flow Templates” on page 690.
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 698 or “Configuring Traffic Sampling” on page 675.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

Chapter 39

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 Properties on page 777
- Sending cflowd Records to the Flow Collector Interface on page 780
- Enabling Flow Collection Mode and Interface on page 780
- Example: Flow Collector Interface Configuration on page 780

Configuring Flow Collection Properties

This section describes the following tasks for configuring flow collection:

- Configuring Flow Collector Destinations on page 777
- Configuring a Packet Analyzer on page 777
- Configuring File Formats on page 778
- Configuring Interface Mappings on page 778
- Configuring Transfer Logs on page 779
- Configuring Retry Attempts on page 779

Configuring Flow Collector Destinations

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:

```
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:

```
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:

```
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:

```
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:

```

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 the Flow Collector Interface

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:

```
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.

Enabling Flow Collection Mode and Interface

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:

```
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:

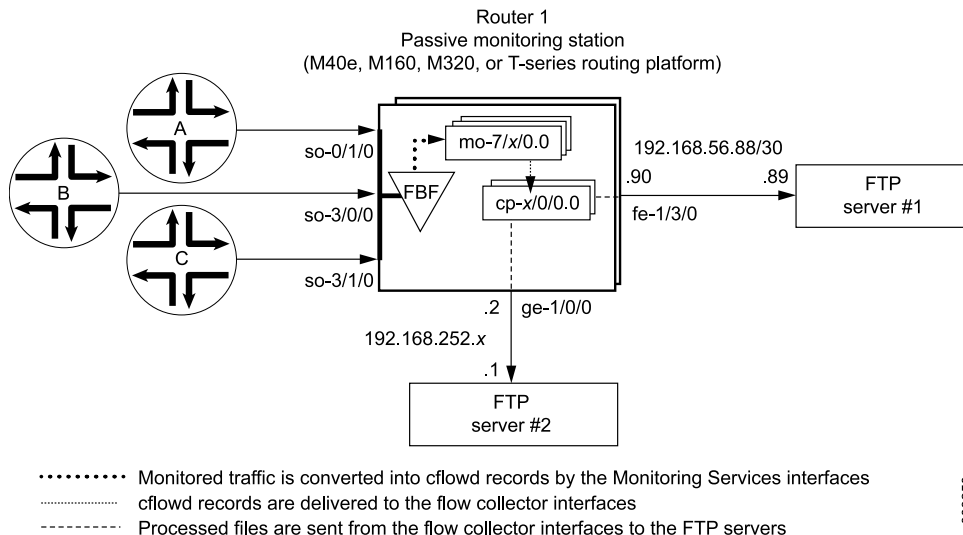
```
cp-fpc/pic/port {
  ...
}
```

Example: Flow Collector Interface Configuration

Figure 10 on page 781 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 10: Flow Collector Interface Topology Diagram



```

Router 1 [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-w2oZdbZDHmF3001"; # SECRET-DATA
            }
            "ftp://user@192.168.252.1//tmp/collect2/" { # The secondary FTP server.
                password "$9$elbvL7-dsgaGVwGjkP3nOBI"; # SECRET-DATA
            }
        }
    }
}

```

```

}
file-specification { # Define sets of flow collector characteristics here.
  def-spec {
    name-format "default-allInt-0-%D_%T-%I_%N.bcp.bi.gz";
    data-format flow-compressed; # The default compressed output format.
  } # When no overrides are specified, a collector uses default transfer values.
  f1 {
    name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
    data-format flow-compressed; # The default compressed output format.
    transfer timeout 1800 record-level 1000000; # Here are configured values.
  }
}
interface-map { # Allows you to map interfaces to flow collector interfaces.
  file-specification def-spec; # Flows generated for default traffic are sent to the
  collector cp-7/0/0; # default flow collector interface "cp-7/0/0".
  so-0/1/0.0 { # Flows generated for the so-0/1/0 interface are sent
    collector cp-6/0/0; # to cp-6/0/0, and the file-specification used is
  } # "default."
  so-3/0/0.0 { # Flows generated for the so-3/0/0 interface are sent
    file-specification f1; # to cp-6/0/0, and the file-specification used is "f1."
    collector cp-6/0/0;
  }
  so-3/1/0.0; # Because no settings are defined, flows generated for this
} # interface use interface cp-7/0/0 and the default file specification.
transfer-log-archive { # Sends flow collector interface log files to an FTP server.
  filename-prefix so_3_0_0_log;
  maximum-age 15;
  archive-sites {
    "ftp://user@192.168.56.89//tmp/transfers/" {
      password "$9$IFaEyevMXNVsWLsgaU.m6/C";
    }
  }
}
}
}
}

```

Chapter 40

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	<code>analyzer-address <i>address</i>;</code>
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	<i>address</i> —IP address for packet analyzer.
Usage Guidelines	See “Configuring a Packet Analyzer” on page 777.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

analyzer-id

Syntax	<code>analyzer-id <i>name</i>;</code>
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	<i>name</i> —Identifier for packet analyzer.
Usage Guidelines	See “Configuring a Packet Analyzer” on page 777.
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 779.

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 778.

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

data-format

Syntax	<code>data-format <i>format</i>;</code>
Hierarchy Level	[edit services flow-collector file-specification variant <i>variant-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the data format for a specific file format variant.
Options	<i>format</i> —Data format. Specify <code>flow-compressed</code> as the data format.
Usage Guidelines	See “Configuring File Formats” on page 778.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destinations

Syntax	<pre>destinations { ftp:url { password "<i>password</i>"; } }</pre>
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 Flow Collector Destinations” on page 777.
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 <i>prefix</i> ;
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	<i>prefix</i> —Filename identifier.
Usage Guidelines	See “Configuring Transfer Logs” on page 779.
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 791
- file-specification (Interface Mapping) on page 791

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 778.

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 778.

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 433.

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 795
- **ftp (Transfer Log Files)** on page 796

ftp (Flow Collector Files)

Syntax	<code>ftp:url;</code>
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	<p><i>url</i>—FTP server address. The URL can include the following macros, typed in braces:</p> <ul style="list-style-type: none"> ■ <code>{%D}</code>—Date ■ <code>{%T}</code>—Time when the file is created ■ <code>{%I}</code>—Description string for the logical interface configured using the collector <i>interface-name</i> statement at the [edit services flow-collector interface-map] hierarchy ■ <code>{%N}</code>—Unique, sequential number for each new file created ■ <code>{am_pm}</code>—AM or PM ■ <code>{date}</code>—Current date using the <code>{year}</code> <code>{month}</code> <code>{day}</code> macros ■ <code>{day}</code>—From 01 through 31 ■ <code>{day_abbrev}</code>—Sun through Sat ■ <code>{day_full}</code>—Sunday through Saturday ■ <code>{generation number}</code>—Unique, sequential number for each new file created ■ <code>{hour_12}</code>—From 01 through 12 ■ <code>{hour_24}</code>—From 00 through 23 ■ <code>{ifalias}</code>—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy ■ <code>{minute}</code>—From 00 through 59 ■ <code>{month}</code>—From 01 through 12 ■ <code>{month_abbrev}</code>—Jan through Dec ■ <code>{month_full}</code>—January through December ■ <code>{num_zone}</code>—From -2359 to +2359; this macro is not supported ■ <code>{second}</code>—From 00 through 60 ■ <code>{time}</code>—Time the file is created, using the <code>{hour_24}</code> <code>{minute}</code> <code>{second}</code> macros ■ <code>{time_zone}</code>—Time zone code name of the locale; for example, <code>gmt</code> (this macro is not supported). ■ <code>{year}</code>—In the format YYYY; for example, 1970 ■ <code>{year_abbrev}</code>—From 00 through 99
Usage Guidelines	See “Configuring Flow Collector Destinations” on page 777.

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 779.

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 778.

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 <i>minutes</i> ;
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 <i>minutes</i> —Transfer log file age. Range: 1 through 360
Usage Guidelines	See “Configuring Transfer Logs” on page 779.
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 " <i>format</i> ";
Hierarchy Level	[edit services flow-collector file-specification variant <i>variant-number</i>]
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	<p><i>format</i>—Specify the filename format, within quotation marks. The name format can include the following macros, typed in braces:</p> <ul style="list-style-type: none"> ■ <i>{%D}</i>—Date ■ <i>{%T}</i>—Time when the file is created ■ <i>{%I}</i>—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level ■ <i>{%N}</i>—Unique, sequential number for each new file created ■ <i>{am_pm}</i>—AM or PM ■ <i>{date}</i>—Current date using the <i>{year}</i> <i>{month}</i> <i>{day}</i> macros ■ <i>{day}</i>—From 01 through 31 ■ <i>{day_abbrev}</i>—Sun through Sat ■ <i>{day_full}</i>—Sunday through Saturday ■ <i>{generation number}</i>—Unique, sequential number for each new file created ■ <i>{hour_12}</i>—From 01 through 12 ■ <i>{hour_24}</i>—From 00 through 23 ■ <i>{ifalias}</i>—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level ■ <i>{minute}</i>—From 00 through 59 ■ <i>{month}</i>—From 01 through 12 ■ <i>{month_abbrev}</i>—Jan through Dec ■ <i>{month_full}</i>—January through December ■ <i>{num_zone}</i>—From -2359 through +2359; this macro is not supported ■ <i>{second}</i>—From 00 through 60 ■ <i>{time}</i>—Time the file is created, using the <i>{hour_24}</i> <i>{minute}</i> <i>{second}</i> macros ■ <i>{time_zone}</i>—Time zone code name of the locale; for example, gmt (this macro is not supported). ■ <i>{year}</i>—In the format YYYY; for example, 1970

- {year_abbrev}—From 00 through 99

Usage Guidelines See “Configuring File Formats” on page 778.

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 800
- password (Transfer Log File Servers) on page 800

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	<i>password</i> —FTP server password.
Usage Guidelines	See “Configuring Flow Collector Destinations” on page 777.
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	<i>password</i> —FTP server password.
Usage Guidelines	See “Configuring Transfer Logs” on page 779.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

retry

Syntax	<code>retry number;</code>
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	<i>number</i> —Maximum number of transfer retry attempts. Range: 0 through 10
Usage Guidelines	See “Configuring Retry Attempts” on page 779.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

retry-delay

Syntax	<code>retry-delay seconds;</code>
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	<i>seconds</i> —Amount of time between transfer retry attempts. Range: 0 through 60
Usage Guidelines	See “Configuring Retry Attempts” on page 779.
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 <i>number</i> ; timeout <i>seconds</i> ; }
Hierarchy Level	[edit services flow-collector file-specification variant <i>variant-number</i>]
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 <i>number</i> —Number of flow collection files collected. timeout <i>seconds</i> —Timeout duration.
Usage Guidelines	See “Configuring File Formats” on page 778.
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 " <i>password</i> "; username <i>username</i> ; } } filename-prefix <i>prefix</i> ; maximum-age <i>minutes</i> ; }
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 779.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

username

Syntax	<code>username <i>user-name</i>;</code>
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	<i>username</i> —FTP server username.
Usage Guidelines	See “Configuring Transfer Logs” on page 779.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

variant

Syntax	<pre>variant <i>variant-number</i> { data-format <i>format</i>; name-format <i>format</i>; transfer { record-level <i>number</i>; timeout <i>seconds</i>; } }</pre>
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 778.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

Chapter 41

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 805
- Configuring Dynamic Flow Capture Properties on page 806
- Example: Dynamic Flow Capture Configuration on page 813

Dynamic Flow Capture Architecture

The architecture consists of one or more *control sources* that send requests to a Juniper Networks routing platform 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:

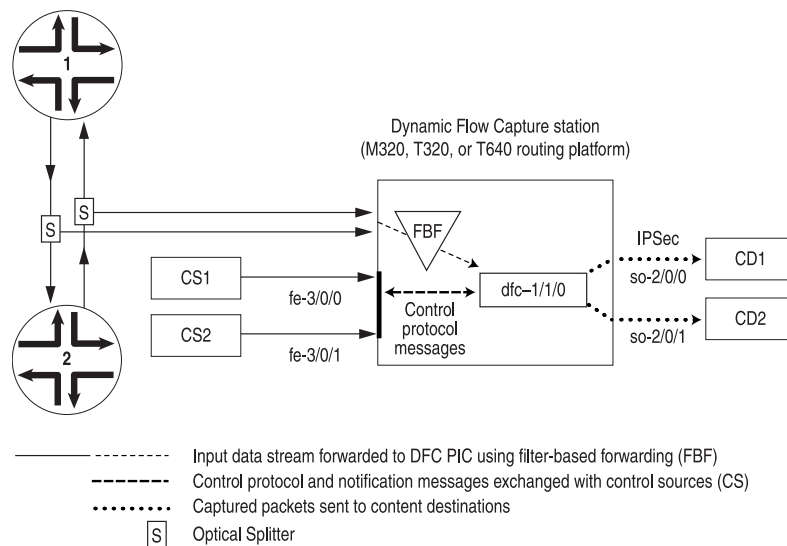
- Control source—A client that monitors electronic data or voice transfer over the network. The control source sends filter requests to the Juniper Networks routing platform using the Dynamic Task Control Protocol (DTCP), specified in draft-cavuto-dtcp-01.txt at <http://www.ietf.org/internet-drafts>. The control source is identified by a unique identifier and an optional list of IP addresses.
- Monitoring platform—A Juniper Networks T-series or M320 routing platform 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 203.



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 11 on page 806 shows a sample topology. The number of control sources and content destinations is arbitrary.

Figure 11: Dynamic Flow Capture Topology



g017075

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 Properties

To configure dynamic flow capture, include the `dynamic-flow-capture` statement at the `[edit services]` hierarchy level:

```
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 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-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 807
- Configuring the Content Destination on page 808
- Configuring the Control Source on page 809
- Configuring the DFC PIC Interface on page 810
- Configuring System Logging on page 811
- Configuring Thresholds on page 811
- Configuring Maximum Duplicates on page 812

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 [ destination ];
    minimum-priority value;
    no-syslog;
    notification-targets [ address 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;
}

```

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-identifier ];
  minimum-priority value;
  no-syslog;
  notification-targets [ address address port port-number ];
  service-port port-number;
  shared-key value;
  source-addresses [ address ];
}
```

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.

Configuring Maximum Duplicates

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 `g-max-duplicates`, `g-duplicates-dropped-periodicity` 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: Dynamic Flow Capture Configuration

The following example configurations show all the parts of a complete dynamic flow capture configuration.

Configuring 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;
  }
}
```

Configuring 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;
    }
  }
}
```

Configuring filter-based forwarding (FBF) to the DFC PIC interface, logical unit 1:

```
interfaces so-1/2/0 {
  encapsulation ppp;
  unit 0 {
    passive-monitor-mode;
    family inet {
```

```

        filter {
            input catch;
        }
    }
}

```

For more information about configuring passive monitoring interfaces, see “Enabling Passive Flow Monitoring” on page 711.

Configuring the firewall filter:

```

firewall {
    filter catch {
        interface-specific;
        term def {
            then {
                count counter;
                routing-instance fbf_inst;
            }
        }
    }
}

```

Configuring a forwarding routing instance:

```

routing-instances fbf_inst {
    instance-type forwarding;
    routing-options {
        static {
            route 0.0.0.0/0 next-hop dfc-0/0/0.1;
        }
    }
}

```

This next hop should point 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.

Configuring routing table groups:

```

routing-options {
    interface-routes {
        rib-group inet common;
    }
    rib-groups {
        common {
            import-rib [ inet.0 fbf_inst.inet.0 ];
        }
    }
    forwarding-table {
        export pplb;
    }
}

```

Configuring 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 42

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 routing platforms, except M120 routers, M160 routers, and TX Matrix platforms. 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.



NOTE: For information about dynamic flow capture, see “Dynamic Flow Capture Configuration Guidelines” on page 805. 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 818
- Configuring Flow-Tap Properties on page 819
- Example: Flow-Tap Configuration on page 821

Flow-Tap Architecture

The architecture consists of one or more *mediation devices* that send requests to a Juniper Networks routing platform 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 routing platform using the DTCP. The clients are not identified for security reasons, but have permissions defined by a set of special login classes.
- **Monitoring platform**—A Juniper Networks M-series or T-series routing platform 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 203.
- **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.

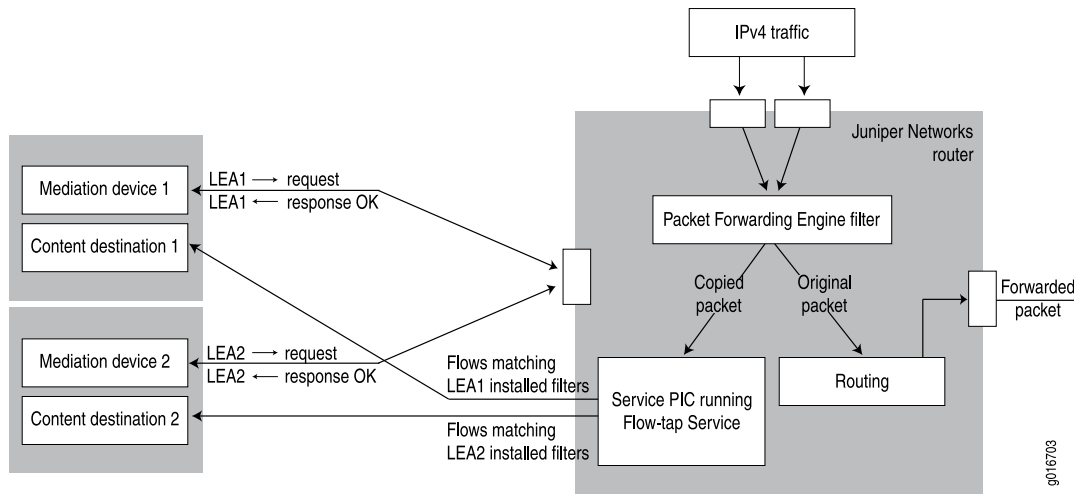
Following is a sample filter configuration; note that it is dynamically generated by the router (no user configuration required):

```
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 12 on page 819 shows a sample topology that uses two mediation devices and two content destinations.

Figure 12: Flow-Tap Topology



Configuring Flow-Tap Properties

This section describes the following tasks for configuring flow-tap service:

- Configuring the Flow-Tap Interface on page 819
- Configuring Security Properties on page 820
- Restrictions on page 820

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;
```

```
    }
}
```

Configuring Security Properties

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

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.

Example: Flow-Tap Configuration

The following example shows all the 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";
      }
    }
  }
}

```


Chapter 43

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	<code>address <i>address</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure an IP address for the flow capture destination.
Options	<i>address</i> —IP address for the content destination.
Usage Guidelines	See “Configuring the Content Destination” on page 808.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

allowed-destinations

Syntax	<code>allowed-destinations [<i>identifier</i>];</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]</code>
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	<i>identifier</i> —Allowed content destination name.
Usage Guidelines	See “Configuring the Control Source” on page 809.
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 [destination];
 minimum-priority value;
 no-syslog;
 notification-targets [address 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;
 }`

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 807.

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 808.

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 <i>identifier</i> { allowed-destinations [<i>destination</i>]; minimum-priority <i>value</i> ; no-syslog; notification-targets [address <i>address</i> port <i>port-number</i>]; service-port <i>port-number</i> ; shared-key <i>value</i> ; source-addresses [<i>address</i>]; }
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Identify the source of the dynamic flow capture request.
Options	<i>identifier</i> —Name of control source. The remaining statements are explained separately.
Usage Guidelines	See “Configuring the Control Source” on page 809.
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 <i>seconds</i> ;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
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	<i>seconds</i> —Period for sending DuplicatesDropped notifications. Default: 30 seconds
Usage Guidelines	See “Configuring Maximum Duplicates” on page 812.
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 [*destination*];
 minimum-priority *value*;
 no-syslog;
 notification-targets [address *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-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 805.

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 <i>interface-name</i> ; }
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 817.
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 <i>seconds</i> ;
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	<i>seconds</i> —Period for sending DuplicatesDropped notifications.
Usage Guidelines	See “Configuring Maximum Duplicates” on page 812.
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	<code>g-max-duplicates <i>number</i>;</code>
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 <code>max-duplicates</code> setting applied at the <code>capture-group</code> level overrides the global setting.
Default	If no value is configured, a default setting of 3 is used.
Options	<i>number</i> —Maximum number of content destinations. Range: 1 through 64
Usage Guidelines	See “Configuring Maximum Duplicates” on page 812.
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	<code>hard-limit <i>bandwidth</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]
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 <code>hard-limit-target</code> value.
Options	<i>bandwidth</i> —Hard limit threshold, in bits per second.
Usage Guidelines	See “Configuring the Content Destination” on page 808.
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	<code>hard-limit-target <i>bandwidth</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]</code>
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	<i>bandwidth</i> —Target value, in bits per second.
Usage Guidelines	See “Configuring the Content Destination” on page 808.
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	<code>input-packet-rate-threshold <i>rate</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Specify a packet rate threshold value that triggers a system log warning message.
Options	<i>rate</i> —Threshold value.
Usage Guidelines	See “Configuring Thresholds” on page 811.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface

Syntax	<code>interface sp-fpc/pic/port.logical-unit;</code>
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	<i>interface-name</i> —Name of the DFC interface.
Usage Guidelines	See “Configuring the Flow-Tap Interface” on page 819.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interfaces

Syntax	<code>interfaces interface-name;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
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	<i>interface-name</i> —Name of the DFC interface.
Usage Guidelines	See “Configuring the DFC PIC Interface” on page 810.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

max-duplicates

Syntax	<code>max-duplicates <i>number</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i>]</code>
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 <code>g-max-duplicates</code> setting.
Default	If no value is configured, a default setting of 3 is used.
Options	<i>number</i> —Maximum number of content destinations. Range: 1 through 64
Usage Guidelines	See “Configuring Maximum Duplicates” on page 812.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<code>g-max-duplicates</code>

minimum-priority

Syntax	<code>minimum-priority <i>value</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]</code>
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify the minimum priority for the control source.
Options	<i>value</i> —Minimum priority value; if not specified, defaults to 0. Range: 0 through 254
Usage Guidelines	See “Configuring the Control Source” on page 809.
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 <i>client-name</i> control-source <i>identifier</i>]
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 811.
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 <i>address</i> port <i>port-number</i>];
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
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 <i>address</i> —Allowed destination IP address. port <i>port-number</i> —Allowed destination UDP port number.
Usage Guidelines	See “Configuring the Control Source” on page 809.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pic-memory-threshold

Syntax	<code>pic-memory-threshold percentage <i>percentage</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i>]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Specify a PIC memory usage percentage that triggers a system log warning message.
Options	<code>percentage <i>percentage</i></code> —PIC memory threshold value.
Usage Guidelines	See “Configuring Thresholds” on page 811.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-port

Syntax	<code>service-port <i>port-number</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]</code>
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Identify the User Datagram Protocol (UDP) port number for control protocol requests.
Options	<code><i>port-number</i></code> —Port number for control protocol request messages.
Usage Guidelines	See “Configuring the Control Source” on page 809.
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 805 or “Flow-Tap Configuration Guidelines” on page 817.
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 <i>value</i> ;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Configure the authentication key value.
Options	<i>value</i> —Secret authentication value shared between a control source and destination.
Usage Guidelines	See “Configuring the Control Source” on page 809.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

soft-limit

Syntax	<code>soft-limit <i>bandwidth</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]</code>
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 <code>syslog</code> statement, a log message will also be generated.
Options	<i>bandwidth</i> —Soft limit threshold, in bits per second.
Usage Guidelines	See “Configuring the Content Destination” on page 808.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

soft-limit-clear

Syntax	<code>soft-limit-clear <i>bandwidth</i>;</code>
Hierarchy Level	<code>[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]</code>
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	<i>bandwidth</i> —Soft-limit clear threshold, in bits per second.
Usage Guidelines	See “Configuring the Content Destination” on page 808.
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 [<i>address</i>];
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
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 routing platform.
Options	<i>address</i> —Allowed IP source address.
Usage Guidelines	See “Configuring the Control Source” on page 809.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

tth

Syntax	tth <i>hops</i> ;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Time-to-live (TTL) value for the IP-IP header.
Options	<i>hops</i> —TTL value.
Usage Guidelines	See “Configuring the Content Destination” on page 808.
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 841
- Link and Multilink Services Configuration Guidelines on page 845
- Summary of Link Services Configuration Statements on page 885

Chapter 44

Link and Multilink Services Overview

The Multilink Protocol (MP) 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 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 311. For more information about voice services, see “Voice Services Configuration Guidelines” on page 383.



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 CoS components are fully supported and are handled normally on M-series and T-series routing platforms, 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 311.

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 45

Link and Multilink Services Configuration Guidelines

To configure multilink and link services logical interface properties, include the following statements:

```
(ml-fpc/pic/port | ls-fpc/pic/port) {  
  unit logical-unit-number {  
    dlc dlci-identifier;  
    drop-timeout milliseconds;  
    encapsulation type;  
    fragment-threshold bytes;  
    interleave-fragments;  
    minimum-links number;  
    mrru bytes;  
    multicast-dlc 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 interface properties, include the `mlfr-uni-nni-bundle-options` statement at the [edit interfaces *ls-fpc/pic/port:channel*] hierarchy level:

```
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 is organized as follows:

- Configuring Multilink and Link Services Logical Interface Properties on page 846
- Configuring Link Services Physical Interface Properties on page 859
- Multilink and Link Services Interface Structure on page 863
- Configuring CoS Components on Link Services PICs on page 866
- Examples: Configuring Multilink Interfaces on page 871
- Examples: Configuring Link Services Interfaces on page 875

Configuring Multilink and Link Services Logical Interface Properties

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 847.

You can configure the following multilink and link services logical interface properties:

- Default Settings for Multilink and Link Services Logical Interfaces on page 847
- Configuring a Link Services Point-to-Point DLCI on page 848
- Configuring a Link Services Multicast-Capable DLCI on page 848
- Configuring a Drop Timeout Period on page 849
- Configuring Logical Interface Encapsulation on page 850
- Configuring a Fragmentation Threshold on page 851
- Configuring Link Services Delay-Sensitive Packet Interleaving on page 851
- Configuring Minimum Links on page 854
- Configuring the MRRU and MTU Values on page 855
- Configuring the Sequence Format on page 856
- Configuring Compressed RTP with MLPPP Encapsulation on page 856
- Configuring Compressed RTP with PPP Encapsulation on page 857

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 Bundles” on page 864.



NOTE: On DS0, E1, or T1 interfaces in LSQ bundles, you can configure the **bandwidth** statement, but the routing platform 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*.

Default Settings for Multilink and Link Services Logical Interfaces

Table 20 on page 847 lists the default settings for multilink and link services statements, together with the other permitted values or value ranges.

Table 20: 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 (Nx64)
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 21 on page 859 for statements that apply to link services physical interfaces only.

Configuring a Link Services Point-to-Point DLCI

For link services interfaces only, you can configure multiple 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.

DLCIs are not supported on multilink interfaces.

Configuring a Link Services Multicast-Capable DLCI

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 Link Services Keepalive Settings on Frame Relay LMI” on page 862.

Configuring a Drop Timeout Period

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 the Link Services Differential Delay” on page 861.

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.

Configuring Logical Interface Encapsulation

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 the Link Services Physical Interface Encapsulation” on page 860.

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.

Configuring a Fragmentation Threshold

By default, the fragmentation threshold parameter is disabled. For interfaces with MLPPP encapsulation only, you can configure a fragmentation threshold to set a maximum size for 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.



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 ($N \times 64$). The default value, 0, results in no fragmentation.

Configuring Link Services Delay-Sensitive Packet Interleaving

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 Fragmentation by Forwarding Class” on page 328.

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 Components on Link Services PICs” on page 866.

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 311.

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 Minimum Links

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 the MRRU and MTU Values

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 Format

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 Compressed RTP with MLPPP Encapsulation

On J-series Services Routers link services interfaces (**ls**), you can configure the compressed Real-Time Transport Protocol (CRTP) with MLPPP encapsulation.



NOTE: Optionally, you can configure CRTP with PPP encapsulation on a J-series Services Router. For more information, see “Configuring Compressed RTP with PPP Encapsulation” on page 857.

You can also configure CRTP with MLPPP on link services IQ interfaces (**lsq**) and voice services interfaces. When you configure CRTP with MLPPP, the configuration syntax for J-series Services Routers is the same as the syntax you use on the AS or MultiServices PIC. The only difference is that you use the **ls-** interface descriptor for J-series Services Routers, and you use the **lsq-** interface descriptor for the AS or MultiServices PIC. For more information, see “Configuring Voice Services Properties” on page 385.

All CRTP traffic is sent on the LFI queue (**q2**). We recommend that the incoming voice traffic be classified to the appropriate queue.

Example: Configuring Compressed RTP with MLPPP Encapsulation

Configure compression on a J-series T1 interface with MLPPP encapsulation. Configure fragmentation for all IP packets larger than 128 bytes.

For detailed information about these configuration statements, see “Configuring Voice Services Properties” on page 385.

```
[edit interfaces]
t1-1/0/0 {
  unit 0 {
    family mlppp {
      bundle ls-1/1/0.1;
    }
  }
}
ls-1/1/0 {
  unit 1 {
    compression {
      rtp {
        port minimum 16384 maximum 32768;
      }
    }
    encapsulation multilink-ppp;
    family inet {
      address 30.1.1.2/24;
    }
    fragment-threshold 128;
  }
}
```

Configuring Compressed RTP with PPP Encapsulation

On J-series Services Routers link services interfaces (**ls**) only, you can configure the compressed Real-Time Transport Protocol (RTP) with PPP encapsulation.



NOTE: Optionally, you can configure CRTP with MLPPP on a J-series Services Router. For more information, see “Configuring Compressed RTP with MLPPP Encapsulation” on page 856.

The configuration syntax for CRTP with PPP encapsulation is much like the syntax you use when you configure CRTP with MLPPP.

To configure CRTP with PPP encapsulation, perform the following steps:

1. On the logical T1 or E1 constituent link, configure the link services interface used for compression by including the **compression-device** statement:

```
compression-device interface-name;
```

The interface name is the link services interface on which the T1 or E1 constituent link is bundled.

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*]
2. Specify how the link services interface handles voice traffic compression by including the following statements:

```
compression {
  rtp {
    f-max-period number;
    queues [ queue-numbers ];
    port {
      minimum port-number;
      maximum port-number;
    }
  }
}
```

You can include these statements 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 detailed information about these statements, see “Configuring Voice Services Properties” on page 385.

Example: Configuring Compressed RTP with PPP Encapsulation

Configure compression on a J-series T1 interface with PPP encapsulation. Configure fragmentation for all IP packets larger than 128 bytes.

```
[edit interfaces]
t1-1/0/0 {
  encapsulation ppp;
  unit 0 {
    compression-device ls-0/0/0.1
  }
}
```

```

ls-0/0/0.1 {
  compression {
    rtp {
      port minimum 16384 maximum 32768;
    }
  }
  family inet {
    address 10.0.2.3;
    # Other IPv4 parameters here.
  }
}

```

Configuring Link Services Physical Interface Properties

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 the section “Default Settings for Link Services Interfaces” on page 859.

You can configure the following link services physical interface properties:

- Default Settings for Link Services Interfaces on page 859
- Configuring the Link Services Physical Interface Encapsulation on page 860
- Configuring Link Services Acknowledgment Timers on page 860
- Configuring the Link Services Differential Delay on page 861
- Configuring Link Services Keepalive Settings on Frame Relay LMI on page 862

For descriptions of link services physical interface properties that also can be configured at the logical unit level, see “Configuring Multilink and Link Services Logical Interface Properties” on page 846.

Default Settings for Link Services Interfaces

Table 21 on page 859 lists the default settings for link services statements, together with the other permitted values or value ranges.

Table 21: 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 (Nx64)

Table 21: Link Services Physical Interface Statements for MLFR FRF.16 (continued)

Option	Default Value	Possible Values
LMI type	itu	ansi, itu
Minimum links	1 link	1 through 8 links
MRRU	1504 bytes	1500 through 4500 bytes
n391 (full status polling counter)	6	1 through 255
n392 (LMI error threshold)	3	1 through 10
n393 (LMI monitored event count)	4	1 through 10
t391 (link integrity verify polling timer)	10	5 through 30
t392 (polling verification timer)	15	5 through 30
Sequence ID format for MLFR	12 bits	12 bits

Configuring the Link Services Physical Interface Encapsulation

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 Logical Interface Encapsulation” on page 850.

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 Link Services Acknowledgment Timers

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 the Link Services Differential Delay

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 end point 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 ls-qfpc/pic/port:channel` command.

Configuring Link Services Keepalive Settings on Frame Relay LMI

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;
```

You can set the following properties:

- `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 `t391`. `n391` specifies the frequency with which these inquiries expect a full status report; for example, an `n391` value of 10 would specify a full status report in response to every tenth inquiry. The intermediate inquiries ask for a keepalive exchange only. The range is 1 through 255, with a default value of 6.
- `n392`—Error threshold. The number of errors required to bring down the link, within the event count specified by `n393`. The range is from 1 through 10, with a default value of 3.
- `n393`—Monitored event count. The range is from 1 through 10, with a default value of 4.

- t391—Keepalive timer. Interval at which the DTE sends out a keepalive response request to the DCE and updates status, depending on the error threshold value. The range is from 5 through 30 seconds, with a default value of 10 seconds.
- t392—Keepalive timer. 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 value of 15 seconds.



NOTE: For the LMI to work properly, you must configure one side of a link services bundle to be a DCE.

Multilink and Link Services Interface Structure

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 J-series and M-series platforms, only one DS3 link is allowed in an MLFR bundle. MLPPP bundles can include two DS3 links.

This section is organized as follows:

- Multilink Services and Link Services PIC Capacities on page 863
- Link Services PIC Capabilities on page 864
- Configuring Bundles on page 864

Multilink Services and Link Services PIC Capacities

Three versions of Multilink Services and three versions of Link Services PICs are available, as shown in Table 22 on page 864. 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 22: 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.

Link Services PIC Capabilities

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 “Configuring a Link Services Interface with Two Links” on page 875.

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 Services Interfaces” on page 875.




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 Bundles

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 13 on page 865).

The physical interface is usually connected to networks capable of supporting MLPPP or MLFR (FRF.15 or FRF.16).

Figure 13: Multilink Interface Configuration



Using the topology in Figure 13 on page 865 as an example, configure a multilink or link services bundle over a T1 connection (for which you have already configured the T1 physical interface) with the following additional configuration statements:

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.

Configuring CoS Components on Link Services PICs

For link services IQ interfaces (`lsq`), JUNOS CoS components are fully supported and are handled normally, 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 311.

For Link Services PIC interfaces (`ls`), CoS works differently for a Link Services PIM installed on a J-series Services Router than it does for a Link Services PIC installed on an M-series or T-series platform, as described in the following sections:

- Link Services CoS on J-series Services Routers on page 866
- Link Services CoS on M-series and T-series Platforms on page 866
- Example: Configuring Link Services CoS Components on page 868

Link Services CoS on J-series Services Routers

Unlike M-series and T-series platforms, J-series Services Routers support per-bundle queuing on link services (`ls`) interfaces. Link services interfaces for J-series Services Routers behave the same way as link services IQ interfaces (`lsq`). (For more information, see “Link Services IQ Interfaces Configuration Guidelines” on page 311.) There are some exceptions, as follows:

- Queue 2 is reserved for voice traffic (LFI) on J-series Services Routers, while all other queues perform fragmentation.
- For FRF.15 on J-series Services Routers, the constituent links bundled on the link services interfaces (`ls`) require a single scheduler with 25 percent transmission rates and buffer sizes for queues 0 through 3. You should assign this scheduler to each constituent link—in other words, to E1 interfaces and T1 interfaces. However, you can configure customized scheduler maps for each associated `ls-fpc/pic/port.logical` interface.

For FRF.16 and MLPPP, link services interfaces (`ls`) on the J-series Services Router work the same way as link services IQ interfaces (`lsq`) on M-series platforms.

Link Services CoS on M-series and T-series Platforms

For Link Services PIC interfaces (`ls`) on M-series and T-series platforms, queue 0 is the only queue that you should configure to receive fragmented packets. Configure all other queues to be higher-priority queues.

Table 23 on page 867 summarizes how CoS queues work on link services (`ls`) interfaces.

Table 23: Link Services CoS Queues

Supported Bundling Type	Queue 0	Higher-Priority Queues
Hash-based load balancing	No	Yes
MLFR FRF.15	Yes	No
MLFR FRF.16	Yes	No
MLPPP	Yes	No

For M-series and T-series platforms, 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 Link Services Delay-Sensitive Packet Interleaving” on page 851 and the *JUNOS Policy Framework Configuration Guide*.

Example: Configuring Link Services CoS Components

Configure CoS on a link services interface and its constituent link interfaces.



NOTE: This example applies to M-series and T-series platforms. For examples that apply to a Link Services PIC installed on a J-series Services Router, see “Link Services IQ Interfaces Configuration Guidelines” on page 311.

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

These examples show only the multilink part of the configuration. To see the T1 configuration options, see the *JUNOS Network Interfaces Configuration Guide*.

The examples in this section show the following configurations:

- Configuring an MLPPP Interface on page 871
- Configuring an MLPPP over ATM 2 Interface on page 872
- Configuring an MLFR FRF.15 Interface on page 874

Configuring an MLPPP Interface

```

[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;
        }
    }
}
}

```

Configuring an MLPPP over ATM 2 Interface

```

[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;
    }
  }
}
...
}

```

Configuring an MLFR FRF.15 Interface

```

[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 Services Interfaces

This example shows only the link services part of the configuration. To see the T1 configuration options, see the *JUNOS Network Interfaces Configuration Guide*.

The examples in this section show the following configurations:

- Configuring a Link Services Interface with Two Links on page 875
- Configuring a Link Services interface with MLPPP on page 876
- Configuring a Link Services PIC with MFR FRF.15 on page 877
- Configuring a Link Services PIC with MLFR FRF.16 on page 878
- Configuring a Link Services PIC and Voice Services Interface with a Combination of Bundle Types on page 878

This configuration initiates the MLFR UNI NNI protocol between Router A and Router B and logically connects link services bundles ls-1/1/0.3 and ls-0/0/0.10, as specified in Table 24 on page 875.

Table 24: Link Services Bundle

Router A	Router B
t1-0/1/0 (ls-1/1/0:3)	t1-0/3/0 (ls-0/0/0:10)
t1-0/1/1 (ls-1/1/0:3)	t1-0/3/1 (ls-0/0/0:10)

For LMI to work properly, you must configure one router to be a DCE.

Configuring a Link Services Interface with Two Links

On Router A:

```
[edit interfaces]
ls-1/1/0:3 {
  dce;
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    dlc1 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;
    }
  }
}

```

On Router B:

```

[edit interfaces]
ls-0/0/0:10 {
  encapsulation multilink-frame-relay-uni-nni;
  unit 0 {
    dci 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;
    }
  }
}

```

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;
    }
  }
}
}

```

**Configuring a Link
Services PIC with MFR
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;
    }
  }
}

```

**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;
      }
    }
  }
}

```

**Configuring a Link
Services PIC and Voice
Services Interface 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 46

Summary of Link Services Configuration Statements

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

acknowledge-retries

Syntax	<code>acknowledge-retries <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> 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	<i>number</i> —Number of retransmission attempts to be made following the expiration of the acknowledgment timer. Range: 1 through 5 Default: 2
Usage Guidelines	See “Configuring Link Services Acknowledgment Timers” on page 860.
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 <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> 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	<i>milliseconds</i> —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 Link Services Acknowledgment Timers” on page 860.
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 <i>ls-fpc/pic/port:channel</i> 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 the Link Services Differential Delay” on page 861.
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 <i>address</i> { destination <i>address</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interface address.
Options	<i>address</i> —Address of the interface. The remaining statements are explained separately.
Usage Guidelines	See “Configuring Multilink and Link Services Logical Interface Properties” on page 846; for a general discussion of address statement options, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

bundle

Syntax	bundle (<i>ml-fpc/pic/port</i> <i>ls-fpc/pic/port</i>);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mlfr-end-to-end], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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	<i>ml-fpc/pic/port</i> —Name of the multilink interface you are linking. <i>ls-fpc/pic/port</i> —Name of the link services interface you are linking.
Usage Guidelines	See “Configuring Bundles” on page 864.
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 <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J-series Services Routers with PPP over Ethernet interfaces, configure the name of the access concentrator (AC).
Options	<i>name</i> —Name of the AC.
Usage Guidelines	See “Configuring Compressed RTP with PPP Encapsulation” on page 857.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

destination

Syntax	destination <i>destination-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
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	<i>destination-address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring Multilink and Link Services Logical Interface Properties” on page 846.
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 <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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 a Drop Timeout Period” on page 849.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dlci

Syntax	dlci <i>dlci-identifier</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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 <i>multipoint-destination</i> statement to specify the DLCI.
Options	<i>dlci-identifier</i> —Data-link connection identifier. Range: 16 through 1022
Usage Guidelines	See “Configuring a Link Services Multicast-Capable DLCI” on page 848; for general information about Frame Relay DLCIs, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

drop-timeout

Syntax	<code>drop-timeout <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit logical-systems <i>logical-system-name</i> interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>]
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	<i>milliseconds</i> —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 a Drop Timeout Period” on page 849.
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 891
- encapsulation (Physical Interface) on page 892

encapsulation (Logical Interface)

Syntax	encapsulation (atm-mlppp-llc multilink-frame-relay-end-to-end multilink-ppp ...);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Logical link-layer encapsulation type.
Options	<p>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, <i>PPP over AAL5</i>.</p> <p>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.</p> <p>multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink and link services interfaces and their constituent T1 or E1 interfaces.</p>
Usage Guidelines	See “Configuring Logical Interface Encapsulation” on page 850; for information about encapsulation statement options used with other interface types, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

encapsulation (Physical Interface)

Syntax	encapsulation (multilink-frame-relay-uni-nni ...);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
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 the Link Services Physical Interface Encapsulation” on page 860; for information about encapsulation statement options used with other interface types, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

family

Syntax	<pre>family family { address address { destination address; } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure protocol family information for the logical interface.
Options	<p><i>family</i>—Protocol family:</p> <ul style="list-style-type: none"> ■ <i>ccc</i>—Circuit cross-connect protocol suite ■ <i>inet</i>—IP version 4 (IPv4) ■ <i>inet6</i>—IP version 6 (IPv6) ■ <i>iso</i>—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite ■ <i>mlfr-end-to-end</i>—Multilink Frame Relay FRF.15 ■ <i>mlfr-uni-nni</i>—Multilink Frame Relay FRF.16 ■ <i>multilink-ppp</i>—Multilink Point-to-Point Protocol ■ <i>mpls</i>—Multiprotocol Label Switching ■ <i>tcc</i>—Translational cross-connect protocol suite ■ <i>tnp</i>—Trivial Network Protocol ■ <i>vpls</i>—Virtual private LAN service <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Link and Multilink Services Configuration Guidelines” on page 845; for a general discussion of <i>family</i> statement options, see the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

fragment-threshold

Syntax	fragment-threshold <i>bytes</i> ;
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit logical-systems <i>logical-system-name</i> interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For multilink and link services interfaces only, set the fragmentation threshold, in bytes.
Options	<i>bytes</i> —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 “Configuring a Fragmentation Threshold” on page 851.
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 <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> 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	<i>milliseconds</i> —The rate at which hello messages are sent. Range: 1 through 180 milliseconds Default: 10 milliseconds
Usage Guidelines	See “Configuring Link Services Acknowledgment Timers” on page 860.
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 <i>JUNOS Network Interfaces Configuration Guide</i> .
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/portunit 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 Link Services Delay-Sensitive Packet Interleaving” on page 851.
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 <i>ls-fpc/pic/port:channel</i> 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 Link Services Keepalive Settings on Frame Relay LMI” on page 862.
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 <i>number</i> ;
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit logical-systems <i>logical-system-name</i> interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>]
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	<i>number</i> —Number of links. Range: 1 through 8 Default: 1
Usage Guidelines	See “Configuring Minimum Links” on page 854.
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);
 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 the Link Services Physical Interface Encapsulation” on page 860.

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

mrru

Syntax	<code>mrru bytes;</code>
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit interfaces (ml-fpc/pic/port <i>ls-fpc/pic/port</i>) unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>ls-fpc/pic/port:channel</i> mlfr-uni-nni-bundle-options], [edit logical-systems <i>logical-system-name</i> interfaces (ml-fpc/pic/port <i>ls-fpc/pic/port</i>) unit <i>logical-unit-number</i>]
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	<i>bytes</i> —MRRU size. Range: 1500 through 4500 bytes Default: 1500 bytes
Usage Guidelines	See “Configuring the MRRU and MTU Values” on page 855.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mtu

Syntax	<code>mtu bytes;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values.
Options	<i>bytes</i> —MTU size. Range: 0 through 5012 bytes Default: 1500 bytes (inet, inet6, and iso families), 1448 bytes (mpls)
Usage Guidelines	See “Configuring the MRRU and MTU Values” on page 855.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i>

multicast-dlci

Syntax	<code>multicast-dlci <i>dlci-identifier</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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	<i>dlci-identifier</i> —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 a Link Services Multicast-Capable DLCI” on page 848.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

n391

Syntax	<code>n391 <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> 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	<i>number</i> —Polling interval. Range: 1 through 255 Default: 6
Usage Guidelines	See “Configuring Link Services Keepalive Settings on Frame Relay LMI” on page 862.
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	<code>n392 number;</code>
Hierarchy Level	<code>[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]</code>
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	<i>number</i> —Error threshold. Range: 1 through 10 Default: 3
Usage Guidelines	See “Configuring Link Services Keepalive Settings on Frame Relay LMI” on page 862.
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	<code>n393 number;</code>
Hierarchy Level	<code>[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services interfaces only, set the Frame Relay monitored event count.
Options	<i>number</i> —Event count. Range: 1 through 255 Default: 6
Usage Guidelines	See “Configuring Link Services Keepalive Settings on Frame Relay LMI” on page 862.
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	<code>red-differential-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> 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	<i>milliseconds</i> —Red differential delay threshold. Range: 1 through 2000 milliseconds Default: 120 milliseconds
Usage Guidelines	See “Configuring the Link Services Differential Delay” on page 861.
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	<code>short-sequence;</code>
Hierarchy Level	[edit interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces (<i>ls-fpc/pic/port</i> <i>ml-fpc/pic/port</i>) unit <i>logical-unit-number</i>]
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 Format” on page 856.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

t391

Syntax	<code>t391 number;</code>
Hierarchy Level	<code>[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services interfaces only, set the Frame Relay link integrity polling interval.
Options	<i>number</i> —Link integrity polling interval. Range: 5 through 30 seconds Default: 10 seconds
Usage Guidelines	See “Configuring Link Services Keepalive Settings on Frame Relay LMI” on page 862.
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	<code>t392 number;</code>
Hierarchy Level	<code>[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services interfaces only, set the Frame Relay polling verification interval.
Options	<i>number</i> —Polling verification interval. Range: 5 through 30 seconds Default: 15 seconds
Usage Guidelines	See “Configuring Link Services Keepalive Settings on Frame Relay LMI” on page 862.
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 845; 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	<code>yellow-differential-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>ls-fpc/pic/port:channel</i> 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	<i>milliseconds</i> —Yellow differential delay threshold. Range: 1 through 2000 milliseconds Default: 72 milliseconds
Usage Guidelines	See “Configuring the Link Services Differential Delay” on page 861.
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 907
- Real-Time Performance Monitoring Configuration Guidelines on page 909
- Summary of Real-Time Performance Monitoring Configuration Statements on page 927

Chapter 47

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 Border Gateway Protocol (BGP) neighbors. You can view the results of the discovery using a Simple Network Management Protocol (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 48

Real-Time Performance Monitoring Configuration Guidelines

To configure Real-Time Performance Monitoring (RPM) services, include the `rpm` statement at the `[edit services]` hierarchy level:

```
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 {
        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;
    }
}
}

```

This chapter includes the following sections:

- Configuring BGP Neighbor Discovery Through RPM on page 910
- Configuring Real-Time Performance Monitoring Properties on page 912
- Examples: Configuring BGP Neighbor Discovery Through RPM on page 922
- Examples: Configuring Real-Time Performance Monitoring on page 923

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 Properties

This section describes the following tasks for configuring RPM properties:

- Configuring the Probe on page 912
- Configuring the Server on page 917
- Configuring the Maximum Number of Probes on page 917
- Configuring RPM Timestamping on page 917
- Configuring TWAMP on page 920



NOTE: RPM does not require an Adaptive Services (AS) or MultiServices Physical Interface Card (PIC), unless you are configuring RPM timestamping.

Configuring the Probe

The `probe` statement defines the owner name and test name used as identifiers for the probe. Together, they 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 917), 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.

- To specify the output interface for sending packets to the forwarding plane, include the **destination-interface** statement at the [edit services rpm probe owner test *test-name*] hierarchy level. This interface must support looping of probe packets to an input interface without adding any encapsulation by using a pair of logical tunnel interfaces.



NOTE: This usage of the **destination-interface** statement is supported only on J-series Services Routers, for which you should specify an It tunnel interface.

The **destination-interface** statement is required to support classification of RPM probes generated by JUNOS software that are based on a user-configured classifier. The classified packets are sent to the output queue on the output interface specified by the scheduler map configured on that interface. For more

information about classifiers and scheduler maps, see the *JUNOS Class of Service Configuration Guide*. For a complete configuration example, see “Examples: Configuring Real-Time Performance Monitoring” on page 923.



CAUTION: Use this feature with caution, because improper configuration can cause packets to be dropped.

On M-series and T-series routing platforms only, 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 917. On J-series Services Routers, you configure hardware timestamping by including the **hardware-timestamp** statement. On all platforms, 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, a Simple Network Management Protocol (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 the Server

The TCP and UDP probes are proprietary to Juniper Networks and require a receiver to receive the probes. To configure the server to receive the probes, include the `probe-server` statement at the `[edit services rpm]` hierarchy level:

```
probe-server {
  tcp {
    destination-interface interface-name;
    port (RPM) number;
  }
  udp {
    destination-interface interface-name;
    port (RPM) number;
  }
}
```

The port number specified for the UDP and TCP server can be 7 or from 49160 through 65535. The `destination-interface` statement specifies the output interface for the RPM server. This interface should be able to support looping of packets to an input interface without adding any encapsulation.



NOTE: This usage of the `destination-interface` statement is supported only on J-series Services Routers, for which you should specify an `lt` tunnel interface.

The `destination-interface` statement is required to support classification of RPM probes generated by JUNOS software based on a user-configured classifier. The classified packets are sent to the output queue on the output interface specified by the scheduler map configured on that interface. For more information about classifiers and scheduler maps, see the *JUNOS Class of Service Configuration Guide*. For a complete configuration example, see “Examples: Configuring Real-Time Performance Monitoring” on page 923.

Configuring the Maximum Number of 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`. The configuration is handled in different ways on M-series and T-series routing platforms and on J-series Services Routers, as described in the following sections:

- Timestamps on M-series and T-series Routing Platforms on page 918
- Timestamps on J-series Services Routers on page 920

Timestamps on M-series and T-series Routing Platforms

If your router network includes an Adaptive Services (AS) or MultiServices PIC, you can enable timestamping of RPM probe messages. The timestamp is applied on the RPM client router (the router that originates the RPM probes). The functionality is not supported on the RPM probe server and applies only to IPv4 traffic. It is supported in the Layer 2 service package on all MultiServices PICs and in the Layer 3 service package on AS and MultiServices PICs.

To configure two-way timestamping on M-series and T-series routing platforms, 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 (**unit 0**) 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.



NOTE: If you configure RPM time stamping on an AS PIC, you cannot configure the **source-address** statement at the `[edit services rpm probe probe-name test test-name]` hierarchy level.

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 0 {
    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*.

Timestamps on J-series Services Routers

On J-series Services Routers, timestamps are applied in the forwarding process (FWDD-RT).

- To configure two-way timestamping, include the **hardware-timestamp** statement at the [edit services rpm probe *probe-owner* test *test-name*] hierarchy level:

```
hardware-timestamp;
```

- To configure one-way timestamping, include both the **hardware-timestamp** and the **one-way-hardware-timestamp** statements at the [edit services rpm probe *probe-owner* test *test-name*] hierarchy level:

```
hardware-timestamp;
one-way-hardware-timestamp;
```

For more information, see the *J-series Services Router Administration Guide*.



NOTE: You cannot include the **source-address** and **hardware-timestamp** or **one-way-hardware-timestamp** statements at the [edit services rpm probe *probe-name* test *test-name*] hierarchy level simultaneously.

Configuring TWAMP

You can configure support for the Two-Way Active Measurement Protocol (TWAMP) on MultiServices PICs running in Layer 2 mode. You can enable TWAMP on all MultiServices PIC models and on all router platforms that support these PICs.

For more information on TWAMP, see the current version of Internet draft *draft-ietf-ippm-twamp.txt*.

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 ];
    }
    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 the TWAMP Interface on page 921
- Configuring the TWAMP Server Properties on page 921

Configuring the TWAMP Interface

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;
```

Configuring the TWAMP Server Properties

You can specify a number of properties, some of which are optional, by including the `server` statement at the `[edit services rpm twamp]` hierarchy level:

```
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.
- To specify the authentication mode, include the `authentication-mode` statement at the `[edit services rpm twamp server]` hierarchy level. By default, the value is `none`, which is the only value currently accepted.
- 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. 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. 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 923.

Examples: Configuring BGP Neighbor Discovery Through RPM

The following example configures 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;
}
```

The following example configures 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;
}
}

```

The following example configures 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

The following example configures 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;

```

The following example shows a typical configuration for classification, including the use of 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;
    }
}
}

```

The following example shows configuration of 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;
    }
  }
}

```

The following example shows the minimum configuration necessary to enable TWAMP:

```

[edit services]
rpm {
  twamp {
    server {
      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.
    }
  }
}

```

The following example shows 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 49

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.

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 920.

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

data-fill

Syntax	<code>data-fill data;</code>
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	<i>data</i> —A hexadecimal value; for example, 0-9, A-F.
Usage Guidelines	The <code>data-fill</code> statement is not valid with the <code>http-get</code> or <code>http-metadata-get</code> probe types. See “Configuring BGP Neighbor Discovery Through RPM” on page 910 or “Configuring the Probe” on page 912.
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.

data-size

Syntax	<code>data-size size;</code>
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	<i>data</i> —The size can be from 0 through 65507 Default: 0
<hr/> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1; text-align: center; margin-right: 10px;">  </div> <div> <p>NOTE: If you configure the hardware timestamp feature (see “Configuring RPM Timestamping” on page 917), 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.</p> </div> </div> <hr/>	
Usage Guidelines	The <code>data-size</code> statement is not valid with the <code>http-get</code> or <code>http-metadata-get</code> probe type. See “Configuring BGP Neighbor Discovery Through RPM” on page 910 or “Configuring the Probe” on page 912.
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-interface

Syntax	<code>destination-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit services rpm probe <i>owner</i> test <i>test-name</i>], [edit services rpm probe-server (tcp udp)]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>On J-series Services Routers only, specify the output (lt-) interface used to send packets to the forwarding plane. This interface must support looping of packets to an input interface without adding any encapsulation.</p> <p>On M-series and T-series routing platforms, specify a services (sp-) interface that adds a timestamp to RPM probe messages. This feature is supported only with <code>icmp-ping</code>, <code>icmp-ping-timestamp</code>, <code>udp-ping</code>, and <code>udp-ping-timestamp</code> probe types. You must also configure the <code>rpm</code> statement on the <code>sp-</code> interface and include the <code>unit 0 family inet</code> statement with a <code>/32</code> address.</p>
Options	<i>interface-name</i> —Name of the output interface (J-series Services Routers) or adaptive services interface (M-series and T-series routers).
Usage Guidelines	See “Configuring the Probe” on page 912, “Configuring the Server” on page 917, or “Configuring RPM Timestamping” on page 917.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	hardware-timestamp, rpm

destination-port

Syntax	<code>destination-port <i>port</i>;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe <i>owner</i> test <i>test-name</i>]
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	<i>port</i> —The port number can be 7 or from 49,160 to 65,535.
Usage Guidelines	See “Configuring BGP Neighbor Discovery Through RPM” on page 910 or “Configuring the Probe” on page 912.
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dscp-code-point

Syntax	<code>dscp-code-point <i>dscp-bits</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code>
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	<p><i>dscp-bits</i>—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:</p> <ul style="list-style-type: none"> ■ 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 the Probe” on page 912.

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.

Description On J-series Services Routers only, enable timestamping of RPM probe messages. 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 917.

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 910 or “Configuring the Probe” on page 912.

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

inactivity-timeout

Syntax	<code>inactivity-timeout seconds;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Inactivity timeout period, in seconds.
Options	<i>seconds</i> —Length of time the session is inactive before it times out. Default: 1800 seconds
Usage Guidelines	See “Configuring TWAMP” on page 920.
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.

logical-system

Syntax	<code>logical-system <i>logical-system-name</i> { [routing-instances <i>instance-name</i>]; }</code>
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	<i>logical-system-name</i> —Logical system name.
Usage Guidelines	See “Configuring BGP Neighbor Discovery Through RPM” on page 910.
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 <i>count</i> ;
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	<i>count</i> —Maximum number of connections. Range: 1 through 2048
Usage Guidelines	See “Configuring TWAMP” on page 920.
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 <i>count</i> ;
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	<i>count</i> —Maximum number of connections.
Usage Guidelines	See “Configuring TWAMP” on page 920.
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-sessions

Syntax	<code>maximum-sessions count;</code>
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	<i>count</i> —Maximum number of sessions. Range: 1 through 2048
Usage Guidelines	See “Configuring TWAMP” on page 920.
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	<code>maximum-sessions-per-connection count;</code>
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	<i>count</i> —Maximum number of sessions.
Usage Guidelines	See “Configuring TWAMP” on page 920.
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.

moving-average-size

Syntax	<code>moving-average-size <i>number</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code>
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	<i>number</i> —Number of samples to be used in calculations. Range: 0 through 255
Usage Guidelines	See “Configuring the Probe” on page 912.
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	<code>one-way-hardware-timestamp;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code>
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 <code>destination-interface</code> statement (M-series and T-series platforms) or the <code>hardware-timestamp</code> statement (J-series Services Routers) to invoke timestamping. This feature is supported only with <code>icmp-ping</code> , <code>icmp-ping-timestamp</code> , <code>udp-ping</code> , and <code>udp-ping-timestamp</code> probe types.
Usage Guidelines	See “Timestamps on J-series Services Routers” on page 920.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<code>destination-interface</code> , <code>hardware-timestamp</code>

port

- port (RPM) on page 937
- port (TWAMP) on page 937

port (RPM)

Syntax	port <i>number</i> ;
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	<i>number</i> —Port number for the probe server. The value can be 7 or 49,160 through 65,535.
Usage Guidelines	See “Configuring the Server” on page 917.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

port (TWAMP)

Syntax	port <i>number</i> ;
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	<i>number</i> —Port number. Range: 1 through 65,535
Usage Guidelines	See “Configuring TWAMP” on page 920.
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 the Probe” on page 912.

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

probe-count

Syntax	<code>probe-count <i>count</i>;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit services rpm probe owner test <i>test-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the number of probes within a test.
Options	<i>count</i> —A value from 1 through 15.
Usage Guidelines	See “Configuring BGP Neighbor Discovery Through RPM” on page 910 or “Configuring the Probe” on page 912.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

probe-interval

Syntax	<code>probe-interval <i>interval</i>;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit services rpm probe owner test <i>test-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the time to wait between sending packets, in seconds.
Options	<i>interval</i> —Number of seconds, from 1 through 255.
Usage Guidelines	See “Configuring BGP Neighbor Discovery Through RPM” on page 910 or “Configuring the Probe” on page 912.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

probe-limit

Syntax	<code>probe-limit <i>limit</i>;</code>
Hierarchy Level	[edit services rpm]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the maximum number of concurrent probes allowed.
Options	<i>limit</i> —A value from 1 through 500. Default: 100.
Usage Guidelines	See “Configuring the Maximum Number of Probes” on page 917.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

probe-server

Syntax	<pre> probe-server { tcp { destination-interface <i>interface-name</i>; port (RPM) <i>number</i>; } udp { destination-interface <i>interface-name</i>; port (RPM) <i>number</i>; } } </pre>
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 the Server” on page 917.
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 <i>type</i> ;
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test <i>test-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the packet and protocol contents of a probe.
Options	<p><i>type</i>—Specify one of the following probe type values:</p> <ul style="list-style-type: none"> ■ 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 910 or “Configuring the Probe” on page 912.
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 <i>instance-name</i> ;
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the routing instance used by the probes.
Options	<p><i>instance-name</i>—A routing instance configured at the [edit routing-instance] hierarchy level.</p> <p>Default: Internet routing table inet.0.</p>
Usage Guidelines	See “Configuring the Probe” on page 912.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

routing-instances

Syntax	<code>routing-instances <i>instance-name</i>;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm bgp logical-system <i>logical-system-name</i>]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	Specify the routing instance used by the probes.
Options	<i>instance-name</i> —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 910.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rpm

Syntax	<code>rpm <i>client</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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	<i>client</i> —Identifier for RPM client router.
Usage Guidelines	See “Configuring RPM Timestamping” on page 917.
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 920.

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 “Configuring Real-Time Performance Monitoring Properties” on page 912.

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

source-address

Syntax	<code>source-address <i>address</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code>
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	<i>address</i> —Valid IP address.
Usage Guidelines	See “Configuring the Probe” on page 912.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

target

Syntax	<code>target (url <i>url</i> address <i>address</i>);</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the destination address used for the probes.
Options	url <i>url</i> —For HTTP probe types, specify a fully formed URL that includes <code>http://</code> in the URL address. address <i>address</i> —For all other probe types, specify an IPv4 address for the target host.
Usage Guidelines	See “Configuring the Probe” on page 912.
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 the Server” on page 917.

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 the Probe” on page 912.

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

test-interval

Syntax	<code>test-interval <i>frequency</i>;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe <i>owner</i> test <i>test-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the time to wait between tests, in seconds.
Options	<i>frequency</i> —Number of seconds, from 0 through 86400.
Usage Guidelines	See “Configuring BGP Neighbor Discovery Through RPM” on page 910 or “Configuring the Probe” on page 912.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

thresholds

Syntax	<code>thresholds thresholds;</code>
Hierarchy Level	<code>[edit services rpm probe owner test test-name]</code>
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, a Simple Network Management Protocol (SNMP) trap (if configured) is generated when a threshold is exceeded.
Options	<p><i>thresholds</i>—Specify one or more threshold measurements. The following options are supported:</p> <ul style="list-style-type: none"> ■ <i>egress-time</i>—Measures maximum source-to-destination time per probe. ■ <i>ingress-time</i>—Measures maximum destination-to-source time per probe. ■ <i>jitter-egress</i>—Measures maximum source-to-destination jitter per test. ■ <i>jitter-ingress</i>—Measures maximum destination-to- source jitter per test. ■ <i>jitter-rtt</i>—Measures maximum jitter per test, from 0 through 60,000,000 microseconds. ■ <i>rtt</i>—Measures maximum round-trip time per probe, in microseconds. ■ <i>std-dev-egress</i>—Measures maximum source-to-destination standard deviation per test. ■ <i>std-dev-ingress</i>—Measures maximum destination-to-source standard deviation per test. ■ <i>std-dev-rtt</i>—Measures maximum standard deviation per test, in microseconds. ■ <i>successive-loss</i>—Measures successive probe loss count, indicating probe failure. ■ <i>total-loss</i>—Measures total probe loss count indicating test failure, from 0 through 15.
Usage Guidelines	See “Configuring the Probe” on page 912.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

traps

Syntax	<code>traps traps;</code>
Hierarchy Level	<code>[edit services rpm probe owner test test-name]</code>
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	<p><code>traps</code>—Specify one or more traps. The following options are supported:</p> <ul style="list-style-type: none"> ■ <code>egress-jitter-exceeded</code>—Generates traps when the jitter in egress time threshold is met or exceeded. ■ <code>egress-std-dev-exceeded</code>—Generates traps when the egress time standard deviation threshold is met or exceeded. ■ <code>egress-time-exceeded</code>—Generates traps when the maximum egress time threshold is met or exceeded. ■ <code>ingress-jitter-exceeded</code>—Generates traps when the jitter in ingress time threshold is met or exceeded. ■ <code>ingress-std-dev-exceeded</code>—Generates traps when the ingress time standard deviation threshold is met or exceeded. ■ <code>ingress-time-exceeded</code>—Generates traps when the maximum ingress time threshold is met or exceeded. ■ <code>jitter-exceeded</code>—Generates traps when the jitter in round-trip time threshold is met or exceeded. ■ <code>probe-failure</code>—Generates traps for successive probe loss thresholds crossed. ■ <code>rtt-exceeded</code>—Generates traps when the maximum round-trip time threshold is met or exceeded. ■ <code>std-dev-exceeded</code>—Generates traps when the round-trip time standard deviation threshold is met or exceeded. ■ <code>test-completion</code>—Generates traps when a test is completed. ■ <code>test-failure</code>—Generates traps when the total probe loss threshold is met or exceeded.
Usage Guidelines	See “Configuring the Probe” on page 912.
Required Privilege Level	<p><code>interface</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>

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 920.

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 920.

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 the Server” on page 917.

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 955
- Tunnel Interfaces Configuration Guidelines on page 959
- Summary of Tunnel Services Configuration Statements on page 975

Chapter 50

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 Multiprotocol Label Switching (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 647 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 JUNOS software supports the tunnel types shown in Table 25 on page 955.

Table 25: Tunnel Interface Types

Interface	Description
gr-0/0/0	Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol over another routing protocol. Within a router, packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then re-encapsulated with another protocol packet to complete the GRE. The GRE interface is an internal interface only and is not associated with a physical 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 25: Tunnel Interface Types *(continued)*

Interface	Description
ip-0/0/0	<p>Configurable IP-over-IP encapsulation (also called IP tunneling) interface. IP tunneling allows the encapsulation of one IP packet over another IP packet.</p> <p>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.</p> <p>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.</p>
ipip	<p>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.</p>
lt-0/0/0	<p>The lt interface on M-series and T-series routing platforms supports configuration of logical systems—the capability to partition a single physical router into multiple logical devices that perform independent routing tasks.</p> <p>On J-series Services Routers, it has a different function: it provides class-of-service (CoS) support for data link switching (DLSw) traffic and real-time performance monitoring (RPM) probe packets.</p> <p>Within a Services 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 DLSW and RPM services.</p>
mt-0/0/0	<p>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.</p> <p>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 routing platform 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.</p>
mtun	<p>Internally generated multicast tunnel interface. This interface is generated by the JUNOS software to handle multicast tunnel services. It is not a configurable interface.</p>

Table 25: Tunnel Interface Types (*continued*)

Interface	Description
pd-0/0/0	<p>Configurable Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop routing platform 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.</p> <p>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.</p>
pe-0/0/0	<p>Configurable Protocol Independent Multicast (PIM) encapsulation interface. In PIM sparse mode, the first-hop routing platform 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.</p> <p>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.</p>
pimd	Internally generated Protocol Independent Multicast (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 Protocol Independent Multicast (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	<p>Configurable virtual loopback tunnel interface. Facilitates VRF table lookup based on MPLS labels. This interface type is supported on M-series and T-series routing platforms, but not on J-series Services Routers.</p> <p>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.</p>

Chapter 51

Tunnel Interfaces Configuration Guidelines

This chapter includes the following tunnel interface configuration tasks and examples:

- Configuring a Unicast Tunnel on page 959
- Restricting a Tunnel to Multicast Traffic on page 964
- Configuring a Logical Tunnel Interface on page 964
- Configuring a Tunnel Interface for Routing Table Lookup on page 966
- Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 967
- Configuring PIM Tunnels on page 968
- Configuring an IPv6-over-IPv4 Tunnel on page 969
- Configuring a Dynamic Tunnel on page 969
- Configuring Tunnel Interfaces on MX-series Routers on page 970
- Example: Configuring Unicast Tunnels on page 970
- Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 971
- Example: Configuring an IPv6-over-IPv4 Tunnel on page 972
- Example: Configuring a Logical Tunnel on page 972

Configuring a Unicast Tunnel

To configure a unicast tunnel, you configure the `gr` interface (to use GRE encapsulation) or the `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 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, 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 “Example: Configuring Unicast Tunnels” on page 970 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 PICs, you can configure additional tunnel attributes, as described in the following sections:

- Configuring a Key Number on GRE Tunnels on page 961
- Enabling Fragmentation on GRE Tunnels on page 962
- Specifying an MTU Setting for the Tunnel on page 962
- Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 963
- Configuring Packet Reassembly on page 963

Configuring a Key Number on GRE Tunnels

For Adaptive Services and MultiServices PICs on M-series and T-series routing platforms, 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 PICs 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 228.

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 961 and “Enabling Fragmentation on GRE Tunnels” on page 962), 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 a Tunnel 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 routing platform 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 a Logical Tunnel Interface

Logical tunnel (lt) interfaces provide quite different services depending on the host routing platform:

- On M-series and T-series routing platforms, 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, It 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 sections:

- Connecting Logical Systems on page 965
- Configuring Logical Tunnels on J-series Platforms on page 966

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 Logical Tunnels on J-series Platforms

On J-series Services Routers, you can configure the `lt` interface to provide class-of-service (CoS) support for data link switching (DLSw) traffic and real-time performance monitoring (RPM) probe packets.

Within a J-series Services 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 DLSW and RPM services. For more information, see the *J-series Services Router Advanced WAN Access Configuration Guide*.



NOTE: The `lt` interface on the J-series Services Router does not support logical systems.

Configuring a Tunnel Interface 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 a Virtual Loopback Tunnel for VRF Table Lookup” on page 967.

For more information about VPNs, see the *JUNOS VPNs Configuration Guide*.

Configuring a Virtual Loopback Tunnel for VRF Table Lookup

To enable egress filtering, you can either configure filtering based on the IP header, or you can configure a virtual loopback tunnel on routers equipped with a Tunnel PIC. Table 26 on page 967 describes each method.

Table 26: Methods for Configuring Egress Filtering

Method	Interface Type	Configuration Guidelines	Comments
Filter traffic based on the IP header	Nonchannelized Point-to-Point Protocol / High Level Data Link Control (PPP/HDLC) core-facing SONET/SDH interfaces	Include the <code>vrf-table-label</code> statement at the <code>[edit routing-instances instance-name]</code> hierarchy level. For more information, see the <i>JUNOS VPNs Configuration Guide</i> .	You cannot include the <code>vrf-table-label</code> statement when configuring the 10-port E1 PIC, aggregated interfaces, Fast Ethernet 12-port and 48-port PIC, Gigabit Ethernet 4-port PIC, or Gigabit Ethernet intelligent queuing (IQ) PIC. 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 <code>vrf-table-label</code> 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 an IPv6-over-IPv4 Tunnel

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 a Unicast Tunnel” on page 959. For an IPv6-over-IPv4 tunnel configuration example, see “Example: Configuring an IPv6-over-IPv4 Tunnel” on page 972.

Configuring a Dynamic Tunnel

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 Border Gateway Protocol (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*.

Example: 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;
  }
}
```

To configure a numbered tunnel interface, include an address under **family inet**:

```
[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;
    }
}

```

To configure MPLS over GRE tunnels, include the `family mpls` statement:

```

[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.

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;
      }
    }
  }
}

```

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 a Logical Tunnel

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 52

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 <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> 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 963.
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 <i>address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel interfaces, specify the remote address of the backup tunnel.
Options	<i>address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring IPSec Tunnel Redundancy” on page 655.
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 <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
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 963.
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 (Address) on page 977
- destination (Routing Instance) on page 977

destination (Address)

Syntax	<code>destination address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel interfaces, specify the remote address of the tunnel.
Options	<i>destination-address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring a Unicast Tunnel” on page 959, “Configuring Traffic Sampling” on page 675, and “Configuring Flow Monitoring” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination (Routing Instance)

Syntax	<code>destination routing-instance-name;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 <code>inet.0</code> .
Usage Guidelines	See “Configuring a Tunnel Interface for Routing Table Lookup” on page 966.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination-networks

Syntax	<code>destination-networks prefix;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Create a tunnel for routes in these destination networks.
Options	<i>prefix</i> —Destination prefix of network.
Usage Guidelines	See “Configuring a Dynamic Tunnel” on page 969.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

do-not-fragment

Syntax	<code>do-not-fragment;</code>
Hierarchy Level	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> 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 963.
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 <i>tunnel-name</i> { destination-networks <i>prefix</i> ; source-address <i>address</i> ; tunnel-type <i>type-of-tunnel</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> 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	<i>tunnel-name</i> —Name of the dynamic tunnel. The statements are explained separately in this chapter.
Usage Guidelines	See “Configuring a Dynamic Tunnel” on page 969.
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 router.
Default	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
Usage Guidelines	See the <i>JUNOS Network Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

key

Syntax	<code>key number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Adaptive Services and MultiServices Physical Interface Cards (PICs) on M-series and T-series routing platforms, identify an individual traffic flow within a tunnel, as defined in RFC 2890, <i>Key and Sequence Number Extensions to GRE</i> .
Options	<i>number</i> —Value of the key. Range: 0 through 4,294,967,295
Usage Guidelines	See “Configuring a Key Number on GRE Tunnels” on page 961.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

multicast-only

Syntax	<code>multicast-only;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 a Tunnel to Multicast Traffic” on page 964.
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	<code>peer-unit <i>unit-number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a peer relationship between two logical systems.
Options	<i>unit-number</i> —Peering logical system unit number.
Usage Guidelines	See “Configuring a Logical Tunnel Interface” on page 964.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

reassemble-packets

Syntax	<code>reassemble-packets;</code>
Hierarchy Level	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>]
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 963.
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 <i>routing-instance-name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 a Tunnel Interface for Routing Table Lookup” on page 966.
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 <i>routing-instance-name</i> { ... }
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an additional routing entity for a router. You can create multiple instances of Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), OSPF version 3 (OSPFv3), and RIP for a router.
Default	Routing instances are disabled for the router.
Options	<i>routing-instance-name</i> —Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately.
Usage Guidelines	See the <i>JUNOS Routing Protocols Configuration Guide</i> and the <i>JUNOS Policy Framework Configuration Guide</i> .
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 <i>logical-system-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure protocol-independent routing properties.
Usage Guidelines	See the <i>JUNOS Routing Protocols Configuration Guide</i> .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

source

Syntax	source <i>source-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel <i>destination-address</i>]
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	<i>source-address</i> —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 959.
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 <i>address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the tunnel source address.
Options	<i>address</i> —Name of the source address.
Usage Guidelines	See “Configuring a Dynamic Tunnel” on page 969.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ttl

Syntax	ttl <i>value</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> 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	<i>value</i> —Time-to-live value. Range: 0 through 255 Default: 64
Usage Guidelines	See “Tunnel Interfaces Configuration Guidelines” on page 959.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

tunnel

Syntax	<pre> tunnel { allow-fragmentation; backup-destination <i>address</i>; destination <i>destination-address</i>; do-not-fragment; key <i>number</i>; routing-instance { destination <i>routing-instance-name</i>; } source <i>source-address</i>; ttl <i>number</i>; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>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).</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Encryption Interfaces Configuration Guidelines” on page 647 and “Tunnel Interfaces Configuration Guidelines” on page 959.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS VPNs Configuration Guide</i>

tunnel-type

Syntax	tunnel-type <i>type</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Select the dynamic tunnel type.
Options	<i>type</i> —Tunnel type. Generic routing encapsulation (GRE) is supported.
Usage Guidelines	See “Configuring a Dynamic Tunnel” on page 969.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

unit

Syntax	<pre> unit <i>logical-unit-number</i> { peer-unit <i>unit-number</i>; reassemble-packets; tunnel { allow-fragmentation; backup-destination <i>address</i>; destination <i>destination-address</i>; do-not-fragment; key <i>number</i>; routing-instance { destination <i>routing-instance-name</i>; } source <i>source-address</i>; ttl <i>number</i>; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
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	<i>logical-unit-number</i> —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 959.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Network Interfaces Configuration Guide</i> for other statements that do not affect services interfaces.

Part 9

Index

- Index on page 991
- Index of Statements and Commands on page 1009

Index

Symbols

#, comments in configuration statements.....	xlvi
(), in syntax descriptions.....	xlvi
< >, in syntax descriptions.....	xlvi
[], in configuration statements.....	xlvi
{ }, in configuration statements.....	xlvi
(pipe), in syntax descriptions.....	xlvi

A

accept	
action.....	675
accounting statement.....	720
usage guidelines.....	710
acknowledge-retries statement.....	885
usage guidelines.....	860
acknowledge-timer statement.....	886
usage guidelines.....	860
action-red-differential-delay statement.....	886
usage guidelines.....	861
activation priority.....	392
activation-priority statement.....	397
usage guidelines.....	392
adaptive-services-pics statement.....	451
usage guidelines.....	780
address statement	
DFC.....	823
usage guidelines.....	808
encryption.....	657
usage guidelines.....	647
flow monitoring.....	721
usage guidelines.....	675
interfaces.....	481
usage guidelines.....	470
link services.....	887
usage guidelines.....	865
NAT.....	145
usage guidelines.....	124
voice services.....	398
usage guidelines.....	385
VoIP	
usage guidelines.....	390
address-range statement	
NAT.....	146

administrative statement	
packet gateway.....	523
advertise-interval statement.....	623
aggregate-export-interval statement.....	721
usage guidelines.....	710
aggregation statement.....	181, 722
usage guidelines.....	174, 688
alert (system logging severity level).....	287, 441, 471
algorithm statement	
packet gateway.....	523
usage guidelines.....	506
ALGs	
application protocols.....	53
configuring.....	54
definition.....	53
allow-fragmentation statement.....	975
usage guidelines.....	963
allow-ip-option statement.....	112
usage guidelines.....	104
allow-multicast statement.....	452
usage guidelines.....	442
allowed-destinations statement.....	824
usage guidelines.....	809
analyzer-address statement.....	787
usage guidelines.....	777
analyzer-id statement.....	787
usage guidelines.....	777
anomaly checklist.....	36
anti-replay-window-size statement	
usage guidelines.....	229
any (system logging severity level).....	287, 441, 471
application layer gateways <i>See</i> ALGs	
application protocol	
definition.....	53
application statement.....	91
usage guidelines.....	54
application-data-inactivity-detection statement	
packet gateway.....	524
application-profile statement.....	419
usage guidelines.....	414
application-protocol statement.....	92
usage guidelines.....	54
application-set statement.....	93
usage guidelines.....	64

application-sets statement	
CoS.....	420
usage guidelines.....	412
IDS.....	182
usage guidelines.....	173
NAT.....	146
usage guidelines.....	130
stateful firewall.....	113
usage guidelines.....	103
applications.....	173
example configuration.....	87
applications statement	
applications hierarchy.....	93
usage guidelines.....	53
CoS.....	420
usage guidelines.....	412
IDS.....	182
usage guidelines.....	173
NAT.....	147
usage guidelines.....	130
stateful firewall.....	113
usage guidelines.....	103
applying service set to interface.....	434
archive-sites statement.....	788
usage guidelines.....	779
AS PIC	
multicast traffic.....	442
redundancy.....	289, 475, 716
attack detection.....	169
audit-observed-events-returns-history statement	
packet gateway.....	524
authentication statement.....	245
usage guidelines.....	209
authentication-algorithm statement	
IKE.....	246
IPSec.....	246
authentication-algorithm statement (IKE)	
usage guidelines.....	212
authentication-algorithm statement (IPSec)	
usage guidelines.....	220
authentication-method statement.....	247
usage guidelines.....	213
autonomous-system-type statement.....	723
usage guidelines.....	688
auxiliary-spi statement.....	247
usage guidelines.....	209
Avaya VoIP product.....	390

B

backup AS PIC.....	475
backup Link Services IQ PIC.....	316
backup-destination statement.....	658, 976
usage guidelines.....	655
backup-interface statement.....	658
usage guidelines.....	654

backup-remote-gateway statement.....	248
usage guidelines.....	228
bandwidth	
and delay buffer allocation.....	331
guaranteed.....	331, 335
base-root	
packet gateway.....	525
base-root statement	
packet gateway.....	525
bearer bandwidth limit.....	392
bearer-bandwidth-limit statement.....	399
usage guidelines.....	392
BGP	
router identifier.....	983
braces, in configuration statements.....	xlvi
brackets	
angle, in syntax descriptions.....	xlvi
square, in configuration statements.....	xlvi
bundle statement.....	399, 887
usage guidelines.....	389, 864
by-destination statement.....	183
usage guidelines.....	176
by-pair statement.....	184
usage guidelines.....	176
by-source statement.....	185
usage guidelines.....	176

C

call admission control.....	392
cancel-graceful statement	
packet gateway.....	526, 527
capture-group statement.....	825
usage guidelines.....	807
cflowd statement.....	724
usage guidelines.....	687
CIR.....	335
circuit-weight statement.....	624
usage guidelines.....	613
cisco-interoperability statement.....	371
usage guidelines.....	315
cleanup-timeout statement	
packet gateway.....	527
usage guidelines.....	503
clear-dont-fragment-bit statement	
GRE tunnel.....	482
IPSec.....	248
usage guidelines.....	228, 472, 962
clear-ike-sas-on-pic-restart statement.....	249
usage guidelines.....	211
clear-ipsec-sas-on-pic-restart statement.....	249
usage guidelines.....	211
client-list statement.....	927
collector statement.....	788
usage guidelines.....	778

collector-pic statement
 usage guidelines.....780
 comments, in configuration statements.....xlvii
 compressed RTP.....856, 857
 example configuration.....857, 858
 compression statement.....400
 usage guidelines.....386, 387
 compression-device statement.....400, 888
 usage guidelines.....389, 856, 857
 configuration
 dynamic flow capture interface.....813
 flow collector interface.....780
 flow-tap application.....821
 connection-idle-timeout statement.....624
 usage guidelines.....615
 content destinations
 DFC.....805
 flow-tap.....818
 content-destination statement.....826
 usage guidelines.....808
 context-indications statement
 packet gateway.....528
 control source
 DFC.....805
 control-association-indications statement
 packet gateway.....529
 control-source statement.....827
 usage guidelines.....809
 controller-address statement
 packet gateway.....530
 usage guidelines.....506
 controller-failure statement
 packet gateway.....530
 controller-port statement
 packet gateway.....531
 usage guidelines.....506
 conventions
 text and syntax.....xlvii
 copy-tos-to-outer-ip-header statement.....976
 usage guidelines.....963
 core-dump statement.....725
 usage guidelines.....683
 CoS
 action statements.....413
 applications.....412
 example configuration.....417
 for tunnels
 GRE TOS bits.....963
 link services interfaces.....328, 330, 866
 link services IQ interfaces.....311
 match conditions.....412
 rules.....410
 scheduler map
 configuration example.....853
 cost statement.....625
 usage guidelines.....613

critical (system logging severity level).....287, 442, 471
 curly braces, in configuration statements.....xlvii
 customer support.....lv
 contacting JTAC.....lv

D

Data inactivity detection.....579
 data link switching *See* DLSw
 data-fill statement.....928
 data-format statement.....789
 usage guidelines.....778
 data-inactivity-detection statement
 packet gateway.....531
 data-size statement.....928
 usage guidelines.....912
 dead peer detection (DPD) protocol.....228
 delay buffer
 calculating.....331, 335
 shaping rate.....331, 335
 delay-buffer-rate statement
 usage guidelines.....331
 delivery-function
 packet gateway.....532
 delivery-function statement
 packet gateway.....532
 description statement
 IPSec.....250
 usage guidelines.....218, 220, 223
 destination statement
 DLSw.....625
 encryption.....659
 usage guidelines.....647, 655
 flow monitoring.....726
 usage guidelines.....675
 link services.....888
 usage guidelines.....865
 tunnel.....977
 usage guidelines.....959, 966
 VoIP
 usage guidelines.....390
 destination-address
 packet gateway.....532
 destination-address statement
 CoS.....421
 usage guidelines.....412
 IDS.....186
 usage guidelines.....173
 IPSec.....250
 usage guidelines.....226
 NAT.....147
 usage guidelines.....130
 packet gateway.....532
 stateful firewall.....114
 usage guidelines.....103

destination-address-range statement		
IDS.....	186	
usage guidelines.....	173	
NAT.....	148	
usage guidelines.....	130	
stateful firewall.....	114	
usage guidelines.....	103	
destination-interface statement.....	626, 929	
usage guidelines.....	912, 918	
destination-networks statement		
usage guidelines.....	969	
destination-pool statement.....	148	
usage guidelines.....	131	
destination-port		
packet gateway.....	533	
destination-port statement		
applications.....	93	
packet gateway.....	533	
RPM.....	930	
usage guidelines.....	59	
destination-prefix statement.....	149, 187	
usage guidelines.....	175	
destination-prefix-ipv6 statement.....	187	
usage guidelines.....	175	
destination-prefix-list statement		
CoS.....	421	
IDS.....	188	
NAT.....	149	
stateful firewall.....	115	
usage guidelines.....	103	
destinations statement.....	789	
usage guidelines.....	777	
detect		
packet gateway.....	533	
detect statement		
packet gateway.....	533	
DFC		
architecture.....	805	
capture group.....	807	
control source configuration.....	809	
destination configuration.....	808	
example configuration.....	813	
interface configuration.....	810	
system logging.....	811	
threshold configuration.....	811	
dh-group statement.....	251	
usage guidelines.....	213	
dial-options statement.....	483	
interfaces		
usage guidelines.....	288	
diffserv statement.....	534	
direction statement.....	252	
usage guidelines.....	207	
disable statement.....	726	
traffic sampling		
usage guidelines.....	677	
disable-mlppp-inner-ppp-pfc statement.....	889	
usage guidelines.....	850	
disable-session-mirroring		
packet gateway.....	534	
disable-session-mirroring statement		
packet gateway.....	534	
discard accounting		
usage guidelines.....	710	
disconnect statement		
packet gateway.....	535	
dlci statement.....	889	
usage guidelines.....	848	
DLCIs		
multicast-capable connections.....	848	
point-to-point connections.....	848	
DLSw		
configuration statements.....	611	
load balancing.....	613	
timers.....	613	
tracing operations.....	617	
dlsw statement.....	626	
usage guidelines.....	612	
dlsw-cos statement.....	627	
usage guidelines.....	615	
do-not-fragment statement.....	978	
usage guidelines.....	963	
documentation set		
comments on.....	lv	
down statement		
packet gateway.....	536	
drop-timeout statement.....	890	
usage guidelines.....	849	
dscp statement.....	422, 537	
usage guidelines.....	413	
dscp-code-point statement.....	931	
usage guidelines.....	912	
DTCP.....	805, 817	
duplicates-dropped-periodicity statement.....	827	
usage guidelines.....	812	
dynamic authentication.....	231	
dynamic CAC configuration.....	392	
dynamic flow capture <i>See</i> DFC		
dynamic route insertion.....	232	
dynamic rules.....	231	
dynamic security associations		
usage guidelines.....	211, 212	
dynamic statement.....	253	
usage guidelines.....	211	
Dynamic Tasking Control Protocol <i>See</i> DTCP		
dynamic tunnels		
destination.....	978	
source.....	984	
dynamic-call-admission-control statement.....	401	
usage guidelines.....	392	
dynamic-flow-capture statement.....	828	

dynamic-tunnels statement.....979
 usage guidelines.....969

E

emergency (system logging severity level).....287, 441, 471
 enable flow collection mode.....780
 encapsulation statement.....402, 891
 usage guidelines.....850
 voice services
 usage guidelines.....388
 encoding statement
 packet gateway.....537
 encryption interface.....647
 applying inbound filter.....653
 example configuration.....653
 applying outbound filter.....652
 example configuration.....651, 652
 configuring inbound filter.....652
 example configuration.....653
 configuring MTU.....648
 encryption statement.....254
 usage guidelines.....210
 encryption-algorithm statement (IKE).....255
 usage guidelines.....213
 encryption-algorithm statement (IPSec).....255
 usage guidelines.....220
 engine-id statement.....727
 engine-type statement.....727
 error (system logging severity level).....287, 442, 472
 ES interfaces
 example configuration.....648
 ES PIC
 apply inbound filter.....653
 PIC redundancy.....654
 redundancy
 example configuration.....655
 tunnel redundancy.....655
 ES-IS.....611
 es-options statement.....659
 usage guidelines.....654
 event policy
 all (tracing flag).....448
 configuration (tracing flag).....448
 database (tracing flag).....448
 events (tracing flag).....448
 policy (tracing flag).....448
 event-timestamp-notification statement
 packet gateway.....538
 explorer-wait-time statement.....627
 usage guidelines.....613
 export-format statement.....728
 usage guidelines.....685
 extension-service statement.....452
 usage guidelines.....443

F

f-max-period statement.....402
 usage guidelines.....386
 facility-override statement.....297, 453, 484
 usage guidelines.....441
 failover-cold statement
 packet gateway.....538
 failover-warm statement
 packet gateway.....539
 failure statement
 packet gateway.....541
 family statement
 encryption.....660
 usage guidelines.....647
 flow monitoring.....729
 usage guidelines.....675
 interfaces.....485
 usage guidelines.....470
 link services.....403, 893
 usage guidelines.....864
 fast-update-filters statement
 packet gateway.....541
 file statement.....732, 733
 traffic sampling output
 usage guidelines.....677, 679
 file-specification statement.....791
 usage guidelines.....778
 filename statement.....733
 filename-prefix statement.....790
 usage guidelines.....779
 files
 logging information output file.....679
 traffic sampling output files.....677
 var/log/sampled file.....679
 var/tmp/sampled.pkts file.....677
 files statement.....734
 usage guidelines.....678
 filter statement
 encryption.....661
 usage guidelines.....653
 flow monitoring.....735
 usage guidelines.....675
 filters
 used with services.....434
 firewall filters
 actions.....675
 in traffic sampling.....675
 service filters.....474
 flow aggregation.....687
 multiple flow servers.....695
 flow collector
 analyzer configuration.....777
 destination configuration.....777
 example configuration.....780
 file format configuration.....778

interface mapping.....	778
transfer log.....	779
flow limiting.....	441
flow monitoring	
example configuration	
multiple port mirroring.....	703
next-hop groups.....	703
load balancing.....	707
overview.....	667
redundancy.....	716
flow server	
replicating flows to multiple servers.....	695
flow-active-timeout statement.....	736
usage guidelines.....	685
flow-collector statement.....	792
usage guidelines.....	776, 780
flow-export-destination statement.....	737
usage guidelines.....	685
flow-inactive-timeout statement.....	737
usage guidelines.....	685
flow-monitoring statement.....	738
flow-tap	
application.....	817
architecture.....	818
interface.....	819
permissions statement.....	820
RADIUS configuration.....	820
restrictions.....	820
security.....	820
flow-tap application	
example configuration.....	821
flow-tap statement.....	829
flow-tap-dtcp statement.....	820
font conventions.....	xlvi
force-entry statement.....	188
usage guidelines.....	175
forwarding classes	
fragmentation.....	328
forwarding-class statement.....	372, 422
usage guidelines.....	328, 413
forwarding-options statement.....	738
usage guidelines.....	671
fragment-threshold statement	
link services.....	894
usage guidelines.....	851
LSQ.....	373
usage guidelines.....	328
voice services.....	404
usage guidelines.....	387
fragmentation	
forwarding classes.....	328
GRE tunnels.....	962
multiclass MLPPP.....	330
fragmentation and reassembly.....	387, 851
example configuration.....	853
fragmentation-map statement.....	373
usage guidelines.....	329
fragmentation-maps statement.....	374
usage guidelines.....	328
Frame Relay connections	
point-to-point connections.....	848
Frame Relay encapsulation	
multicast-capable connections.....	848
FRF.12.....	387
example configuration.....	360
LFI.....	851
LSQ.....	357
FRF.15 and FRF.16.....	845
FRF.16.....	348
configuration example.....	351
from statement	
CoS.....	423
usage guidelines.....	410, 411
IDS.....	189
usage guidelines.....	172, 173
IPSec.....	256
usage guidelines.....	225, 226
NAT.....	150
usage guidelines.....	128, 130
stateful firewall.....	115
usage guidelines.....	102, 103
ftp statement.....	794
usage guidelines.....	777, 779
FTP traffic, sampling.....	682
full-cone NAT.....	129
G	
g-duplicates-dropped-periodicity statement.....	829
usage guidelines.....	812
g-max-duplicates statement.....	830
usage guidelines.....	812
gateway statement	
packet gateway.....	542
gateway-address statement	
packet gateway.....	545
usage guidelines.....	503
gateway-controller statement	
packet gateway.....	546
usage guidelines.....	506
gateway-port statement	
packet gateway.....	546
usage guidelines.....	503
graceful statement	
packet gateway.....	547, 548
graceful-restart	
packet gateway.....	548, 563, 564
graceful-restart statement	
packet gateway.....	548, 563, 564

GRE tunnels	
fragmentation.....	962
key number.....	961
guaranteed rate.....	335
guaranteed-rate statement	
usage guidelines.....	335

H

H.248 properties	
packet	
gateway.....	534, 537, 551, 566, 567, 568, 569, 570, 572
H.248 timers	
packet gateway configuration.....	508
h248-options	
packet gateway.....	549
h248-options statement	
packet gateway.....	549
h248-properties statement.....	551
h248-timers	
packet gateway.....	552
h248-timers statement	
packet gateway.....	552
usage guidelines.....	508
hanging-termination-detection statement	
packet gateway.....	553
hard-limit statement.....	830
usage guidelines.....	808
hard-limit-target statement.....	831
usage guidelines.....	808
hardware requirements.....	3
hardware-timestamp statement.....	932
usage guidelines.....	920
hello-interval statement.....	298
usage guidelines.....	286
hello-timer statement.....	894
usage guidelines.....	860
hide-avps statement.....	298
usage guidelines.....	286
hint statement.....	151
history-size statement.....	932
usage guidelines.....	911, 912
hold-time statement.....	628
host statement.....	299, 453, 486
usage guidelines.....	286, 441, 471
hot-standby statement.....	374

I

icmp-code statement.....	94
usage guidelines.....	57
icmp-type statement.....	95
usage guidelines.....	57
icons defined, notice.....	xlvi

IDS	
action statements.....	174
applications.....	173
example configurations.....	178
match conditions.....	173
rules.....	171
ids-rule-sets statement.....	454
usage guidelines.....	438
ids-rules statement.....	454
usage guidelines.....	438
ignore-entry statement.....	188, 189
usage guidelines.....	175
IKE.....	40, 212
authentication algorithm	
usage guidelines.....	212
authentication method	
usage guidelines.....	213
DH (Diffie-Hellman) group	
usage guidelines.....	213
dynamic SAs.....	212
encryption-algorithm statement	
usage guidelines.....	213
lifetime	
usage guidelines.....	214
mode	
usage guidelines.....	216
policy	
example.....	219
usage guidelines.....	215
policy statement	
usage guidelines.....	215
pre-shared-key statement	
usage guidelines.....	216
proposals statement	
usage guidelines.....	216
IKE security associations	
clearing.....	211
ike statement.....	257
usage guidelines.....	212
ike-access-profile statement.....	454
usage guidelines.....	234, 440
inactivity-delay statement	
packet gateway.....	554
inactivity-duration statement	
packet gateway.....	554
inactivity-timeout	
packet gateway.....	553
inactivity-timeout statement.....	95, 486, 933
packet gateway.....	553
usage guidelines.....	62, 470
inactivity-timer	
packet gateway.....	555
inactivity-timer statement	
packet gateway.....	555
info (system logging severity level).....	287, 442, 472

initial-average-ack-delay		
packet gateway.....	555	
initial-average-ack-delay statement		
packet gateway.....	555	
usage guidelines.....	508	
initiate-dead-peer-detection statement.....	258	
usage guidelines.....	229	
input statement		
flow monitoring.....	739	
interfaces.....	487	
usage guidelines.....	434, 473	
input-interface-index statement.....	740	
input-packet-rate-threshold statement.....	831	
usage guidelines.....	811	
inside and outside interfaces.....	436	
inside-service-interface statement		
usage guidelines.....	434	
interchassis LSQ failover.....	314	
interface preservation.....	319	
interface statement.....	741	
DLSw.....	629	
encryption		
usage guidelines.....	647	
flow monitoring		
usage guidelines.....	700	
flow-tap.....	832	
usage guidelines.....	819	
packet gateway.....	556	
usage guidelines.....	507	
service interface pool.....	605	
interface style service sets.....	437	
interface-map statement.....	796	
usage guidelines.....	778	
interface-service statement.....	455	
usage guidelines.....	434	
interfaces		
naming.....	468	
interfaces statement		
DFC.....	832	
usage guidelines	810	
encryption.....	661	
usage guidelines.....	647	
flow monitoring.....	743	
usage guidelines.....	675	
interfaces hierarchy.....	487	
usage guidelines.....	467	
link services.....	404, 895	
usage guidelines.....	845	
tunnel.....	979	
usage guidelines.....	959	
interim-ah-scheme statement		
packet gateway.....	556	
usage guidelines.....	506	
interleave-fragments statement.....	895	
usage guidelines.....	851, 853	
Internet Key Exchange. <i>See</i> IKE		
intrachassis LSQ failover.....	316	
intrusion detection		
example configurations.....	178	
rule set.....	171	
tasks.....	169	
IP addresses		
sampling traffic from single IP addresses.....	681	
ip-flow-stop-detection-statement		
packet gateway.....	557	
IPSec		
action statements.....	227	
authentication algorithm		
usage guidelines.....	220	
authentication statement		
usage guidelines.....	209	
direction		
usage guidelines.....	207	
dynamic authentication.....	231	
dynamic endpoints interface configuration.....	234	
dynamic rules.....	231	
dynamic security associations		
usage guidelines.....	211	
encryption		
usage guidelines.....	210	
encryption-algorithm statement		
usage guidelines.....	220	
ES PIC.....	647	
example configuration.....	236	
inbound traffic.....	653	
outbound traffic.....	651	
IKE.....	40	
lifetime of SA		
usage guidelines.....	221	
match conditions.....	226	
minimum configurations		
dynamic SA	205	
manual SA	205	
overview.....	40	
perfect-forward-secrecy statement		
usage guidelines.....	223	
policy.....	222	
overview.....	222	
policy statement		
usage guidelines.....	222	
proposal statement		
usage guidelines.....	220	
proposals statement		
usage guidelines.....	223	
protocol statement (dynamic SA)		
usage guidelines.....	222	
protocol statement (manual SA)		
usage guidelines.....	208	
rule sets.....	224	
security associations.....	40	
security parameter index		
usage guidelines.....	209	

service set dynamic endpoints	
configuration.....	234
traffic.....	649
ipsec statement.....	258
usage guidelines.....	220
ipsec-inside-interface	
usage guidelines.....	231
ipsec-inside-interface statement.....	259
usage guidelines.....	226
ipsec-interface-id statement	
usage guidelines.....	234
ipsec-sa statement.....	662
usage guidelines.....	647
ipsec-vpn-options statement.....	455
usage guidelines.....	439
ipsec-vpn-rule-sets statement.....	456
usage guidelines.....	438
ipsec-vpn-rules statement.....	456
usage guidelines.....	438
ipv4-template statement.....	743
IPv6	
transition	
configured tunnel.....	969
ipv6-multicast-filter statement	
usage guidelines.....	127
ipv6-multicast-interfaces statement.....	151
IPv6-over-IPv4 tunnel	
example configuration.....	972
standards supported.....	969
IS-IS	
standards supported.....	609
tracing operations.....	639

K

key statement.....	980
usage guidelines.....	961

L

L2TP	
access profile.....	284
attribute-value pairs.....	286
example configuration.....	292, 393
redundancy.....	289
timers.....	286
l2tp statement	
usage guidelines	279
l2tp-access-profile statement.....	299
usage guidelines.....	284
l2tp-interface-id statement	
usage guidelines.....	288
l2tp-profile statement	
usage guidelines.....	284
label-position statement.....	744

latch-deadlock-delay statement	
packet gateway.....	557
lawful intercept architecture.....	818
learn-sip-register statement.....	96
usage guidelines.....	62
LFI.....	353, 357, 387, 851
example configuration.....	356, 360, 853
lifetime-seconds statement (IKE).....	259
usage guidelines.....	214
lifetime-seconds statement (IPSec).....	259
usage guidelines.....	221
limiting flows per service set.....	441
link fragmentation and interleaving <i>See</i> LFI	
link PIC redundancy.....	319
link services interfaces	
CoS components.....	328, 330, 866
example configuration.....	868, 875
interleave fragments.....	851
example configuration.....	853
link services IQ interfaces.....	356
CoS components.....	311
example configuration.....	346, 351
link state replication.....	319
link-layer overhead.....	329
link services protocols.....	841
link state replication	
LSQ PICs.....	319
link-layer overhead	
link services IQ interfaces.....	329
link-layer-overhead statement.....	375
usage guidelines.....	325, 329, 338
link-loss statement	
packet gateway.....	558
lmi-type statement.....	896
usage guidelines.....	862
load balancing	
on monitoring interfaces.....	707
load-balance statement.....	629
usage guidelines.....	613
local-certificate statement.....	260
usage guidelines.....	217
local-dump statement.....	744
usage guidelines.....	697
local-gateway address statement.....	300
usage guidelines.....	285
local-gateway statement.....	456
usage guidelines.....	439
local-id statement.....	260
usage guidelines.....	218
local-mac statement.....	630
usage guidelines.....	619
local-peer statement.....	630
usage guidelines.....	612, 614
log output	
adaptive services.....	447
traffic sampling.....	679

log-prefix statement.....	300, 457, 488
usage guidelines.....	286, 441, 471
logging statement.....	190, 457
usage guidelines.....	174
logical interfaces	
multicast-capable connections.....	848
logical tunnels.....	964
example configuration.....	972
logical-system statement.....	933
usage guidelines.....	911
loopback tunnels.....	967
LSQ bandwidth	
oversubscribing.....	331
LSQ failover	
interchassis.....	314
stateful intrachassis.....	316
stateless intrachassis.....	316
LSQ PICs.....	319
redundancy.....	316
lsq-failure-options statement.....	375
usage guidelines.....	314

M

manual security association.....	207
manual statement.....	261
usage guidelines.....	207
manuals	
comments on.....	lv
map statement.....	631
match direction usage in service sets.....	436
match-direction statement	
CoS.....	423
usage guidelines.....	412
IDS.....	190
usage guidelines.....	172
IPSec.....	261
usage guidelines.....	225
NAT.....	152
usage guidelines.....	128
stateful firewall.....	116
usage guidelines.....	103
max-burst-size statement	
gate in packet gateway.....	559
RTP/RTCP gate in packet gateway.....	560
max-concurrent-calls statement	
packet gateway.....	561
max-duplicates statement.....	833
usage guidelines.....	812
max-flows statement.....	458
usage guidelines.....	441
max-packets-per-second statement.....	745
usage guidelines.....	676
maximum-age statement.....	797
usage guidelines.....	779
maximum-connections statement.....	934

maximum-connections-per-client statement.....	934
maximum-contexts statement.....	405
usage guidelines.....	386
maximum-fuf-percentage statement.....	562
maximum-inactivity-time	
packet gateway.....	562
maximum-inactivity-time statement	
packet gateway.....	562
maximum-net-propagation-delay statement	
packet gateway.....	563
usage guidelines.....	508
maximum-send-window statement.....	301
usage guidelines.....	285
maximum-sessions statement.....	935
maximum-sessions-per-connection statement.....	935
maximum-terms statement	
packet gateway.....	564
maximum-waiting-delay statement	
packet gateway.....	565
usage guidelines.....	508
Media Gateway Controller	
definition.....	390
list configuration.....	391
media service	
packet gateway configuration.....	507
media statement.....	565
packet gateway.....	565
media-service statement	
packet gateway.....	566
usage guidelines.....	507
mediation devices	
flow-tap.....	818
mg-maximum-pdu-size statement	
packet gateway.....	566
mg-originated-pending-limit statement.....	567
mg-provisional-response-timer-value statement.....	567
mg-segmentation-timer statement.....	568
mgc-maximum-pdu-size statement	
packet gateway.....	568
mgc-originated-pending-limit statement.....	569
mgc-provisional-response-timer-value statement.....	569
mgc-segmentation-timer statement.....	570
minimum links	
link services interfaces.....	854
multilink interfaces.....	854
minimum-links statement.....	896
usage guidelines.....	852, 854
minimum-priority statement.....	833
usage guidelines.....	809
MLFR and MLPPP.....	845
mlfr-uni-nni-bundle-options statement.....	897
usage guidelines.....	859, 862
MLPPP.....	343, 353
configuration example.....	346
example configuration.....	356

mode statement.....	262
usage guidelines.....	216
monitor statement.....	570
packet gateway.....	570
monitoring statement.....	746
usage guidelines.....	684
moving-average-size statement.....	936
usage guidelines.....	912
MPLS	
packets	
passive flow monitoring.....	712
mpls-ipv4-template statement.....	747
mpls-template statement.....	747
mrru statement.....	898
usage guidelines.....	855
mss statement.....	191
usage guidelines.....	177
mtu statement.....	898
usage guidelines.....	855
multicast filters	
for IPv6 NAT.....	127
multicast traffic	
AS PIC.....	442
multicast tunnels.....	964
multicast-address statement	
usage guidelines.....	615
multicast-capable connections	
Frame Relay encapsulation.....	848
multicast-dlci statement.....	899
usage guidelines.....	848
multicast-only statement.....	980
usage guidelines.....	964
multiclass MLPPP	
fragmentation.....	330
multilink bundles	
fractional T1.....	353
example configuration.....	356, 357, 360
FRF.12.....	357
example configuration.....	360
MLPPP.....	353
example configuration.....	356
NxT1.....	343, 348
configuration example.....	346, 351
multilink interfaces	
example configuration.....	871
minimum links.....	854
multilink protocols.....	841
multilink-class statement.....	376
usage guidelines.....	330
multilink-max-classes statement.....	376
usage guidelines.....	330
multiservice-options statement.....	748
MultiServices PIC	
hardware requirements.....	27

N

n391 statement.....	899
usage guidelines.....	862
n392 statement.....	900
usage guidelines.....	862
n393 statement.....	900
usage guidelines.....	862
name-format statement.....	798
usage guidelines.....	778
NAT	
action statements.....	131
address configuration.....	124
applications.....	130
example configuration.....	133
match conditions.....	130
rule sets.....	127
twice NAT	
description.....	39
type.....	129
nat-pool statement	
packet gateway.....	571
usage guidelines.....	507
nat-rule-sets statement.....	458
usage guidelines.....	438
nat-rules statement.....	458
usage guidelines.....	438
nat-type statement.....	152
usage guidelines.....	129
neighbor discovery using IPv6 NAT.....	127
network-operator-id	
packet gateway.....	571
network-operator-id statement	
packet gateway.....	571
next-hop groups.....	698
next-hop statement.....	748
next-hop groups	
usage guidelines.....	700
usage guidelines.....	698
next-hop style service sets.....	437
next-hop-group statement.....	749
usage guidelines.....	698, 700
next-hop-service statement.....	459
usage guidelines.....	434
no-anti-replay statement.....	262
usage guidelines.....	229
no-core-dump statement.....	725
usage guidelines.....	683
no-dscp-bit-mirroring	
packet gateway.....	574
no-dscp-bit-mirroring statement	
packet gateway.....	574
no-filter-check statement.....	750
usage guidelines.....	699
no-fragmentation statement.....	377
usage guidelines.....	328

no-local-dump statement.....	744
usage guidelines.....	697
no-preempt statement.....	631, 633
usage guidelines.....	619
no-stamp statement.....	766
usage guidelines.....	678
no-syslog statement.....	767, 834
usage guidelines.....	811
no-termination-request statement.....	377
usage guidelines.....	314
no-translation statement.....	153
usage guidelines.....	131
no-world-readable statement.....	774
usage guidelines.....	678
normal-mg-execution-time statement.....	572
normal-mgc-execution-time statement.....	572
notice (system logging severity level).....	287, 442, 472
notice icons defined.....	xlvi
Notification Behavior.....	573
notification-behavior statement.....	573
packet gateway.....	573
notification-rate-limit statement	
packet gateway.....	573
notification-regulation statement	
packet gateway.....	574
notification-targets statement.....	834
usage guidelines.....	809
NxT1 bundles	
FRF.16.....	348
configuration example.....	351
MLPPP.....	343
configuration example.....	346
O	
one-way-hardware-timestamp statement.....	936
usage guidelines.....	918, 920
open-timeout statement.....	488
usage guidelines.....	470
option-refresh-rate statement.....	751
output files	
logging information output file.....	679
traffic sampling output files.....	677
output statement.....	489, 752
usage guidelines.....	434, 473
output-interface-index statement.....	755
outside-service-interface statement	
usage guidelines.....	434
overload control	
packet gateway configuration.....	514
overload-control statement	
packet gateway.....	575
overload-pool statement.....	153
usage guidelines.....	131
overload-prefix statement.....	154
usage guidelines.....	131

oversubscription.....	331
-----------------------	-----

P

packet gateway	
configuration.....	503
hint.....	151
packet gateway controller (PGC)	
gateway controller configuration.....	506
parentheses, in syntax descriptions.....	xlvi
passive flow monitoring.....	667
MPLS packets.....	712
passive-monitor-mode statement.....	755
usage guidelines.....	711
password statement.....	800
usage guidelines.....	777, 779
peak-data-rate statement	
gate in the packet gateway.....	576
packet gateway.....	576
RTP/RTCP gate in the packet gateway.....	577
peer statement.....	632
peer-unit statement.....	981
usage guidelines.....	964
per-unit-scheduler statement.....	378
usage guidelines.....	853
perfect-forward-secrecy statement.....	263
usage guidelines.....	223
performance, monitoring.....	912
pgcp statement	
NAT.....	154
pgcp-rule-sets statement	
usage guidelines.....	438
pgcp-rules statement	
service-set.....	460
usage guidelines.....	438
PIC types for services.....	3
pic-memory-threshold statement.....	835
usage guidelines.....	811
PIM	
tunnels.....	968
PIR.....	331
platforms, supported.....	4
point-to-point connections	
Frame Relay encapsulation.....	848
policy statement	
IKE.....	264
IPSec.....	264
policy statement (IKE)	
usage guidelines.....	215
policy statement (IPSec)	
usage guidelines.....	222
pool statement.....	155
service interface pool.....	606
usage guidelines.....	124
pop-all-labels statement.....	756
usage guidelines.....	713

port mirroring.....	698
port statement	
cflowd	
usage guidelines.....	688
flow monitoring.....	757
NAT.....	156
usage guidelines.....	124
RPM.....	937
TWAMP.....	937
voice services.....	405
usage guidelines.....	386
port-mirroring statement.....	758
usage guidelines.....	699
ports-per-session statement	
packet gateway.....	156
post-service-filter statement.....	489
usage guidelines.....	434
ppp-access-profile statement.....	301
usage guidelines.....	284
ppp-profile statement	
usage guidelines.....	284
pre-shared-key statement.....	265
usage guidelines.....	216
preempt statement.....	633
usage guidelines.....	619
preserve-interface statement.....	378
usage guidelines.....	319
primary statement.....	379, 490
usage guidelines.....	476
priority statement.....	634
usage guidelines.....	619
probe statement.....	938
usage guidelines.....	912
probe-count statement.....	939
usage guidelines.....	912
probe-interval statement.....	939
usage guidelines.....	912
probe-limit statement.....	940
usage guidelines.....	917
probe-server statement.....	940
usage guidelines.....	917
probe-type statement.....	941
usage guidelines.....	912
probes, for monitoring traffic.....	912
procedural overview.....	33
promiscuous statement	
usage guidelines.....	612
proposal statement (IKE).....	266
usage guidelines.....	212
proposal statement (IPSec).....	267
usage guidelines.....	220
proposals statement (IKE).....	267
usage guidelines.....	216
proposals statement (IPSec).....	267
usage guidelines.....	223

protocol statement	
applications.....	97
usage guidelines.....	56
IPSec.....	268
usage guidelines.....	208, 222
protocols statement	
DLSW.....	635

Q

queue-limit-percentage statement	
packet gateway.....	578
usage guidelines.....	514
queues statement.....	406
usage guidelines.....	386

R

random-allocation statement.....	156
rate statement.....	759
usage guidelines.....	676, 699
reachability-cache-timeout statement.....	635
usage guidelines.....	614
Real-Time Performance Monitoring <i>See</i> RPM	
Real-Time Transport Protocol.....	856, 857
example configuration.....	857, 858
reassemble-packets statement.....	981
usage guidelines.....	963
receive-initial-pacing statement	
usage guidelines.....	615
receive-options-packets statement.....	759
usage guidelines.....	711
receive-ttl-exceeded statement.....	760
usage guidelines.....	711
receive-window statement.....	302
usage guidelines.....	285
reconnect statement	
packet gateway.....	578
red-differential-delay statement.....	901
usage guidelines.....	861
redundancy	
AS PIC.....	475
flow monitoring.....	716
L2TP.....	289
redundancy-group statement.....	637
redundancy-options statement.....	379, 490
usage guidelines.....	476
reflexive reverse statement.....	424
usage guidelines.....	414
remote-gateway statement.....	268
usage guidelines.....	228
remote-id statement.....	269
usage guidelines.....	218
remote-mac statement.....	637
remote-peer statement.....	638
usage guidelines.....	612

remotely-controlled statement.....	157
report-service-change statement.....	579
request-timestamp statement.....	579
required-depth statement.....	760
usage guidelines.....	713
retransmit-interval statement.....	302
usage guidelines.....	286
retry statement.....	801
usage guidelines.....	780, 812
retry-delay statement.....	801
usage guidelines.....	780, 812
RFC 2890.....	961
router identifier.....	983
routing-instance statement.....	941, 982
packet gateway.....	580
usage guidelines.....	507
usage guidelines.....	966
routing-instances statement.....	942, 982
usage guidelines.....	912
routing-options statement.....	983
rpc-program-number statement.....	98
usage guidelines.....	63
RPM.....	907, 909
example configuration.....	923
rpm statement.....	942
usage guidelines.....	918
rtcp statement	
media in packet gateway.....	580
rtp statement.....	406, 581
packet gateway.....	581
usage guidelines.....	386
rule content	
packet gateway configuration.....	513
rule set	
packet gateway configuration.....	513
rule statement	
CoS.....	425
usage guidelines.....	410
IDS.....	192
usage guidelines.....	171
IPSec.....	270
usage guidelines.....	225
NAT.....	158
usage guidelines.....	128
packet gateway.....	581
usage guidelines.....	513
stateful firewall.....	117
usage guidelines.....	102
rule-set statement	
CoS.....	426
usage guidelines.....	410
IDS.....	193
usage guidelines.....	171
IPSec.....	271
usage guidelines.....	224
NAT.....	159
usage guidelines.....	127
packet gateway.....	582
usage guidelines.....	513
stateful firewall.....	118
usage guidelines.....	102
run-length statement.....	761
usage guidelines.....	676, 699
S	
sample (firewall filter action).....	675
sampled file.....	679
sampled.pkts file.....	677
sampling	
logical interface.....	676
monitoring interface.....	683
sampling rate.....	676
sampling statement.....	762
usage guidelines.....	675
scheduler map	
CoS	
configuration example.....	853
SDK applications	
service set extensions.....	443
SDK services	
enabling.....	445
secondary statement.....	380, 491
usage guidelines.....	476
security associations	
clearing.....	211
segmentation	
packet gateway.....	582
segmentation statement	
packet gateway.....	582
send cflowd records to flow collector.....	780
send-notification-on-delay statement	
packet gateway.....	583
server statement.....	943
service filters.....	474
service interface configuration.....	434
service packages.....	29
service rules configuration.....	438
service sets	
example configuration.....	449
overview.....	28
service statement.....	491
usage guidelines.....	473
service-change statement	
packet gateway.....	584
service-change-type statement	
packet gateway.....	585
service-domain statement.....	492
usage guidelines.....	436

- service-filter statement
 - firewall
 - usage guidelines.....474
 - interfaces.....492
 - usage guidelines.....434
- service-interface
 - packet gateway.....585
- service-interface statement.....303, 460
 - packet gateway.....585
 - usage guidelines.....285, 434
- service-interface-pools statement
 - service interface pool.....606
- service-port statement.....835
 - usage guidelines.....809
- service-set statement.....461, 493
 - usage guidelines.....433, 473
- service-state statement.....587
 - packet gateway.....586
 - virtual interface in the packet gateway.....587
 - VPF in the packet gateway.....586
- services configuration overview.....33
- services PICs.....3
- services statement
 - CoS.....426
 - usage guidelines.....409
 - DFC.....836
 - usage guidelines.....806
 - flow monitoring
 - usage guidelines.....674
 - flow-monitoring.....764
 - IDS.....193
 - usage guidelines.....169
 - interfaces.....494
 - usage guidelines.....471
 - IPSec.....271
 - usage guidelines.....203
 - L2TP
 - usage guidelines.....286
 - NAT.....159
 - usage guidelines.....123
 - packet gateway.....587
 - RPM.....943
 - usage guidelines.....909
 - service sets.....462
 - usage guidelines.....441
 - stateful firewall.....118, 304
 - usage guidelines.....101
- services-options statement.....495
 - usage guidelines.....286, 470, 471
- session-limit statement.....194
 - usage guidelines.....176
- session-mirroring
 - packet gateway.....588
- session-mirroring statement
 - packet gateway.....588
- shaping-rate statement
 - usage guidelines.....853
- shared-key statement.....836
 - usage guidelines.....809
- short-sequence statement.....901
 - usage guidelines.....856
- SIP configuration.....62
- sip-call-hold-timeout statement.....98
 - usage guidelines.....62
- sip-text statement.....427
 - usage guidelines.....414
- sip-video statement.....427
 - usage guidelines.....414
- sip-voice statement.....428
 - usage guidelines.....414
- size statement.....765
 - usage guidelines.....679
- snmp-command statement.....99
 - usage guidelines.....63
- soft-limit statement.....837
 - usage guidelines.....808
- soft-limit-clear statement.....837
 - usage guidelines.....808
- SONET interfaces
 - sampling SONET interfaces.....680
- source statement.....662, 983
 - usage guidelines.....655, 959
- source-address
 - packet gateway.....588
- source-address statement
 - CoS.....428
 - usage guidelines.....412
 - flow monitoring.....766
 - usage guidelines.....684
 - IDS.....195
 - usage guidelines.....173
 - IPSec.....272
 - usage guidelines.....226
 - NAT.....160
 - usage guidelines.....130
 - packet gateway.....588
 - RPM.....944
 - usage guidelines.....912
 - stateful firewall.....119
 - usage guidelines.....103
 - tunnel services.....984
 - usage guidelines.....969
- source-address-range statement
 - IDS.....195
 - usage guidelines.....173
 - NAT.....160
 - usage guidelines.....130
 - stateful firewall.....119
 - usage guidelines.....103
- source-addresses statement.....838
 - usage guidelines.....809

source-pool statement.....	161
usage guidelines.....	131
source-port	
packet gateway.....	589
source-port statement.....	99
packet gateway.....	589
RPM	
usage guidelines.....	59
source-prefix statement.....	161, 196
usage guidelines.....	175
source-prefix-ipv6 statement.....	196
usage guidelines.....	175
source-prefix-list statement	
CoS.....	429
IDS.....	197
NAT.....	162
stateful firewall.....	120
usage guidelines.....	103
spi statement.....	272
usage guidelines	209
stamp option.....	679
stamp statement.....	766
usage guidelines.....	678
state-loss statement	
packet gateway.....	589
stateful firewall	
action statements.....	104
anomalies.....	36
applications.....	103
example configuration.....	106
match conditions.....	103
rules.....	102
stateful-firewall-rule-sets statement.....	463
usage guidelines.....	438
stateful-firewall-rules statement.....	463
usage guidelines.....	438
statement-name statement.....	522, 540
stop-detection-on-drop statement	
packet gateway.....	590
support, technical <i>See</i> technical support	
sustained-data-rate statement	
gate in packet gateway.....	591
RTP/RTCP gate in packet gateway.....	592
syn-cookie statement.....	197
usage guidelines.....	174
syntax conventions.....	xlvi
syslog statement	
CoS.....	429
usage guidelines.....	413
flow monitoring.....	767
IDS.....	198
usage guidelines.....	175
interfaces.....	496
usage guidelines.....	471
IPSec.....	273
usage guidelines.....	230

L2TP	
usage guidelines.....	286
NAT.....	162
usage guidelines.....	131
service sets.....	464
usage guidelines.....	441
stateful firewall.....	120
usage guidelines.....	104

T

t391 statement.....	902
usage guidelines.....	862
t392 statement.....	902
usage guidelines.....	862
target statement.....	944
target-url statement	
usage guidelines.....	912
tcp statement.....	945
technical support	
contacting JTAC.....	lv
telephony products for J-series Services Routers.....	390
template statement.....	768
template-refresh-rate statement.....	769
term statement	
CoS.....	430
usage guidelines.....	410
IDS.....	199
usage guidelines.....	171
IPSec.....	274
usage guidelines.....	225
NAT.....	163
usage guidelines.....	128
stateful firewall.....	121
usage guidelines.....	102
test statement.....	946
usage guidelines.....	912
test-interval statement.....	947
usage guidelines.....	912
TGM550 VoIP product.....	390
then statement	
CoS.....	431
usage guidelines.....	410, 411
IDS.....	201
usage guidelines.....	172
IPSec.....	275
usage guidelines.....	225
NAT.....	164
usage guidelines.....	128
stateful firewall.....	122
usage guidelines.....	102, 103
threshold statement.....	202
usage guidelines.....	175, 177
thresholds statement.....	948
usage guidelines.....	912
TIM VoIP product.....	390

- time-to-live threshold.....63
- timex statement
 - packet gateway.....593
- timestamp option.....679
- tmax-retransmission-delay statement
 - packet gateway.....595
 - usage guidelines.....508
- trace-options
 - server (tracing flag).....448
 - timer-events (tracing flag).....448
- traceoptions statement
 - DLSw
 - usage guidelines.....617
 - flow monitoring.....770
 - usage guidelines.....679
 - IPSec.....276
 - usage guidelines.....236
 - L2TP.....307
 - usage guidelines.....290
 - packet gateway.....596
 - usage guidelines.....516
 - services.....465
 - usage guidelines.....780
- tracing flags
 - event policy
 - all.....448
 - configuration.....448
 - database.....448
 - events.....448
 - policy.....448
 - server.....448
 - timer-events.....448
- tracing operations
 - adaptive services.....446
 - DLSw.....617
- track statement.....641
- traffic.....649
 - inbound (decryption).....653
 - IPSec, configuring.....649
 - monitoring.....912
 - outbound (encryption).....651
- traffic sampling
 - configuring.....675
 - disabling.....677, 726
 - example configurations.....680
 - flow aggregation.....687
 - FTP traffic.....682
 - logging information output file.....679
 - output files.....677
 - SONET interfaces.....680
 - traffic from single IP addresses.....681
- traffic-control-profiles statement
 - usage guidelines.....331, 335
- traffic-management statement
 - packet gateway.....594
- transfer statement.....802
 - usage guidelines.....778
- transfer-log-archive statement.....802
 - usage guidelines.....779
- translated statement.....165
 - usage guidelines.....131
- translation-type statement.....166
 - usage guidelines.....131
- transport statement
 - NAT.....167
- traps statement.....949
 - usage guidelines.....912
- trigger-link-failure statement.....380
 - usage guidelines.....314
- trusted-ca statement.....466
 - usage guidelines.....440
- ttl statement.....984
 - DFC.....838
 - usage guidelines.....808
 - usage guidelines.....959
- ttl-threshold statement.....100
 - usage guidelines.....63
- tunnel interfaces
 - configuration statements.....959, 964, 967
 - dynamic tunnels.....969
 - example configuration.....970
 - logical tunnels.....964
 - loopback tunnels.....967
 - multicast tunnels.....964
 - PIM tunnels.....968
 - unicast tunnels.....959
- tunnel statement.....663, 985
 - encryption
 - usage guidelines.....647
 - redundancy
 - usage guidelines.....655
 - unicast
 - usage guidelines.....959
- tunnel-group statement.....309
 - usage guidelines.....284
- tunnel-mtu statement.....277
 - usage guidelines.....230
- tunnel-timeout statement.....310
 - usage guidelines.....286
- tunnel-type statement.....986
 - usage guidelines.....969
- tunnels
 - definition.....955
 - GRE
 - fragmentation of.....962
 - key number.....961
 - interface types.....955
- twamp statement.....950
- twamp-server statement.....950
- twice NAT.....39

type-of-service statement.....	642
usage guidelines.....	615

U

udp statement.....	951
unicast tunnels.....	959
unit statement	
encryption.....	664
usage guidelines.....	647
flow monitoring.....	771
usage guidelines.....	675
interfaces.....	497
usage guidelines.....	467
link services.....	407, 903
usage guidelines.....	845
tunnel.....	987
usage guidelines.....	959
Universal Unique Identifier.....	64
up statement	
packet gateway.....	597
username statement.....	803
usage guidelines.....	779
uuid statement.....	100
usage guidelines.....	64

V

var/log/sampled file.....	679
var/tmp/sampled.pkts file.....	677
variant statement.....	803
usage guidelines.....	778
version statement.....	772
usage guidelines.....	688
version9 statement.....	773
virtual interface	
packet gateway configuration.....	507
virtual loopback tunnel	
configuration guidelines.....	966
VRF table lookup	
example configuration.....	971
virtual-interface statement	
packet gateway.....	598
usage guidelines.....	507
virtual-interface-down statement	
packet gateway.....	599
virtual-interface-indications statement	
packet gateway.....	600
virtual-interface-up statement	
packet gateway.....	601
voice services	
bundles.....	389
encapsulation.....	388
interface type.....	385
voice services configuration	
usage guidelines.....	384

voice services interfaces	
interleave fragments.....	387
VoIP	
configuration on J-series Services Routers.....	390
interface configuration.....	390

W

warm standby	
AS PIC.....	475
LSQ PIC.....	316
warm statement	
packet gateway.....	601
warm-standby statement.....	381
warning (system logging severity	
level).....	287, 442, 472
wildcard-response-service-change	
packet gateway.....	602
wildcard-response-service-change statement	
packet gateway.....	602
world-readable statement.....	774
usage guidelines.....	678

Y

yellow-differential-delay statement.....	904
usage guidelines.....	861

Index of Statements and Commands

A

accounting statement.....	720
acknowledge-retries statement.....	885
acknowledge-timer statement.....	886
action-red-differential-delay statement.....	886
activation-priority statement.....	397
adaptive-services-pics statement.....	451
address statement	
DFC.....	823
encryption.....	657
flow monitoring.....	721
interfaces.....	481
link services.....	887
NAT.....	145
voice services.....	398
address-range statement	
NAT.....	146
administrative statement	
packet gateway.....	523
advertise-interval statement.....	623
aggregate-export-interval statement.....	721
aggregation statement.....	181, 722
algorithm statement	
packet gateway.....	523
allow-fragmentation statement.....	975
allow-ip-option statement.....	112
allow-multicast statement.....	452
allowed-destinations statement.....	824
analyzer-address statement.....	787
analyzer-id statement.....	787
application statement.....	91
application-data-inactivity-detection statement	
packet gateway.....	524
application-profile statement.....	419
application-protocol statement.....	92
application-set statement.....	93
application-sets statement	
CoS.....	420
IDS.....	182
NAT.....	146
stateful firewall.....	113
applications statement	
applications hierarchy.....	93
CoS.....	420

IDS.....	182
NAT.....	147
stateful firewall.....	113
archive-sites statement.....	788
audit-observed-events-returns-history statement	
packet gateway.....	524
authentication statement.....	245
authentication-algorithm statement	
IKE.....	246
IPSec.....	246
authentication-method statement	247
autonomous-system-type statement.....	723
auxiliary-spi statement.....	247

B

backup-destination statement.....	658, 976
backup-interface statement.....	658
backup-remote-gateway statement	248
base-root statement	
packet gateway.....	525
bearer-bandwidth-limit statement.....	399
bundle statement.....	399, 887
by-destination statement.....	183
by-pair statement.....	184
by-source statement.....	185

C

cancel-graceful statement	
packet gateway.....	526, 527
capture-group statement.....	825
cflowd statement.....	724
circuit-weight statement.....	624
cisco-interoperability statement.....	371
cleanup-timeout statement	
packet gateway.....	527
clear-dont-fragment-bit statement	
GRE tunnel.....	482
IPSec.....	248
clear-ike-sas-on-pic-restart statement.....	249
clear-ipsec-sas-on-pic-restart statement.....	249
client-list statement.....	927
collector statement.....	788
compression statement.....	400
compression-device statement.....	400, 888

connection-idle-timeout statement.....	624
content-destination statement.....	826
context-indications statement	
packet gateway.....	528
control-association-indications statement	
packet gateway.....	529
control-source statement.....	827
controller-address statement	
packet gateway.....	530
controller-failure statement	
packet gateway.....	530
controller-port statement	
packet gateway.....	531
copy-tos-to-outer-ip-header statement.....	976
core-dump statement.....	725
cost statement.....	625

D

data-fill statement.....	928
data-format statement.....	789
data-inactivity-detection statement	
packet gateway.....	531
data-size statement.....	928
delivery-function statement	
packet gateway.....	532
description statement	
IPSec	250
destination statement	
DLSw.....	625
encryption.....	659
flow monitoring.....	726
link services.....	888
tunnel.....	977
destination-address statement	
CoS.....	421
IDS.....	186
IPSec.....	250
NAT.....	147
packet gateway.....	532
stateful firewall.....	114
destination-address-range statement	
IDS.....	186
NAT.....	148
stateful firewall.....	114
destination-interface statement.....	626, 929
destination-pool statement.....	148
destination-port statement	
applications.....	93
packet gateway.....	533
destination-prefix statement.....	149, 187
destination-prefix-ipv6 statement.....	187
destination-prefix-list statement	
CoS.....	421
IDS.....	188

NAT.....	149
stateful firewall.....	115
destinations statement.....	789
detect statement	
packet gateway.....	533
dh-group statement	251
dial-options statement.....	483
diffserv statement.....	534
direction statement.....	252
disable statement.....	726
disable-mlppp-inner-ppp-pfc statement.....	889
disable-session-mirroring statement	
packet gateway.....	534
disconnect statement	
packet gateway.....	535
dlci statement.....	889
dlsd statement.....	626
dlsd-cos statement.....	627
do-not-fragment statement.....	978
down statement	
packet gateway.....	536
drop-timeout statement.....	890
dscp statement.....	422, 537
dscp-code-point statement.....	931
duplicates-dropped-periodicity statement.....	827
dynamic route insertion.....	232
dynamic statement	253
dynamic-call-admission-control statement.....	401
dynamic-flow-capture statement.....	828
dynamic-tunnels statement.....	979

E

encapsulation statement.....	402, 891
encoding statement	
packet gateway.....	537
encryption statement	254
encryption-algorithm statement (IPSec).....	255
engine-id statement.....	727
engine-type statement.....	727
es-options statement.....	659
event-timestamp-notification statement	
packet gateway.....	538
explorer-wait-time statement.....	627
export-format statement.....	728
extension-service statement.....	452

F

f-max-period statement.....	402
facility-override statement.....	297, 453, 484
failover-cold statement	
packet gateway.....	538
failover-warm statement	
packet gateway.....	539

failure statement	
packet gateway.....	541
family statement	
encryption.....	660
flow monitoring.....	729
interfaces.....	485
link services.....	403, 893
fast-update-filters statement	
packet gateway.....	541
file statement.....	732, 733
file-specification statement.....	791
filename statement.....	733
filename-prefix statement.....	790
files statement.....	734
filter statement	
encryption.....	661
flow monitoring.....	735
flow-active-timeout statement.....	736
flow-collector statement.....	792
flow-export-destination statement.....	737
flow-inactive-timeout statement.....	737
flow-monitoring statement.....	738
flow-tap statement.....	829
force-entry statement.....	188
forwarding-class statement.....	372, 422
forwarding-options statement.....	738
fragment-threshold statement	
link services.....	894
LSQ.....	373
voice services.....	404
fragmentation-map statement.....	373
fragmentation-maps statement.....	374
from statement	
CoS.....	423
IDS.....	189
IPSec.....	256
NAT.....	150
stateful firewall.....	115
ftp statement.....	794

G

g-duplicates-dropped-periodicity statement.....	829
g-max-duplicates statement.....	830
gateway statement	
packet gateway.....	542
gateway-address statement	
packet gateway.....	545
gateway-controller statement	
packet gateway.....	546
gateway-port statement	
packet gateway.....	546
graceful statement	
packet gateway.....	547, 548
graceful-restart statement	
packet gateway.....	548, 563, 564

H

h248-options statement	
packet gateway.....	549
h248-properties statement.....	551
h248-timers statement	
packet gateway.....	552
hanging-termination-detection statement	
packet gateway.....	553
hard-limit statement.....	830
hard-limit-target statement.....	831
hardware-timestamp statement.....	932
hello-interval statement.....	298
hello-timer statement.....	894
hide-avps statement.....	298
hint statement.....	151
history-size statement.....	932
hold-time statement.....	628
host statement.....	299, 453, 486
hot-standby statement.....	374

I

icmp-code statement.....	94
icmp-type statement.....	95
ids-rules statement.....	454
ignore-entry statement.....	188, 189
ike statement.....	257
ike-access-profile statement	454
inactivity-delay statement	
packet gateway.....	554
inactivity-duration statement	
packet gateway.....	554
inactivity-timeout statement.....	95, 486, 933
packet gateway.....	553
inactivity-timer statement	
packet gateway.....	555
initial-average-ack-delay statement	
packet gateway.....	555
initiate-dead-peer-detection statement.....	258
input statement	
flow monitoring.....	739
interfaces.....	487
input-interface-index statement.....	740
input-packet-rate-threshold statement.....	831
interface statement.....	741
DLSw.....	629
flow-tap.....	832
packet gateway.....	556
service interface pool.....	605
interface-map statement.....	796
interface-service statement.....	455
interfaces statement	
DFC.....	832
encryption.....	661
flow monitoring.....	743
interfaces hierarchy.....	487

link services.....	404, 895
tunnel.....	979
interim-ah-scheme statement	
packet gateway.....	556
interleave-fragments statement.....	895
ipsec statement.....	258
ipsec-inside-interface statement.....	259
ipsec-sa statement.....	662
ipsec-vpn-options statement.....	455
ipsec-vpn-rules statement.....	456
ipv4-template statement.....	743
ipv6-multicast-interfaces statement.....	151

K

key statement.....	980
--------------------	-----

L

l2tp-access-profile statement.....	299
label-position statement.....	744
latch-deadlock-delay statement	
packet gateway.....	557
learn-sip-register statement.....	96
lifetime-seconds statement (IKE).....	259
link-layer-overhead statement.....	375
link-loss statement	
packet gateway.....	558
lmi-type statement.....	896
load-balance statement.....	629
local-certificate statement.....	260
local-dump statement.....	744
local-gateway address statement.....	300
local-gateway statement.....	456
local-id statement.....	260
local-mac statement.....	630
local-peer statement.....	630
log-prefix statement.....	300, 457, 488
logging statement.....	190, 457
logical-system statement.....	933
lsq-failure-options statement.....	375

M

manual statement.....	261
map statement.....	631
match-direction statement	
CoS.....	423
IDS.....	190
IPSec.....	261
NAT.....	152
stateful firewall.....	116
max-burst-size statement	
gate in packet gateway.....	559
max-concurrent-calls statement	
packet gateway.....	561

max-duplicates statement.....	833
max-flows statement.....	458
max-packets-per-second statement.....	745
maximum-age statement.....	797
maximum-connections statement.....	934
maximum-connections-per-client statement.....	934
maximum-contexts statement.....	405
maximum-fuf-percentage statement.....	562
maximum-inactivity-time statement	
packet gateway.....	562
maximum-net-propagation-delay statement	
packet gateway.....	563
maximum-send-window statement.....	301
maximum-sessions statement.....	935
maximum-sessions-per-connection statement.....	935
maximum-terms statement	
packet gateway.....	564
maximum-waiting-delay statement	
packet gateway.....	565
media statement.....	565
media-service statement	
packet gateway.....	566
mg-maximum-pdu-size statement	
packet gateway.....	566
mg-originated-pending-limit statement.....	567
mg-provisional-response-timer-value statement.....	567
mg-segmentation-timer statement.....	568
mgc-maximum-pdu-size statement	
packet gateway.....	568
mgc-originated-pending-limit statement.....	569
mgc-provisional-response-timer-value statement.....	569
mgc-segmentation-timer statement.....	570
minimum-links statement.....	896
minimum-priority statement.....	833
mlfr-uni-nni-bundle-options statement.....	897
mode statement.....	262
monitor statement.....	570
monitoring statement.....	746
moving-average-size statement.....	936
mpls-ipv4-template statement.....	747
mpls-template statement.....	747
mrru statement.....	898
mss statement.....	191
mtu statement.....	898
multicast-dlci statement.....	899
multicast-only statement.....	980
multilink-class statement.....	376
multilink-max-classes statement.....	376
multiservice-options statement.....	748

N

n391 statement.....	899
n392 statement.....	900
n393 statement.....	900
name-format statement.....	798

nat-pool statement	
packet gateway.....	571
nat-rules statement.....	458
nat-type statement.....	152
network-operator-id statement	
packet gateway.....	571
next-hop statement.....	748
next-hop-group statement.....	749
next-hop-service statement.....	459
no-anti-replay statement	262
no-dscp-bit-mirroring statement	
packet gateway.....	574
no-filter-check statement.....	750
no-fragmentation statement.....	377
no-preempt statement.....	631, 633
no-stamp statement.....	766
no-syslog statement.....	767, 834
no-termination-request statement.....	377
no-translation statement.....	153
no-world-readable statement.....	774
normal-mg-execution-time statement.....	572
normal-mgc-execution-time statement.....	572
notification-behavior statement	
packet gateway.....	573
notification-rate-limit statement	
packet gateway.....	573
notification-regulation statement	
packet gateway.....	574
notification-targets statement.....	834

O

one-way-hardware-timestamp statement.....	936
open-timeout statement.....	488
option-refresh-rate statement.....	751
output statement.....	489, 752
output-interface-index statement.....	755
overload-control statement	
packet gateway.....	575
overload-pool statement.....	153
overload-prefix statement.....	154

P

passive-monitor-mode statement.....	755
password statement.....	800
peak-data-rate statement	
packet gateway.....	576
peer statement.....	632
peer-unit statement.....	981
per-unit-scheduler statement.....	378
perfect-forward-secrecy statement	263
pgcp statement	
NAT.....	154
pgcp-rules statement	
service-set.....	460

pic-memory-threshold statement.....	835
policy statement	
IKE	264
pool statement.....	155
service interface pool.....	606
pop-all-labels statement.....	756
port statement	
flow monitoring.....	757
NAT.....	156
RPM.....	937
TWAMP.....	937
voice services.....	405
port-mirroring statement.....	758
ports-per-session statement	
packet gateway.....	156
post-service-filter statement.....	489
ppp-access-profile statement.....	301
pre-shared-key statement	265
preserve-interface statement.....	378
primary statement.....	379, 490
priority statement.....	634
probe statement.....	938
probe-count statement.....	939
probe-interval statement.....	939
probe-limit statement.....	940
probe-server statement.....	940
probe-type statement.....	941
proposal statement (IKE).....	266
proposal statement (IPSec).....	267
proposals statement (IKE).....	267
protocol statement	
applications.....	97
IPSec.....	268
protocols statement	
DLSW.....	635

Q

queue-limit-percentage statement	
packet gateway.....	578
queues statement.....	406

R

rate statement.....	759
reachability-cache-timeout statement.....	635
reassemble-packets statement.....	981
receive-options-packets statement.....	759
receive-ttl-exceeded statement.....	760
receive-window statement.....	302
reconnect statement	
packet gateway.....	578
red-differential-delay statement.....	901
redundancy-group statement.....	637
redundancy-options statement.....	379, 490
reflexive reverse statement.....	424

remote-gateway statement	268	service-state statement.....	587
remote-id statement	269	packet gateway.....	586
remote-mac statement.....	637	services statement	
remote-peer statement.....	638	CoS.....	426
remotely-controlled statement.....	157	DFC.....	836
request-timestamp statement.....	579	flow-monitoring.....	764
required-depth statement.....	760	IDS.....	193
retransmit-interval statement.....	302	interfaces.....	494
retry statement.....	801	IPSec.....	271
retry-delay statement.....	801	NAT.....	159
routing-instance statement.....	941, 982	packet gateway.....	587
packet gateway.....	580	RPM.....	943
routing-instances statement	942, 982	service sets.....	462
routing-options statement.....	983	stateful firewall.....	118, 304
rpc-program-number statement.....	98	services-options statement.....	495
rpm statement	942	session-limit statement.....	194
rtp statement.....	406, 581	session-mirroring statement	
rule statement		packet gateway.....	588
CoS.....	425	shared-key statement.....	836
IDS.....	192	short-sequence statement.....	901
IPSec.....	270	sip-call-hold-timeout statement.....	98
NAT.....	158	sip-text statement.....	427
packet gateway.....	581	sip-video statement.....	427
stateful firewall.....	117	sip-voice statement.....	428
rule-set statement		size statement.....	765
CoS.....	426	snmp-command statement.....	99
IDS.....	193	soft-limit statement.....	837
IPSec.....	271	soft-limit-clear statement.....	837
NAT.....	159	source statement.....	662, 983
packet gateway.....	582	source-address statement	
stateful firewall.....	118	CoS.....	428
run-length statement.....	761	flow monitoring.....	766
		IDS.....	195
		IPSec.....	272
		NAT.....	160
		packet gateway.....	588
		RPM.....	944
		stateful firewall.....	119
		tunnel services.....	984
		source-address-range statement	
		IDS.....	195
		NAT.....	160
		stateful firewall.....	119
		source-addresses statement.....	838
		source-pool statement.....	161
		source-port statement.....	99
		packet gateway.....	589
		source-prefix statement.....	161, 196
		source-prefix-ipv6 statement.....	196
		source-prefix-list statement	
		CoS.....	429
		IDS.....	197
		NAT.....	162
		stateful firewall.....	120
		spi statement	272

S

sampling statement.....	762
secondary statement.....	380, 491
segmentation statement	
packet gateway.....	582
send-notification-on-delay statement	
packet gateway.....	583
server statement.....	943
service statement.....	491
service-change statement	
packet gateway.....	584
service-change-type statement	
packet gateway.....	585
service-domain statement.....	492
service-filter statement	
interfaces.....	492
service-interface statement.....	303, 460
packet gateway.....	585
service-interface-pools statement	
service interface pool.....	606
service-port statement.....	835
service-set statement.....	461, 493

state-loss statement	
packet gateway.....	589
stateful-firewall-rules statement.....	463
statement-name statement.....	522, 540
stop-detection-on-drop statement	
packet gateway.....	590
sustained-data-rate statement	
gate in packet gateway.....	591
RTP/RTCP gate in packet gateway.....	592
syn-cookie statement.....	197
syslog statement	
CoS.....	429
flow monitoring.....	767
IDS.....	198
interfaces.....	496
IPSec.....	273
NAT.....	162
service sets.....	464
stateful firewall.....	120

T

t391 statement.....	902
t392 statement.....	902
target statement.....	944
tcp statement.....	945
template statement.....	768
template-refresh-rate statement.....	769
term statement	
CoS.....	430
IDS.....	199
IPSec.....	274
NAT.....	163
stateful firewall.....	121
test statement.....	946
test-interval statement.....	947
then statement	
CoS.....	431
IDS.....	201
IPSec.....	275
NAT.....	164
stateful firewall.....	122
threshold statement.....	202
thresholds statement.....	948
timerx statement	
packet gateway.....	593
tmax-retransmission-delay statement	
packet gateway.....	595
traceoptions statement	
flow monitoring.....	770
IPSec.....	276
L2TP.....	307
packet gateway.....	596
services.....	465
track statement.....	641

traffic-management statement	
packet gateway.....	594
transfer statement.....	802
transfer-log-archive statement.....	802
translated statement.....	165
translation-type statement.....	166
transport statement	
NAT.....	167
traps statement.....	949
trigger-link-failure statement.....	380
trusted-ca statement	466
ttl statement.....	984
DFC.....	838
ttl-threshold statement.....	100
tunnel statement.....	663, 985
tunnel-group statement.....	309
tunnel-mtu statement.....	277
tunnel-timeout statement.....	310
tunnel-type statement.....	986
twamp statement.....	950
twamp-server statement.....	950
type-of-service statement.....	642

U

udp statement.....	951
unit statement	
encryption.....	664
flow monitoring.....	771
interfaces.....	497
link services.....	407, 903
tunnel.....	987
up statement	
packet gateway.....	597
username statement.....	803
uuid statement.....	100

V

variant statement.....	803
version statement.....	772
version9 statement.....	773
virtual-interface statement	
packet gateway.....	598
virtual-interface-down statement	
packet gateway.....	599
virtual-interface-indications statement	
packet gateway.....	600
virtual-interface-up statement	
packet gateway.....	601

W

warm statement	
packet gateway.....	601
warm-standby statement.....	381

wildcard-response-service-change statement
 packet gateway.....602

Y

yellow-differential-delay statement.....904