JUNOS® Software

Security Configuration Guide

Release 10.0

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

Revision 01
Published: 2009-10-08

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 Security Configuration Guide
Release 10.0
Copyright © 2009, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Revision History
October 2009—Revision 01

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.

SOFTWARE LICENSE

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions. Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details. For complete product documentation, please see the Juniper Networks website at www.juniper.net/techpubs.

END USER LICENSE AGREEMENT

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Abbreviated Table of Contents

		About This Guide	xli
Part 1		**Introduction to JUNOS Software**	
	Chapter 1	Introducing JUNOS Software for SRX Series Services Gateways	3
	Chapter 2	Introducing JUNOS Software for J Series Services Routers	33
Part 2		**Security Zones and Interfaces**	
	Chapter 3	Security Zones and Interfaces	47
Part 3		**Security Policies**	
	Chapter 4	Security Policies	69
	Chapter 5	Security Policy Address Books and Address Sets	97
	Chapter 6	Security Policy Schedulers	105
	Chapter 7	Security Policy Applications	115
Part 4		**Application Layer Gateways**	
	Chapter 8	Application Layer Gateways (ALGs)	157
Part 5		**User Authentication**	
	Chapter 9	Firewall User Authentication	305
	Chapter 10	Infranet Authentication	335
Part 6		**Virtual Private Networks**	
	Chapter 11	Internet Protocol Security (IPsec)	355
	Chapter 12	Public Key Cryptography for Certificates	401
	Chapter 13	Dynamic Virtual Private Networks (VPNs)	433
	Chapter 14	NetScreen-Remote VPN Client	477
Part 7		**Intrusion Detection and Prevention**	
	Chapter 15	IDP Policies	501
	Chapter 16	IDP Signature Database	571
	Chapter 17	IDP Application Identification	589

	Chapter 18	IDP SSL Inspection	601
	Chapter 19	Performance and Capacity Tuning	607
	Chapter 20	IDP Logging	609
Part 8		**Unified Threat Management**	
	Chapter 21	Unified Threat Management Overview	615
	Chapter 22	Antispam Filtering	619
	Chapter 23	Full File-based Antivirus Protection	639
	Chapter 24	Express Antivirus Protection	685
	Chapter 25	Content Filtering	707
	Chapter 26	Web Filtering	723
Part 9		**Attack Detection and Prevention**	
	Chapter 27	Attack Detection and Prevention	761
Part 10		**Chassis Cluster**	
	Chapter 28	Chassis Cluster	867
Part 11		**Network Address Translation**	
	Chapter 29	Network Address Translation	963
Part 12		**GPRS**	
	Chapter 30	General Packet Radio Service	995
Part 13		**Index**	
		Index	1015

Table of Contents

About This Guide **xli**

J Series and SRX Series Documentation and Release Notes xli
Objectives xlii
Audience xlii
Supported Routing Platforms xlii
How to Use This Manual xlii
Document Conventions xliv
Documentation Feedback xlvi
Requesting Technical Support xlvi
Self-Help Online Tools and Resources xlvi
Opening a Case with JTAC xlvii

Part 1 Introduction to JUNOS Software

Chapter 1 Introducing JUNOS Software for SRX Series Services Gateways 3

Stateful and Stateless Data Processing Overview 3
Understanding Flow-Based Processing 4
Zones and Policies 5
Flows and Sessions 5
Understanding Packet-Based Processing 6
Changing Session Characteristics 7
Controlling Session Termination 7
Disabling TCP Packet Security Checks 8
Setting the Maximum Segment Size for All TCP Sessions 8
SRX5600 and SRX5800 Services Gateways Overview 8
Following the Data Path for a Unicast Session 12
Session Lookup and Packet Match Criteria 12
Understanding Session Creation: First-Packet Processing 12
Understanding Fast-Path Processing 15
Data Path Debugging 18
CLI Configuration 19
Network Processor Bundling 19
Understanding Packet Processing 19
Services Processing Units 20
Interface Changes 20
Network Processor Bundling Limitations 21

SRX3400 and SRX3600 Services Gateways Overview 21
Central Point and Combo-Mode Support 22
Load Distribution in Combo Mode 22
Sharing Processing Power and Memory in Combo Mode 23
Following the Data Path for a Unicast Session 23
Session Lookup and Packet Match Criteria 24
Understanding Session Creation: First Packet Processing 24
Understanding Fast-Path Processing 26
Packet Flow and Session Management in SRX210 Services Gateways 26
Flow Processing and Session Management 26
First-Packet Processing 27
Session Creation 27
Fast-Path Processing 27
Obtaining Information About Sessions and Terminating Them 28
Obtaining Information About Sessions 28
Displaying Global Session Parameters 28
Displaying a Summary of Sessions 29
Displaying Session and Flow Information About Sessions 29
Displaying Session and Flow Information About a Specific Session 29
Using Filters to Display Session and Flow Information 29
Information Provided in Session Log Entries 30
Terminating Sessions 32
Terminating All Sessions 32
Terminating a Specific Session 32
Using Filters to Specify the Sessions to Be Terminated 32

Chapter 2 Introducing JUNOS Software for J Series Services Routers 33

Stateful and Stateless Data Processing 33
Flow-Based Processing 34
Zones and Policies 35
Flows and Sessions 35
Packet-Based Processing 36
Changing Session Characteristics 37
Controlling Session Termination 38
Disabling TCP Packet Security Checks 38
Accommodating End-to-End TCP Communication 39
Following the Data Path 39
Part 1—Forwarding Processing 40
Part 2—Session-Based Processing 40
Session Lookup 40
First-Packet Path Processing 41
Fast-Path Processing 42
Part 3—Forwarding Features 42
Understanding Secure and Router Contexts 42
Secure and Router Context Support On Different Device Types 43
Secure Context 43
Router Context 44

Part 2 Security Zones and Interfaces

Chapter 3 Security Zones and Interfaces 47

Understanding Security Zones 47
Functional Zone 48
Security Zone 48
Related Topics 49
Creating Security Zones 49
J-Web Configuration 50
CLI Configuration 50
Related Topics 50
Configuring Security Zones—Quick Configuration 51
Configuring Host Inbound Traffic 53
System Services 54
J-Web Configuration 55
CLI Configuration 56
Related Topics 58
Configuring Protocols 58
J-Web Configuration 59
CLI Configuration 59
Related Topics 60
Configuring the TCP-Reset Parameter 60
J-Web Configuration 61
CLI Configuration 61
Related Topics 61
Understanding Security Zone Interfaces 61
Understanding Interface Ports 62
Related Topics 62
Configuring Interfaces—Quick Configuration 62
Configuring a Gigabit Ethernet Interface—Quick Configuration 63

Part 3 Security Policies

Chapter 4 Security Policies 69

Security Policies Overview 69
Understanding Policies 70
Understanding Policy Rules 71
Understanding Policy Elements 73
Understanding Policy Configuration 74
Related Topics 74
Understanding Policy Ordering 74
Related Topics 76
Configuring Policies—Quick Configuration 76

Configuring Policies 80
J-Web Configuration 81
CLI Configuration 84
Related Topics 84
Verifying Policy Configuration 85
Example: Configuring Security Policies—Detailed Configuration 86
Configuring a Policy to Permit Traffic 86
J-Web Configuration 87
CLI Configuration 89
Related Topics 89
Configuring a Policy to Deny Traffic 90
J-Web Configuration 90
CLI Configuration 91
Related Topics 91
Reordering Policies After They Have Been Created 92
Related Topics 92
Troubleshooting Policy Configuration 92
Checking Commit Failure 93
Verifying Commit 93
Debugging Policy Lookup 93
Monitoring Policy Statistics 94

Chapter 5 Security Policy Address Books and Address Sets 97

Address Books and Address Sets Overview 97
Understanding Address Books 97
Understanding Address Sets 98
Configuring Addresses and Address Sets—Quick Configuration 100
Configuring Addresses 100
Configuring Address Sets 101
Configuring Address Books 102
J-Web Configuration 102
CLI Configuration 103
Related Topics 103
Verifying Address Book Configuration 103

Chapter 6 Security Policy Schedulers 105

Configuring a Scheduler—Quick Configuration 106
Configuring Schedulers 107
J-Web Configuration 107
CLI Configuration 109
Related Topics 110
Associating a Policy to a Scheduler 110
J-Web Configuration 110
CLI Configuration 112
Related Topics 112
Verifying Scheduled Policies 112

Chapter 7 Security Policy Applications 115

Policy Application Sets Overview 116
Related Topics 116
Understanding the ICMP Predefined Policy Application 117
Handling ICMP Unreachable Errors 121
Related Topics 121
Understanding Internet-Related Predefined Policy Applications 121
Related Topics 123
Understanding Microsoft Predefined Policy Applications 123
Related Topics 126
Understanding Dynamic Routing Protocols Predefined Policy Applications 126
Related Topics 127
Understanding Streaming Video Predefined Policy Applications 127
Related Topics 127
Understanding Sun RPC Predefined Policy Applications 128
Related Topics 128
Understanding Security and Tunnel Predefined Policy Applications 129
Related Topics 129
Understanding IP-Related Predefined Policy Applications 130
Related Topics 130
Understanding Instant Messaging Predefined Policy Applications 130
Related Topics 131
Understanding Management Predefined Policy Applications 131
Related Topics 132
Understanding Mail Predefined Policy Applications 133
Related Topics 133
Understanding UNIX Predefined Policy Applications 133
Related Topics 134
Understanding Miscellaneous Predefined Policy Applications 134
Related Topics 135
Understanding Custom Policy Applications 135
Custom Application Mappings 136
Related Topics 136
Configuring Applications and Application Sets—Quick Configuration 136
Configuring Global Custom Application Settings 137
Configuring Custom Application Terms 141
Example: Configuring Applications and Application Sets 142
J-Web Configuration 143
CLI Configuration 144
Related Topics 144
Example: Adding a Custom Policy Application 144
J-Web Configuration 145
CLI Configuration 145
Related Topics 146

Example: Modifying a Custom Policy Application 146
J-Web Configuration 146
CLI Configuration 147
Related Topics 147
Example: Defining a Custom Internet Control Message Protocol Application 147
J-Web Configuration 148
CLI Configuration 149
Related Topics 149
Understanding Policy Application Timeouts 149
Application Timeout Configuration and Lookup 150
Contingencies 150
Related Topics 152
Setting a Policy Application Timeout 152
Related Topics 153

Part 4 Application Layer Gateways

Chapter 8 Application Layer Gateways (ALGs) 157

Understanding Application Layer Gateways 159
Related Topics 159
Configuring Application Layer Gateways—Quick Configuration 159
Understanding the H.323 ALG 162
Related Topics 163
Configuring the H.323 ALG—Quick Configuration 164
Setting H.323 Endpoint Registration Timeout 165
J-Web Configuration 166
CLI Configuration 166
Related Topics 166
Setting H.323 Media Source Port Range 167
J-Web Configuration 167
CLI Configuration 167
Related Topics 168
Configuring H.323 Denial of Service (DoS) Attack Protection 168
J-Web Configuration 168
CLI Configuration 169
Related Topics 169
Allowing Unknown H.323 Message Types 169
J-Web Configuration 170
CLI Configuration 170
Related Topics 171
Verifying the H.323 Configuration 171
Verifying H.323 Counters 171
Related Topics 172

Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone 172
J-Web Configuration 173
CLI Configuration 175
Related Topics 175
Passing H.323 ALG Traffic to a Gatekeeper in the External Zone 176
J-Web Configuration 176
CLI Configuration 180
Related Topics 181
Using NAT and the H.323 ALG to Enable Outgoing Calls 181
CLI Configuration 182
Related Topics 183
Using NAT and the H.323 ALG to Enable Incoming Calls 183
CLI Configuration 184
Related Topics 185
Understanding the Avaya H.323 ALG 185
Avaya H.323-Specific Features 185
Call Flow Details in Avaya H.323 186
Related Topics 187
Understanding the SIP ALG 187
SIP ALG Operation 188
SDP Session Descriptions 189
Pinhole Creation 190
SIP ALG Request Methods Overview 191
Related Topics 193
Configuring the SIP ALG—Quick Configuration 193
Understanding SIP ALG Call Duration and Timeouts 196
Related Topics 197
Setting SIP Call Duration and Inactive Media Timeout 197
J-Web Configuration 198
CLI Configuration 199
Related Topics 199
Configuring SIP Denial of Service (DoS) Attack Protection 199
J-Web Configuration 200
CLI Configuration 200
Related Topics 200
Allowing Unknown SIP Message Types 200
J-Web Configuration 201
CLI Configuration 202
Related Topics 202
Disabling SIP Call ID Hiding 202
J-Web Configuration 202
CLI Configuration 203
Related Topics 203
Retaining SIP Hold Resources 203
J-Web Configuration 204
CLI Configuration 204
Related Topics 204
Understanding SIP with Network Address Translation (NAT) 204
Outgoing Calls 206
Incoming Calls 206
Forwarded Calls 207

Call Termination 207
Call Re-INVITE Messages 207
Call Session Timers 207
Call Cancellation 207
Forking 208
SIP Messages 208
SIP Headers 208
SIP Body 210
SIP NAT Scenario 210
Classes of SIP Responses 212
Related Topics 213
Understanding Incoming SIP Call Support Using the SIP Registrar 213
Related Topics 215
Configuring Interface Source NAT for Incoming SIP Calls 216
CLI Configuration 216
Related Topics 217
Configuring a Source NAT Pool for Incoming SIP Calls 218
J-Web Configuration 218
CLI Configuration 222
Related Topics 223
Configuring Static NAT for Incoming SIP Calls 223
J-Web Configuration 224
CLI Configuration 227
Related Topics 227
Configuring the SIP Proxy in the Private Zone 228
CLI Configuration 228
Related Topics 229
Configuring the SIP Proxy in the Public Zone 230
J-Web Configuration 230
CLI Configuration 234
Related Topic 235
Configuring a Three-Zone SIP Scenario 235
J-Web Configuration 236
CLI Configuration 242
Related Topics 243
Verifying the SIP Configuration 243
Verifying the SIP ALG 244
Related Topics 244
Verifying SIP Calls 244
Related Topics 245
Verifying SIP Call Detail 245
Related Topics 246
Verifying SIP Transactions 246
Related Topics 246
Verifying SIP Counters 247
Related Topics 248
Verifying the Rate of SIP Messages 248
Related Topics 248

Understanding the SCCP ALG 249
SCCP Security 250
SCCP Components 250
SCCP Client 251
CallManager 251
Cluster 251
SCCP Transactions 251
Client Initialization 252
Client Registration 252
Call Setup 253
Media Setup 253
SCCP Control Messages and RTP Flow 253
SCCP Messages 254
Related Topics 255
Configuring the SCCP ALG—Quick Configuration 255
Setting SCCP Inactive Media Timeout 257
J-Web Configuration 258
CLI Configuration 258
Related Topics 258
Allowing Unknown SCCP Message Types 258
J-Web Configuration 259
CLI Configuration 260
Related Topics 260
Configuring SCCP Denial of Service (DoS) Attack Protection 260
J-Web Configuration 260
CLI Configuration 261
Related Topics 261
Configuring Call Manager/TFTP Server in the Private Zone 261
CLI Configuration 262
Verifying the SCCP Configuration 263
Verifying the SCCP ALG 263
Related Topics 263
Verifying SCCP Calls 263
Related Topics 264
Verifying SCCP Call Details 264
Related Topics 265
Verifying SCCP Counters 265
Related Topics 266
Understanding the MGCP ALG 266
MGCP Security 267
Entities in MGCP 267
Endpoint 267
Connection 268
Call 268
Call Agent 268
Commands 268
Response Codes 271
Related Topics 272
Configuring the MGCP ALG—Quick Configuration 272
Understanding MGCP ALG Call Duration and Timeouts 274
Related Topics 275

Setting MGCP Call Duration 275
J-Web Configuration 276
CLI Configuration 276
Related Topics 276
Setting MGCP Inactive Media Timeout 276
J-Web Configuration 277
CLI Configuration 277
Related Topics 277
Setting the MGCP Transaction Timeout 278
J-Web Configuration 278
CLI Configuration 278
Related Topics 279
Configuring MGCP Denial of Service (DoS) Attack Protection 279
J-Web Configuration 279
CLI Configuration 280
Related Topics 280
Allowing Unknown MGCP Message Types 280
J-Web Configuration 281
CLI Configuration 282
Related Topics 282
Configuring a Media Gateway in Subscribers' Homes 282
J-Web Configuration 283
CLI Configuration 288
Related Topics 289
Configuring Three-Zone ISP-Hosted Service Using Source and Static NAT 289
CLI Configuration 290
Related Topics 292
Verifying the MGCP Configuration 293
Verifying the MGCP ALG 293
Related Topics 293
Verifying MGCP Calls 293
Related Topics 294
Verifying MGCP Endpoints 294
Related Topics 294
Verifying MGCP Counters 295
Related Topics 295
Understanding the RPC ALG 296
Sun RPC ALG 297
Typical RPC Call Scenario 297
Sun RPC Services 298
Customizing Sun RPC Services 298
Microsoft RPC ALG 299
MS RPC Services in Security Policies 299
Predefined Microsoft RPC Services 299
Related Topics 300

Disabling and Enabling RPC ALG 300
J-Web Configuration 300
CLI Configuration 301
Related Topics 301
Verifying the RPC ALG Tables 301
Display the Sun RPC Port Mapping Table 301
Display the Microsoft RPC UUID Mapping Table 302
Related Topics 302

Part 5 User Authentication

Chapter 9 Firewall User Authentication 305

Firewall User Authentication Overview 305
Authentication, Authorization, and Accounting (AAA) Servers 306
Types of Firewall User Authentication 306
Related Topics 307
Understanding Authentication Schemes 307
Pass-Through Authentication 307
Web Authentication 308
Related Topics 310
Configuring for Pass-Through Authentication 310
J-Web Configuration 311
CLI Configuration 313
Related Topics 314
Configuring for Web Authentication 315
J-Web Configuration 315
CLI Configuration 318
Related Topics 319
Understanding Client Groups for Firewall Authentication 319
J-Web Configuration 320
CLI Configuration 321
J-Web Configuration 321
CLI Configuration 321
Related Topics 322
Configuring for External Authentication Servers 322
J-Web Configuration 322
CLI Configuration 325
Related Topics 326
Understanding SecurID User Authentication 326
Related Topics 326
Configuring the SecurID Server 327
Configuring SecurID as the External Authentication Server 327
CLI Configuration 327
Deleting the Node Secret File 328
Related Topics 329

Displaying the Authentication Table 329
J-Web Configuration 329
CLI Configuration 329
Related Topics 329
Understanding Banner Customization 330
Related Topics 330
Customizing a Banner 330
J-Web Configuration 331
CLI Configuration 331
Related Topics 332
Configuring Firewall Authentication—Quick Configuration 332
Verifying Firewall User Authentication 333

Chapter 10 Infranet Authentication 335

Unified Access Control Overview 335
Communications Between the JUNOS Enforcer and the Infranet Controller 336
JUNOS Enforcer Policy Enforcement 337
Communications Between the JUNOS Enforcer and a Cluster of Infranet Controllers 338
Communications Between the JUNOS Enforcer and the Infranet Agent 339
Related Topics 339
Enabling Unified Access Control on SRX Series and J Series Devices 340
Support and Limitations 341
CLI Configuration 341
Related Topics 342
Configuring the SRX Series and J Series Devices as a JUNOS Enforcer 342
CLI Configuration 342
Configuring the SRX Series Device as a JUNOS Enforcer using IPsec — Example 343
Related Topics 350
Configuring the JUNOS Enforcer Failover Options 350
CLI Configuration 351
Related Topics 351
Enabling the JUNOS Enforcer Test-Only Mode 351

Part 6 Virtual Private Networks

Chapter 11 Internet Protocol Security (IPsec) 355

Virtual Private Networks (VPNs) Overview 355
Security Associations (SAs) 356
IPsec Key Management 357
Manual Key 357
AutoKey IKE 357
Diffie-Hellman Exchange 358

IPsec Security Protocols 359
Authentication Header (AH) Protocol 359
Encapsulating Security Payload (ESP) Protocol 359
IPsec Tunnel Negotiation 360
Phase 1 of IKE Tunnel Negotiation 360
Phase 2 of IKE Tunnel Negotiation 362
Distributed VPN in SRX Series Services Gateways 363
Related Topics 364
Understanding IKE and IPsec Packet Processing 364
Packet Processing in Tunnel Mode 364
IKE Packet Processing 365
IPsec Packet Processing 367
Related Topics 369
Configuring an IPsec Tunnel—Overview 369
Configuring VPN Global Settings (Standard VPNs) 370
J-Web Configuration 371
CLI Configuration 372
Configuring IPSec VPN Global Settings in J-Web (Standard VPNs) 372
Configuring an IKE Phase 1 Proposal (Standard and Dynamic VPNs) 373
J-Web Configuration 374
CLI Configuration 374
Configuring Auto Tunnel in J-Web (Standard VPN) 375
Configuring an IKE Policy (Standard and Dynamic VPNs) 384
J-Web Configuration 384
CLI Configuration 385
Configuring an IKE Gateway (Standard and Dynamic VPNs) 385
J-Web Configuration 386
CLI Configuration 386
Configuring an IPsec Phase 2 Proposal (Standard and Dynamic VPNs) 387
J-Web Configuration 387
CLI Configuration 388
Configuring an IPsec Policy (Standard and Dynamic VPNs) 388
J-Web Configuration 389
CLI Configuration 389
Configuring IPsec AutoKey (Standard and Dynamic VPNs) 390
J-Web Configuration 390
CLI Configuration 391
Configuring the Manual Tunnel using J-Web 391
Configuring Hub-and-Spoke VPNs 393
Creating Hub-and-Spoke VPNs 394
Configuring the IPSec Tunnel on the Hub 394
Configuring Spoke 1 397
Configuring Spoke 2 398

Chapter 12 Public Key Cryptography for Certificates 401

Understanding Public Key Cryptography 402
Related Topics 402
Understanding Certificates 402
Certificate Signatures 403
Certificate Verification 403
Internet Key Exchange 404
Related Topics 404
Understanding Certificate Revocation Lists 405
Related Topics 405
Understanding Public Key Infrastructure 405
PKI Hierarchy for a Single CA Domain or Across Domains 406
PKI Management and Implementation 407
Related Topics 407
Understanding Self-Signed Certificates 408
About Generating Self-Signed Certificates 408
Related Topics 409
Understanding Automatically Generated Self-Signed Certificates 409
Related Topics 409
Understanding Manually Generated Self-Signed Certificates 410
Related Topics 410
Using Digital Certificates 410
Obtaining Digital Certificates Online 411
Obtaining Digital Certificates Manually 411
Verifying the Validity of a Certificate 412
Deleting a Certificate 412
Generating a Public-Private Key Pair 412
CLI Operation 413
Related Topics 413
Configuring a Certificate Authority Profile 413
CLI Configuration 414
Related Topics 414
Enrolling a CA Certificate Online 415
CLI Operation 415
Related Topics 416
Enrolling a Local Certificate Online 416
CLI Configuration 416
Related Topics 417
Generating a Local Certificate Request Manually 418
CLI Operation 419
Related Topics 419
Loading CA and Local Certificates Manually 420
CLI Operation 420
Related Topics 421
Re-enrolling Local Certificates Automatically 421
CLI Configuration 422
Related Topics 422

Manually Loading a CRL onto the Device 423
CLI Operation 423
Related Topics 423
Verifying Certificate Validity 424
CLI Operation 424
Related Topics 425
Checking Certificate Validity Using CRLs 425
J-Web Configuration 426
CLI Configuration 426
Related Topics 426
Using Automatically Generated Self-Signed Certificates 427
J-Web Configuration 427
CLI Configuration 427
Related Topics 428
Manually Generating Self-Signed Certificates 428
J-Web Configuration 429
CLI Configuration 429
Related Topics 429
Deleting Certificates 430
CLI Operation 430
Related Topics 431
Deleting a Loaded CRL 431
CLI Operation 431
Related Topics 432

Chapter 13 Dynamic Virtual Private Networks (VPNs) 433

Dynamic VPN Overview 433
Connecting to the Remote Access Server for the First Time (Pre-IKE Phase) 435
Connecting to the Remote Access Server for Subsequent Sessions (Pre-IKE Phase) 438
Establishing an IPsec VPN Tunnel (IKE Phase) 439
Related Topics 442
Access Manager Client-Side Reference 442
Client-Side System Requirements 442
Client-Side Files 443
Client-Side Registry Changes 445
Client-Side Error Messages 446
Client-Side Logging and Connection Information 448
Configuring a Dynamic VPN—Overview 449
Configuring an IKE Phase 1 Proposal— (Dynamic VPNs) 451
Configuring an IKE Policy—(Dynamic VPNs) 454
Configuring an IKE Gateway—Quick Configuration (Dynamic VPNs) 456
Configuring an IPsec Phase 2 Proposal—Quick Configuration (Dynamic VPNs) 460
Configuring an IPsec Policy—Quick Configuration (Dynamic VPNs) 463
Configuring an IPsec Autokey—Quick Configuration (Dynamic VPNs) 465
Creating a Client Configuration—Quick Configuration (Dynamic VPNs) 469

Creating a Client Configuration (Dynamic VPNs) 472
J-Web Configuration 472
CLI Configuration 473
Configuring Global Client Download Settings—Quick Configuration (Dynamic VPNs) 473
Configuring Global Client Download Settings (Dynamic VPNs) 475
J-Web Configuration 476
CLI Configuration 476

Chapter 14 NetScreen-Remote VPN Client 477

System Requirements for NetScreen-Remote Client Installation 477
Installing the NetScreen-Remote Client on a PC or Laptop 478
Starting NetScreen-Remote Client Installation 478
Starting Installation from a CD-ROM 479
Starting Installation from a Network Share Drive 479
Starting Installation from a Website 479
Completing NetScreen-Remote Client Installation 480
Configuring the Firewall on the Router 483
Firewall Configuration Overview 483
Configuring a Security Zone 483
Configuring a Tunnel Interface 484
Configuring an Access Profile for XAuth 485
Configuring an IKE Gateway 485
Configuring Policies 486
Configuring the PC or Laptop 486
Creating a New Connection 487
Creating the Preshared Key 490
Defining the IPsec Protocols 491
Logging In to the NetScreen Remote Client 496

Part 7 Intrusion Detection and Prevention

Chapter 15 IDP Policies 501

IDP Policies Overview 502
IDP Policy Terms 502
Working with IDP Policies 503
Understanding IDP Policy Rulebases 503
Understanding Application-level DDoS Rulebase 504
Related Topics 504
IPS Rulebase 505
Exempt Rulebase 505
Related Topics 506
Understanding IDP Policy Rules 506
Related Topics 507
Understanding IDP Rule Match Conditions 507
Related Topics 508

Understanding IDP Rule Objects 508
Zone Objects 508
Address or Network Objects 508
Application or Service Objects 508
Attack Objects 509
Signature Attack Objects 509
Protocol Anomaly Attack Objects 509
Compound Attack Objects 509
Attack Object Groups 510
Related Topics 510
Understanding IDP Rule Actions 510
Related Topics 512
Understanding IDP Rule IP Actions 512
Related Topics 513
Understanding IDP Rule Notifications 514
Related Topics 514
Defining Rules for an IPS Rulebase 515
J-Web Configuration 516
CLI Configuration 517
Related Topics 519
Defining Rules for an Exempt Rulebase 519
J-Web Configuration 520
CLI Configuration 521
Related Topics 521
IDP Policies—Quick Configuration 522
Configuring IDP Policies—Quick Configuration 522
Adding a New IDP Policy—Quick Configuration 523
Adding an IPS Rulebase—Quick Configuration 525
Adding an Exempt Rulebase—Quick Configuration 528
Inserting a Rule in the Rulebase 530
CLI Configuration 531
Related Topics 531
Deactivating and Reactivating Rules in a Rulebase 532
CLI Configuration 532
Related Topics 533
Understanding Application Sets 533
Related Topics 533
Configuring Applications or Services for IDP 533
CLI Configuration 534
Related Topics 535
Configuring Application Sets for IDP 535
CLI Configuration 536
Related Topics 536
Enabling IDP in a Security Policy 536
J-Web Configuration 537
CLI Configuration 539
Related Topics 540
Understanding IDP Terminal Rules 540
Related Topics 541

Setting Terminal Rules in Rulebases 541
J-Web Configuration 542
CLI Configuration 543
Related Topics 544
Understanding Custom Attack Objects 544
Attack Name 545
Severity 545
Service or Application Binding 545
Protocol or Port Bindings 549
Time Bindings 550
Scope 551
Count 551
Attack Properties—Signature Attacks 551
Attack Context 552
Attack Direction 553
Attack Pattern 553
Protocol-Specific Parameters 553
Sample Signature Attack Definition 556
Attack Properties—Protocol Anomaly Attacks 557
Attack Direction 557
Test Condition 557
Sample Protocol Anomaly Attack Definition 557
Attack Properties—Compound or Chain Attacks 558
Scope 558
Order 558
Reset 558
Expression (Boolean expression) 559
Member Index 559
Sample Compound Attack Definition 560
Related Topics 560
Understanding Protocol Decoders 560
CLI Configuration Example 561
Multiple IDP Detector Support 561
Configuring Signature-Based Attacks 562
CLI Configuration 564
Related Topics 565
Configuring Protocol Anomaly-Based Attacks 565
CLI Configuration 566
Related Topics 567
Specifying IDP Test Conditions for a Specific Protocol 567
Related Topics 568
Configuring DSCP in an IDP Policy 568
CLI Configuration 569
Related Topics 570

Chapter 16 IDP Signature Database 571

Understanding the IDP Signature Database 571
Related Topics 572
Using Predefined Policy Templates 572
CLI Configuration 574
Related Topics 575
Understanding Predefined Attack Objects and Groups 575
Predefined Attack Objects 576
Predefined Attack Object Groups 576
Related Topics 577
Updating the Signature Database Overview 577
Related Topics 578
Updating the Signature Database Manually 578
CLI Configuration 579
Related Topics 580
Configuring a Security Package Update—Quick Configuration 580
Updating the Signature Database Automatically 582
CLI Configuration 583
Related Topics 584
Understanding the Signature Database Version 584
Related Topics 584
Verifying the Signature Database 584
Verifying the Policy Compilation and Load Status 585
Verifying the Signature Database Version 586

Chapter 17 IDP Application Identification 589

IDP Application Identification Support on Different Device Types 589
Understanding Application Identification 590
Related Topics 590
Understanding Service and Application Bindings 591
Related Topics 592
Understanding Application System Cache 593
Related Topics 593
Configuring IDP Policies for Application Identification 593
CLI Configuration 594
Related Topics 595
Disabling Application Identification 595
CLI Configuration 595
Related Topics 596
Setting Memory and Session Limits 596
CLI Configuration 597
Related Topics 598
Verifying Application Identification 598
Verifying the Application System Cache 598
Verifying Application Identification Counters 599

Chapter 18 IDP SSL Inspection 601

IDP SSL Overview 601
Supported Ciphers 602
Key Exchange 603
Server Key Management and Policy Configuration 603
Displaying Keys and Servers 604
Adding Keys and Servers 604
Deleting Keys and Servers 605
Configuring SSL Inspection 605

Chapter 19 Performance and Capacity Tuning 607

Performance and Capacity Tuning for IDP Overview 607
Configuring Session Capacity for IDP 608
CLI Configuration 608

Chapter 20 IDP Logging 609

Understanding IDP Logging 609
Related Topics 610
Configuring Log Suppression Attributes 610
CLI Configuration 611
Related Topics 611
IDP Attack Log Information Usage on the Infranet Controller 611
IDP Log Message Filtering to the Infranet Controller 612
Configuration Information 612

Part 8 Unified Threat Management

Chapter 21 Unified Threat Management Overview 615

Custom Objects 616
Platform Support and Licensing 616

Chapter 22 Antispam Filtering 619

Using Server-Based Spam Filtering 619
Configuring Server-Based Spam Filtering 620
Configuration Overview 620
J-Web Configuration 621
J-Web Point and Click CLI Configuration 623
CLI Configuration 625
Using Local List Spam Filtering 626

Configuring Local List Spam Filtering 627
Configuration Overview 627
J-Web Configuration 628
J-Web Point and Click CLI Configuration 631
CLI Configuration 634
Understanding Spam Message Handling 636
Blocking Detected Spam 636
Tagging Detected Spam 636
Verifying Antispam Configurations 637
Using J-Web for Antispam Monitoring 637
Using the CLI for Antispam Monitoring 637

Chapter 23 Full File-based Antivirus Protection 639

Understanding the Internal Scan Engine and Scan Options 640
Scan Type Support 640
Scan Mode Support 641
Enable and Disable Scanning on a Per-Protocol Basis 641
File Extension Scanning 641
Intelligent Prescreening 642
Signature Database Support 642
Content Size Limits 642
Decompression Layer Limit 643
Scanning Timeout 643
Global, Profile-Based, and Policy-Based Scan Settings 644
Scan Session Throttling 645
Understanding Updating the Antivirus Scanner Pattern Database 645
Updating Antivirus Patterns 645
Example: Automatic Update 646
Example: Manual Update 646
Database Download Process 647
Understanding the Scanning of Application Protocols 647
HTTP Scanning 648
HTTP Scanning: Process Overview 649
HTTP Trickling 649
MIME White List 649
URL White List 650
Script-only Scanning 650
FTP Scanning 651
FTP Scanning: Process Overview 651
SMTP Scanning 651
SMTP Scanning: Process Overview 652
Mail Message Replacement 652
Sender Notification 652
Subject Tagging 653
POP3 Scanning 653
POP3 Scanning: Process Overview 653
Mail Message Replacement 654
Sender Notification 654
Subject Tagging 654

IMAP Scanning 655
IMAP Scanning: Process Overview 655
Mail Message Replacement 655
Sender Notification 656
Subject Tagging 656
IMAP Antivirus Scanning Limitations 656
Configuring Full Antivirus Protection 657
Configuration Overview 657
J-Web Configuration 658
J-Web Point and Click CLI Configuration 664
CLI Configuration 671
Understanding Virus-Detected Notification Options 676
Protocol-Only Notifications 676
E-mail Notification 676
Custom Message Notification 677
Fallback Options 677
Understanding Scan Result Handing 679
Verifying Antivirus Configurations 679
General Scan Engine Status 679
Session Status 680
Verifying Antivirus Scan Results Using J-Web 680
Verifying Antivirus Scan Results Using the CLI 682

Chapter 24 Express Antivirus Protection 685

Packet-Based Scanning Versus File-Based Scanning 686
Express Antivirus Limitations 686
Expanded MIME Decoding Support 686
Result Handling 687
Updating Antivirus Patterns 687
Example: Automatic Update 688
Example: Manual Update 688
Intelligent prescreening Support 688
Configuring Express Antivirus 688
Configuration Overview 689
J-Web Configuration 689
J-Web Point and Click CLI Configuration 695
CLI Configuration 701

Chapter 25 Content Filtering 707

Types of Content Filters 707
Content Filtering Protocol Support 708
HTTP Support 708
FTP Support 709
E-mail Support 709
Example: Content Filter CLI 709
Configuring Content Filtering 710
Configuration Overview 710
J-Web Configuration 711

J-Web Point and Click CLI Configuration 715
CLI Configuration 719
Verifying Content Filtering Configurations 722
Using J-Web to Monitor Content Filtering 722
Using the CLI to Monitor Content Filtering 722

Chapter 26 Web Filtering 723

Understanding Integrated Web Filtering 724
Integrated Web Filtering: Process Overview 725
Integrated Web Filtering Cache 725
Web Filtering Profiles 726
Profile Matching Precedence 726
Configuring Integrated Web Filtering 727
Configuration Overview 727
J-Web Configuration 728
J-Web Point and Click CLI Configuration 732
CLI Configuration 736
Understanding Redirect Web Filtering 739
Redirect Web Filtering: Process Overview 740
Configuring Redirect Web Filtering 740
Configuration Overview 740
J-Web Configuration 741
J-Web Point and Click CLI Configuration 745
CLI Configuration 749
Understanding Local Web Filtering 751
Local Web Filtering: Process Overview 752
Web Filtering Profiles 752
Profile Matching Precedence 753
Configuring Local Web Filtering 753
Verifying Web Filtering Configurations 755
Using J-Web to Monitor Web Filtering 755
Using the CLI to Monitor Web Filtering 755
What Happens When Both Web Filtering and Antivirus Scanning Are Employed? 756
WebTrends Enhanced Log File Format (WELF) for UTM Features 756
WELF Overview 756
CLI Configuration 757

Part 9 Attack Detection and Prevention

Chapter 27 Attack Detection and Prevention 761

Reconnaissance Deterrence Overview 763
Related Topics 764
Understanding IP Address Sweeps 764
Related Topics 765

Blocking IP Address Sweeps 765
J-Web Configuration 765
CLI Configuration 766
Related Topics 766
Understanding Port Scanning 766
Related Topics 767
Blocking Port Scans 768
J-Web Configuration 768
CLI Configuration 768
Related Topics 768
Understanding Network Reconnaissance Using IP Options 769
Uses for IP Packet Header Options 769
Screen Options for Detecting IP Options Used For Reconnaissance 771
Related Topics 771
Detecting Packets That Use IP Options for Reconnaissance 771
J-Web Configuration 771
CLI Configuration 773
Understanding Operating System Probes 773
TCP Headers with SYN and FIN Flags Set 773
TCP Headers With FIN Flag and Without ACK Flag 774
TCP Header Without Flags Set 775
Related Topics 775
Blocking Packets with SYN and FIN Flags Set 776
J-Web Configuration 776
CLI Configuration 777
Related Topics 777
Blocking Packets with FIN Flag/No ACK Flag Set 777
J-Web Configuration 777
CLI Configuration 778
Related Topics 778
Blocking Packets with No Flags Set 778
J-Web Configuration 778
CLI Configuration 779
Related Topics 779
Understanding Attacker Evasion Techniques 779
FIN Scan 780
Non-SYN Flags 780
IP Spoofing 782
IP Source Route Options 782
Related Topics 784
Thwarting a FIN Scan 784
Related Topics 785
Setting TCP SYN Checking 785
J-Web Configuration 786
CLI Configuration 786
Related Topics 786
Setting Strict Syn Checking 786
J-Web Configuration 787
CLI Configuration 787
Related Topics 787

Blocking IP Spoofing 787
J-Web Configuration 788
CLI Configuration 788
Related Topics 788
Blocking Packets with Either a Loose or Strict Source Route Option Set 789
J-Web Configuration 789
CLI Configuration 790
Related Topics 790
Detecting Packets with Either a Loose or Strict Source Route Option Set 790
J-Web Configuration 790
CLI Configuration 791
Related Topics 791
Suspicious Packet Attributes Overview 791
Related Topics 792
Understanding ICMP Fragment Protection 792
Related Topics 793
Blocking Fragmented ICMP Packets 793
J-Web Configuration 794
CLI Configuration 794
Related Topics 794
Understanding Large ICMP Packet Protection 794
Related Topics 796
Blocking Large ICMP Packets 796
J-Web Configuration 796
CLI Configuration 797
Related Topics 797
Understanding Bad IP Option Protection 797
Related Topics 798
Detecting and Blocking IP Packets with Incorrectly Formatted Options 798
J-Web Configuration 799
CLI Configuration 799
Related Topics 799
Understanding Unknown Protocol Protection 799
Related Topics 800
Dropping Packets Using an Unknown Protocol 800
J-Web Configuration 801
CLI Configuration 801
Related Topics 801
Understanding IP Packet Fragment Protection 802
Related Topics 803
Dropping Fragmented IP Packets 803
J-Web Configuration 803
CLI Configuration 804
Related Topics 804
Understanding SYN Fragment Protection 804
Related Topics 805
Dropping IP Packets Containing SYN Fragments 805
J-Web Configuration 806
CLI Configuration 806
Related Topics 806

Denial-of-Service Attack Overview 807
Related Topics 807
Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview 807
Related Topics 807
Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview 808
Understanding Application-level DDoS Module 808
Understanding Application-ddos Definition 809
Understanding AppDDoS Rule 810
Understanding AppDDoS IP-action 811
Understanding Session Action 811
Related Topics 811
Enabling Protection Against Application-level Distributed Denial-of-Service (AppDDoS) Attack 812
CLI Configuration 813
Related Topics 814
Firewall DoS Attacks Overview 815
Related Topics 815
Understanding Session Table Flood Attacks 815
Source-Based Session Limits 816
Destination-Based Session Limits 817
Related Topics 817
Setting Source-Based Session Limits 818
J-Web Configuration 818
CLI Configuration 819
Related Topics 819
Setting Destination-Based Session Limits 819
J-Web Configuration 820
CLI Configuration 820
Related Topics 821
Understanding SYN-ACK-ACK Proxy Flood Attacks 821
Related Topics 821
Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack 821
J-Web Configuration 822
CLI Configuration 822
Related Topics 822
Network DoS Attacks Overview 823
Related Topics 823
Understanding SYN Flood Attacks 823
SYN Flood Protection 824
SYN Flood Options 826
Related Topics 828
Example: SYN Flood Protection 829
J-Web Configuration 831
CLI Configuration 834
Related Topics 835
Enabling SYN Flood Protection 835
Related Topics 835
Understanding SYN Cookie Protection 836
Related Topics 837

Enabling SYN Cookie Protection 837
J-Web Configuration 838
CLI Configuration 838
Related Topics 839
Understanding ICMP Flood Attacks 839
Related Topics 840
Enabling ICMP Flood Protection 840
J-Web Configuration 841
CLI Configuration 841
Related Topics 841
Understanding UDP Flood Attacks 842
Related Topics 843
Enabling UDP Flood Protection 843
J-Web Configuration 843
CLI Configuration 844
Related Topics 844
Understanding Land Attacks 844
Related Topics 845
Enabling Protection Against a Land Attack 845
J-Web Configuration 846
CLI Configuration 846
Related Topics 846
OS-Specific DoS Attacks Overview 847
Related Topics 847
Understanding Ping of Death Attacks 847
Related Topics 848
Enabling Protection Against a Ping of Death Attack 848
J-Web Configuration 848
CLI Configuration 849
Related Topics 849
Understanding Teardrop Attacks 849
Related Topics 851
Enabling Protection Against a Teardrop Attack 851
J-Web Configuration 851
CLI Configuration 852
Related Topics 852
Understanding WinNuke Attacks 852
Related Topics 853
Enabling Protection Against a WinNuke Attack 853
J-Web Configuration 854
CLI Configuration 854
Related Topics 854
Configuring Firewall Screen Options—Quick Configuration 855
Verifying Application Security Information Using Trace Options 860
Setting Security Trace Options 861
J-Web Configuration 861
CLI Configuration 862
Example: Show Security Traceoptions Output 863
Verifying Application Security Flow Information 864

Part 10 Chassis Cluster

Chapter 28 Chassis Cluster 867

Understanding Chassis Cluster 868
Related Topics 869
Understanding Chassis Cluster Formation 869
Related Topics 869
Understanding Redundancy Groups 870
About Redundancy Groups 870
Redundancy Group 0: Routing Engines 871
Redundancy Groups 1 Through 128 871
Redundancy Group Interface Monitoring 874
Redundancy Group IP Address Monitoring 875
Related Topics 877
Understanding Monitoring 877
SPU Monitoring 877
Flowd Monitoring 878
Cold-Sync Monitoring 878
Related Topics 879
Understanding Redundant Ethernet Interfaces 879
Related Topics 880
Understanding the Control Plane 881
About the Control Link 882
SRX Series Devices 882
J Series Devices 882
About Dual Control LInks 882
About Heartbeats 883
About Control Link Failure and Recovery 884
Related Topics 885
Understanding the Data Plane 885
About Session RTOs 886
About the Fabric Data Link 886
About Data Forwarding 887
About Fabric Data Link Failure and Recovery 887
Related Topics 888
Understanding Failover 888
About Redundancy Group Failover 888
About Manual Failover 889
Hardware Setup for SRX Series Chassis Clusters 890
Connecting Dual Control Links 893
Hardware Setup for J Series Chassis Clusters 894
What Happens When You Enable Chassis Cluster 895
Node Interfaces on SRX Series Chassis Clusters 896
Node Interfaces on J Series Chassis Clusters 901
Management Interfaces 903
Fabric Interface 903
Control Interfaces 903
Related Topics 904

Creating an SRX Series Chassis Cluster—Overview 904
Related Topics 906
Creating a J Series Chassis Cluster—Overview 906
Related Topics 908
Configuring the Control Ports 909
CLI Configuration 909
Related Topics 910
Setting the Node ID and Cluster ID 910
CLI Configuration 910
Related Topics 911
Configuring the Management Interface 911
CLI Configuration 912
Related Topics 912
Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration 912
Related Topics 915
Configuring Redundant Ethernet Interfaces—Quick Configuration 915
Configuring Chassis Cluster Information 919
CLI Configuration 920
Related Topics 920
Configuring the Fabric 920
CLI Configuration 921
Related Topics 921
Configuring Redundancy Groups 921
CLI Configuration 922
Related Topics 923
Configuring Redundant Ethernet Interfaces 923
CLI Configuration 923
Related Topics 924
Configuring Control Link Recovery 924
CLI Configuration 925
Related Topics 925
Configuring Interface Monitoring 925
CLI Configuration 926
Related Topics 926
Configuring a Dampening Time Between Back-to-Back Redundancy Group Failovers 926
CLI Configuration 926
Related Topics 926
Initiating a Manual Redundancy Group Failover 927
CLI Configuration 927
Configuring Conditional Route Advertising 929
CLI Configuration 930
Related Topics 931
Verifying the Chassis Cluster Configuration 931
Verifying the Chassis Cluster 932
Related Topics 932
Verifying Chassis Cluster Interfaces 932
Verifying Chassis Cluster Statistics 933
Verifying Chassis Cluster Status 936
Verifying Chassis Cluster Redundancy Group Status 937

SNMP Failover Traps 938
Related Topics 939
Upgrading Devices in a Chassis Cluster 939
Upgrading Each Device Separately 939
Low-Impact ISSU Chassis Cluster Upgrades 939
Automatic Failback 941
Troubleshooting ISSU Failure 941
Mismatched Control Link Statistics During an ISSU 941
Related Topics 941
Disabling Chassis Cluster 942
Related Topics 942
Chassis Cluster Configuration Scenarios 942
Active/Passive Chassis Cluster Scenario 942
CLI Configuration 943
J-Web Configuration 945
Active/Passive Chassis Cluster with IPsec Tunnel Scenario 947
CLI Configuration 949
J-Web Configuration 952
Asymmetric Routing Chassis Cluster Scenario 954
Case 1: Failures in the Trust Zone reth 955
Case 2: Failures in the Untrust Zone Interfaces 956
CLI Configuration 956
J-Web Configuration 958

Part 11 Network Address Translation

Chapter 29 Network Address Translation 963

NAT Overview 964
Source NAT 964
Destination NAT 965
Static NAT 965
Related Topics 965
Understanding NAT Rule Sets and Rules 965
NAT Rule Sets 965
NAT Rules 966
Rule Processing 967
Related Topics 967
Understanding Static NAT 967
Related Topics 968
Static NAT Configuration Overview 968
Static NAT Rules 968
Related Topic 969
Example: Configuring Static NAT 969
CLI Configuration 969
Understanding Destination NAT 970
Related Topics 970

Destination NAT Configuration Overview 970
Destination NAT Address Pool 971
Destination NAT Rules 971
Related Topic 972
Example: Configuring Destination NAT 972
CLI Configuration 972
Configuring Destination NAT—Quick Configuration 973
Understanding Source NAT 974
Related Topics 975
Source NAT Configuration Overview 975
Source NAT Pools 975
Pool Utilization Alarms 976
Persistent Addresses 976
Source NAT Rules 976
Related Topics 977
Example: Configuring Source NAT Pools 977
CLI Configuration 977
Understanding NAT Source Pools with PAT 978
Port Ranges 978
Address Persistent 979
Related Topics 979
Disabling Port Randomization for Source NAT 979
CLI Configuration 979
Understanding NAT Source Pools Without PAT 979
Source Pool Utilization Alarm 980
Related Topics 980
Example: Configuring Source NAT 980
CLI Configuration 981
Configuring Source NAT—Quick Configuration 982
Example: Configuring a Persistent Address 984
CLI Configuration 985
Related Topics 985
Understanding Persistent NAT 985
Related Topics 986
Understanding Session Traversal Utilities for NAT (STUN) Protocol 986
Related Topics 987
Persistent NAT Configuration Overview 987
Related Topics 988
Example: Configuring Persistent NAT with Source NAT Address Pool 988
CLI Configuration 988
Example: Configuring Persistent NAT with Interface NAT 989
CLI Configuration 990
Configuring Proxy ARP 991
CLI Configuration 991
Verifying NAT Configuration 991
CLI Configuration 991

Part 12 GPRS

Chapter 30 General Packet Radio Service 995

Understanding GPRS 995
Understanding GPRS Tunneling Protocol on SRX Series Devices 996
Gp and Gn Interfaces 996
Gi Interface 997
Operational Modes 998
Understanding Policy-Based GPRS Tunneling Protocol 998
Configuring Policies to Enable GTP Inspection 999
CLI Configuration 999
Understanding GPRS Tunneling Protocol Inspection Object 1000
Creating a GTP Inspection Object 1000
CLI Configuration 1000
Understanding GTP Message Filtering 1001
Message-Length Filtering 1001
Setting GTP Message Lengths 1001
CLI Configuration 1001
Message-Type Filtering 1002
Permitting and Denying Message Types 1002
CLI Configuration 1002
Supported Message Types 1002
Message-Rate Limiting 1004
Setting a Rate Limit 1005
CLI Configuration 1005
Sequence Number Validation 1005
Enabling Sequence Number Validation 1005
CLI Configuration 1005
IP Fragmentation 1005
Understanding GTP Information Elements 1006
Access Point Name Filtering 1006
Setting an APN and a Selection Mode 1007
IMSI Prefix Filtering 1007
Setting a Combined IMSI Prefix and APN Filter 1008
Supported R6 Information Elements 1008
3GPP R6 IE Removal 1011
R6 Removal 1011
Understanding GGSN Redirection 1011

Part 13 Index

Index 1015

About This Guide

This preface provides the following guidelines for using the *JUNOS Software Security Configuration Guide*:

- J Series and SRX Series Documentation and Release Notes on page xli
- Objectives on page xlii
- Audience on page xlii
- Supported Routing Platforms on page xlii
- How to Use This Manual on page xlii
- Document Conventions on page xliv
- Documentation Feedback on page xlvi
- Requesting Technical Support on page xlvi

J Series and SRX Series Documentation and Release Notes

For a list of related J Series documentation, see http://www.juniper.net/techpubs/software/junos-jseries/index-main.html .

For a list of related SRX Series documentation, see http://www.juniper.net/techpubs/hardware/srx-series-main.html.

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

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

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

Objectives

This guide describes how to use and configure key security features on J Series Services Routers and SRX Series Services Gateways running JUNOS Software. It provides conceptual information, suggested workflows, and examples where applicable.

Audience

This manual is designed for anyone who installs, sets up, configures, monitors, or administers a J Series Services Router or an SRX Series Services Gateway running JUNOS Software. The manual is intended for the following audiences:

- Customers with technical knowledge of and experience with networks and network security, the Internet, and Internet routing protocols
- Network administrators who install, configure, and manage Internet routers

Supported Routing Platforms

This manual describes features supported on J Series Services Routers and SRX Series Services Gateways running JUNOS Software.

How to Use This Manual

This manual and the other manuals in this set explain how to install, configure, and manage:

- JUNOS Software for J Series Services Routers
- JUNOS Software for SRX Series Services Gateways

Table 1 on page xlii identifies the tasks required to configure and manage these devices and shows where to find task information and instructions.

Table 1: Tasks and Related Documentation

Task	Related Documentation
Basic Device Installation and Setup	
■ Reviewing safety warnings and compliance statements ■ Installing hardware and establishing basic connectivity ■ Initially setting up a device	J Series Services Routers: ■ *J Series Services Routers Quick Start* ■ *J Series Services Routers Hardware Guide* ■ *JUNOS Software Release Notes* SRX Series Services Gateways: the appropriate *Services Gateway Getting Started Guide*

Table 1: Tasks and Related Documentation *(continued)*

Task	Related Documentation
Migration from ScreenOS or JUNOS Software (Legacy Services) to JUNOS Software (if necessary)	
■ Migrating from JUNOS Software (legacy services) Release 8.3 or later to JUNOS Software ■ Migrating from ScreenOS Release 5.4 or later to JUNOS Software.	*JUNOS Software Migration Guide* (J Series Services Routers only)
Context—Changing to Secure Context or Router Context	
Changing the device from one context to another and understanding the factory default settings	*JUNOS Software Administration Guide*
Interface Configuration	
Configuring device interfaces	■ *JUNOS Software Interfaces and Routing Configuration Guide* ■ *JUNOS Software CLI Reference*
Deployment Planning and Configuration	
■ Understanding and gathering information required to design network firewalls and IPsec VPNs ■ Implementing a JUNOS Software firewall from a sample scenario ■ Implementing a policy-based IPsec VPN from a sample scenario	*JUNOS Software Design and Implementation Guide* (J Series Services Routers only)
Security Configuration	
Configuring and managing the following security services: ■ Stateful firewall policies ■ Zones and their interfaces and address books ■ IPsec VPNs ■ Firewall screens ■ Interface modes: Network Address Translation (NAT) mode and Router mode ■ Public Key Cryptography (PKI) ■ Application Layer Gateways (ALGs) ■ Chassis clusters ■ Intrusion Detection and Prevention (IDP)	■ *JUNOS Software Security Configuration Guide* ■ *JUNOS Software CLI Reference*
Routing Protocols and Services Configuration	
■ Configuring routing protocols, including static routes and the dynamic routing protocols RIP, OSPF, BGP, and IS-IS ■ Configuring class-of-service (CoS) features, including traffic shaping and policing ■ Configuring packet-based stateless firewall filters (access control lists) to control access and limit traffic rates ■ Configuring MPLS to control network traffic patterns	■ *JUNOS Software Interfaces and Routing Configuration Guide* ■ *JUNOS Software CLI Reference*
WAN Acceleration Module Installation (Optional)	

Table 1: Tasks and Related Documentation *(continued)*

Task	Related Documentation
Installing and initially configuring a WXC Integrated Services Module (ISM 200)	*WXC Integrated Services Module Installation and Configuration Guide* (J Series Services Routers only)
User and System Administration	
■ Administering user authentication and access ■ Monitoring the device, routing protocols, and routing operations ■ Configuring and monitoring system alarms and events, real-time performance (RPM) probes, and performance ■ Monitoring the firewall and other security-related services ■ Managing system log files ■ Upgrading software ■ Diagnosing common problems	*JUNOS Software Administration Guide*
User Interfaces	
■ Understanding and using the J-Web interface ■ Understanding and using the CLI configuration editor	■ *J Series Services Routers Quick Start* (J Series Services Routers only) ■ *JUNOS Software Administration Guide*

Document Conventions

Table 2 on page xliv defines the notice icons used in this guide.

Table 2: 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 3 on page xlv defines the text and syntax conventions used in this guide.

Table 3: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configure mode, type the configure command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host> show chassis alarms` `No alarms currently active`
Italic text like this	■ Introduces important new terms. ■ Identifies book names. ■ Identifies RFC and Internet draft titles.	■ A policy *term* is a named structure that defines match conditions and actions. ■ *JUNOS System Basics Configuration Guide* ■ RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	■ To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. ■ The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric *metric*>;
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast | multicast (*string1* | *string2* | *string3*)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [*community-ids*]

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

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

Table 3: Text and Syntax Conventions *(continued)*

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

Documentation Feedback

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

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

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at http://www.juniper.net/customers/support/downloads/710059.pdf .
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/ .
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/

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

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

Opening a Case with JTAC

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

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

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

Part 1
Introduction to JUNOS Software

- Introducing JUNOS Software for SRX Series Services Gateways on page 3
- Introducing JUNOS Software for J Series Services Routers on page 33

Chapter 1

Introducing JUNOS Software for SRX Series Services Gateways

JUNOS Software for SRX Series Services Gateways integrates the world-class network security and routing capabilities of Juniper Networks. JUNOS Software for the SRX Series Services Gateways includes a wide range of packet-based filtering, class-of-service (CoS) classifiers, and traffic-shaping features as well as a rich, extensive set of flow-based security features including policies, screens, NAT, and other flow-based services.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Stateful and Stateless Data Processing Overview on page 3
- SRX5600 and SRX5800 Services Gateways Overview on page 8
- Network Processor Bundling Limitations on page 21
- SRX3400 and SRX3600 Services Gateways Overview on page 21
- Packet Flow and Session Management in SRX210 Services Gateways on page 26
- Obtaining Information About Sessions and Terminating Them on page 28

Stateful and Stateless Data Processing Overview

Traffic that enters and exits an SRX Series Services Gateway running JUNOS Software is processed according to features you configure, such as packet filters, security policies, and screens. For example, the software can determine:

- Whether the packet is allowed into the device
- Which firewall screens to apply to the packet
- The route the packet takes to reach its destination
- Which class of service (CoS) to apply to the packet, if any
- Whether to apply Network Address Translation (NAT) to translate the packet's IP address
- Whether the packet requires an Application Layer Gateway (ALG)

Packets that enter and exit an SRX Series device undergo both packet-based and flow-based processing.

- Flow-based packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that were established for the first packet of the packet stream, which is referred to as a flow.

  For the distributed processing architecture of the SRX Series Services Gateway, all flow-based processing occurs on the SPU and sampling is multi-thread aware. Packet sequencing is maintained for the sampled packets.

- Packet-based, or stateless, packet processing treats packets discretely. Each packet is assessed individually for treatment.

  For the distributed processing architecture of the SRX Series Services Gateway, some packet-based processing, such as traffic shaping, occurs on the NPU. Some packet-based processing, such as application of classifiers to a packet, occurs on the SPU.

Understanding Flow-Based Processing

A packet undergoes flow-based processing after packet-based filters and some screens have been applied to it. All flow-based processing for a single flow occurs on a single SPU. An SPU processes the packets of a flow according to the security features and other services configured for the session.

Figure 1 on page 4 shows a conceptual view of how flow-based traffic processing occurs on an SRX Series Services Gateway.

Figure 1: Traffic Flow for Flow-Based Processing

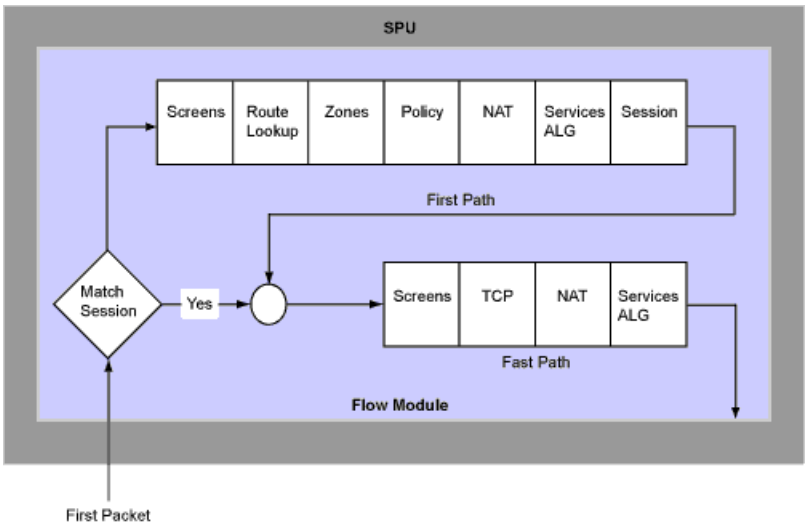

A flow is a stream of related packets that meet the same matching criteria and share the same characteristics. JUNOS Software treats packets belonging to the same flow in the same manner.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, if it requires an Application Layer Gateway (ALG), if Network Address Translation (NAT) is applied to translate the packet's source and/or destination IP address—are assessed for the first packet of a flow.

To determine if a flow exists for a packet, the NPU attempts to match the packet's information to that of an existing session based on the following match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique session token number for a given zone and virtual router

Zones and Policies

The security policy to be used for the first packet of a flow is cached in a flow table for use with the same flow and closely related flows. Security policies are associated with zones. A zone is a collection of interfaces that define a security boundary. A packet's incoming zone, as determined by the interface through which it arrived, and its outgoing zone, as determined by the forwarding lookup, together determine which policy is used for packets of the flow.

Flows and Sessions

Flow-based packet processing, which is stateful, requires the creation of sessions. A session is created for the first packet of a flow for the following purposes:

- To store most of the security measures to be applied to the packets of the flow
- To cache information about the state of the flow

  For example, logging and counting information for a flow is cached in its session. (Some stateful firewall screens rely on threshold values that pertain to individual sessions or across all sessions.)
- To allocate required resources for the flow for features such as Network Address Translation (NAT)
- To provide a framework for features such as Application Layer Gateways (ALGs) and firewall features

Most packet processing occurs in the context of a flow, including:

- Management of policies, NAT, zones, and most screens
- Management of ALGs and authentication

Understanding Packet-Based Processing

A packet undergoes packet-based processing when it is removed from the queue from its input interface and before it is added to the queue on its output interface.

Packet-based processing applies stateless firewall filters, class-of-service (CoS) features, and some screens to discrete packets.

- When a packet arrives at an interface on the SRX Series device, sanity checks, packet-based filters, some CoS features, and some screens are applied to it.
- Before a packet leaves the device, any packet-based filters, some CoS features, and some screens associated with the interface are applied to the packet.

Filters and CoS features are typically associated with one or more interfaces to influence which packets are allowed to transit the system and to apply special actions to packets as necessary.

Here are the kinds of packet-based features that you can configure and apply to transit traffic.

- **Stateless firewall filters**—Also referred to as access control lists (ACLs), stateless firewall filters control access and limit traffic rates. They statically evaluate the contents of packets transiting the device from a source to a destination, or packets originating from or destined for the Routing Engine. A stateless firewall filter evaluates every packet, including fragmented packets.

  You can apply a stateless firewall filter to an input or output interface, or to both. A filter contains one or more terms, and each term consists of two components—match conditions and actions. By default, a packet that does not match a firewall filter is discarded.

  You can plan and design stateless firewall filters to be used for various purposes—for example, to limit traffic to certain protocols, IP source or destination addresses, or data rates. Stateless firewall filters are executed on the SPU.

- **Class-of-service (CoS) features**—CoS features allow you to classify and shape traffic. CoS features are executed on the SPU.

  - **Behavior aggregate (BA) classifiers**—These classifiers operate on packets as they enter the device. Using behavior aggregate classifiers, the device aggregates different types of traffic into a single forwarding class to receive the same forwarding treatment. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Service (DiffServ) value.

  - **Traffic shaping**—You can shape traffic by assigning service levels with different delay, jitter, and packet loss characteristics to particular applications served by specific traffic flows. Traffic shaping is especially useful for real-time applications, such as voice and video transmission.

- **Certain screens**—Some screens, such as denial-of-service (DoS) screens, are applied to a packet outside the flow process. They are executed on the NPU.

For details on specific stateless firewall filters and CoS features, see the *JUNOS Software Interfaces and Routing Configuration Guide* and the *JUNOS Software CLI Reference*.

Changing Session Characteristics

Sessions are created, based on routing and other classification information, to store information and allocate resources for a flow. Sessions have characteristics, some of which you can change, such as when they are terminated. For example, you might want to ensure that a session table is never entirely full to protect against an attacker's attempt to flood the table and thereby prevent legitimate users from starting sessions.

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 30 minutes. The default timeout for UDP is 1 minute. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 10 seconds.

If no traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse. You can affect the life of a session in the following ways:

- You can specify circumstances to terminate sessions by using any of the following methods:
  - Age out sessions based on how full the session table is.
  - Set an explicit timeout for aging out TCP sessions.
  - Configure a TCP session to be invalidated when it receives a TCP RST (reset) message.
- You can configure sessions to accommodate other systems as follows:
  - Disable TCP packet security checks.
  - Change the maximum segment size.

The following sections show you how to modify a session's characteristics. For details, see the *JUNOS Software CLI Reference*.

Controlling Session Termination

JUNOS Software terminates sessions normally in certain situations—for example, after receiving a TCP FINish Close or receiving a RST (reset) message, when encountering Internet Control Message Protocol (ICMP) errors for UDP, and when no matching traffic is received before the service timeout. When sessions are terminated, their resources are freed up for use for other sessions.

To control when sessions are terminated, you configure the SRX Series device to age out sessions after a certain period of time, when the number of sessions in the session table reaches a specified percentage, or both.

- To terminate sessions based on a timeout value or the number of sessions in the session table:
  - You can use the following set security flow command to specify the number of seconds in tens of seconds after which a session is invalidated. The following command ages out sessions after 20 seconds:

    set security flow aging early-ageout 2

- To configure an explicit timeout value, use the following command. This set security flow command removes a TCP session from the session table after 280 seconds.

  set security flow tcp-session tcp-initial-timeout 280

- To cause any session that receives a TCP RST message to be invalidated, use the following command:

  set security flow tcp-session rst-invalidate-session

Disabling TCP Packet Security Checks

The JUNOS Software provides a mechanism to disable security checks on TCP packets to ensure interoperability with hosts and devices with faulty TCP implementations. The following set security flow command disables TCP SYN checks and TCP sequence checks on all TCP sessions.

```
set security flow tcp-session no-syn-check
set security flow tcp-session no-sequence-check
```

Setting the Maximum Segment Size for All TCP Sessions

This command allows you to specify the maximum segment size in TCP SYN packets used during session establishment. Decreasing the maximum segment size helps to limit packet fragmentation and to protect against packet loss that can occur when a packet must be fragmented to meet the MTU size but the packet's DF-bit (don't fragment) is set.

- The following command configures the TCP maximum segment size to 1400 bytes for all TCP sessions:

  set security flow tcp-mss all-tcp 1400

SRX5600 and SRX5800 Services Gateways Overview

JUNOS Software for SRX Series Services Gateways is a distributed parallel processing high throughput and high performance system. The distributed parallel processing architecture of the SRX Series Services Gateways includes multiple processors to manage sessions and run security and other services processing. This architecture provides greater flexibility and allows for high throughput and fast performance.

NOTE: In SRX3400, SRX3600, SRX5600, and SRX5800 devices, IKE negotiations involving NAT traversal do not work if the IKE peer is behind a NAT device that will change the source IP address of the IKE packets during the negotiation. For example, if the NAT device is configured with DIP, it changes the source IP because the IKE protocol switches the UDP port from 500 to 4500.

The SRX5600 and SRX5800 Services Gateways include I/O cards (IOC) and Services Processing Cards (SPCs) that each contain processing units that process a packet as it traverses the device. A Network Processing Unit (NPU) runs on an IOC. An IOC has one or more NPUs. One or more Services Processing Units (SPUs) run on an SPC.

These processing units have different responsibilities. All flow-based services for a packet are executed on a single SPU. Otherwise, however, the lines are not clearly divided in regard to the kinds of services that run on these processors. (For details on flow-based processing, see "Understanding Flow-Based Processing" on page 4.) For example:

- An NPU processes packets discretely. It performs sanity checks and applies some screens that are configured for the interface, such as denial-of-service (DoS) screens, to the packet.
- An SPU manages the session for the packet flow and applies security features and other services to the packet. It also applies packet-based stateless firewall filters, classifiers, and traffic shapers to the packet.
- The system uses one processor as a central point (CP) to take care of arbitration and allocation of resources and distribute sessions in an intelligent way. The central point assigns an SPU to be used for a particular session when the first packet of its flow is processed.

These discrete, cooperating parts of the system, including the central point, each store the information identifying whether a session exists for a stream of packets and the information against which a packet is matched to determine if it belongs to an existing session.

This architecture allows the device to distribute processing of all sessions across multiple SPUs. It also allows an NPU to determine if a session exists for a packet, to check the packet, and to apply screens to it. How a packet is handled depends on whether it is the first packet of a flow.

If the packet matches an existing flow, processing for the packet is assessed in the context of its flow state. The SPU maintains the state for each session, and the settings are then applied to the rest of the packets in the flow. If the packet does not match an existing flow, it is used to create a flow state and a session is allocated for it.

Figure 2 on page 10 illustrates the path the first packet of a flow takes as it enters the device: the NPU determines that no session exists for the packet, and the NPU sends the packet to the central point; the central point selects the SPU to set up the session for the packet and process it, and it sends the packet to that SPU. The SPU processes the packet and sends it to the NPU for transmission from the device. (This high-level description does not address application of features to a packet.)

Figure 2: First-Packet Processing

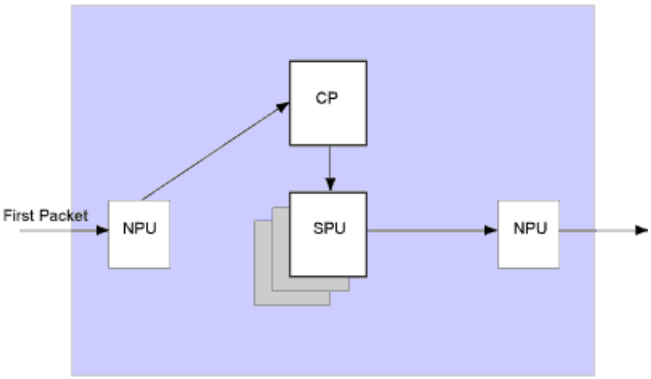

For details on session creation for the first packet in a flow, see “Understanding Session Creation: First-Packet Processing” on page 12.

After the first packet of a flow has traversed the system and a session has been established for it, it undergoes fast-path processing.

Subsequent packets of the flow also undergo fast-path processing; in this case, after each packet enters the session and the NPU finds a match for it in its session table, the NPU forwards the packet to the SPU that manages its session.

Figure 3 on page 10 illustrates fast-path processing. This is the path a packet takes when a flow has already been established for its related packets. (It is also the path that the first packet of a flow takes after the session for the flow that the packet initiated has been set up.) After the packet enters the device, the NPU finds a match for the packet in its session table, and it forwards the packet to the SPU that manages the packet’s session. Note that the packet bypasses interaction with the central point.

Figure 3: Fast-Path Processing

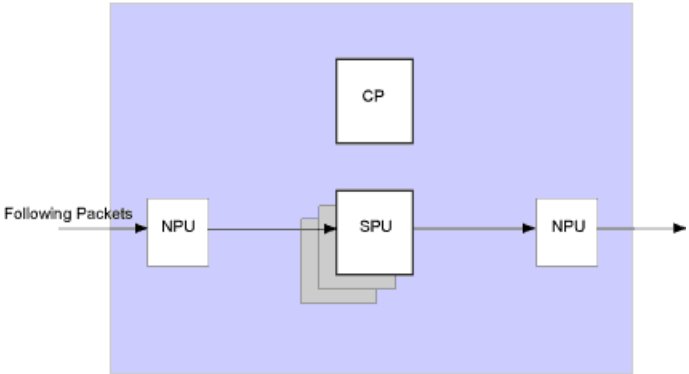

This section explains how a session is created and the process a packet undergoes as it transits the SRX Series device.

Here is an overview of the main components involved in setting up a session for a packet and processing the packets both discretely and as part of a flow as they transit the SRX5600 and SRX5800 devices:

- **Network Processing Units (NPUs)**—NPUs reside on I/O cards. They handle packet sanity checking and application of some screens. NPUs maintain session tables that they use to determine if a session exists for an incoming packet or for reverse traffic.

  The NPU session table contains an entry for a session if the session is established on an SPU for a packet that had previously entered the device via the interface and was processed by this NPU. The SPU installs the session in the NPU table when it creates the session.

  An NPU determines if a session exists for a packet by checking the packet information against its session table. If the packet matches an existing session, the NPU sends the packet and the metadata for it to the SPU. If there is no session, the NPU sends the packet to the central point for SPU assignment.

- **Services Processing Units (SPUs)**—The main processors of the SRX5600 and SRX5800 devices reside on Services Processing Cards (SPCs). They establish and manage traffic flows and perform most of the packet processing on a packet as it transits the device. Each SPU maintains a hash table for fast session lookup. The SPU applies stateless firewall filters, classifiers, and traffic shapers to traffic. An SPU performs all flow-based processing for a packet and most packet-based processing. Each multicore SPU processes packets independently with minimum interaction among SPUs on the same or different SPC. All packets that belong to the same flow are processed by the same SPU.

  The SPU maintains a session table with entries for all sessions that it established and whose packets it processes. When an SPU receives a packet from an NPU, it checks its session table to ensure that the packet belongs to it. It also checks its session table when it receives a packet from the central point (CP) and a message to establish a session for that packet to verify that there is not an existing session for the packet.

- **Central point (CP)**—The SRX Series device uses one processor as a central point to take care of arbitration and allocation of resources and distribute sessions in an intelligent way to avoid multiple SPUs from wrongly handling the same flow.

  The central point's main function is to delegate session processing to one of the SPUs. If the session has not yet been established, the central point selects an SPU to establish the session for the flow, based on load- balancing criteria. If the session already exists, the central point forwards packets for that flow to the SPU hosting it. It also redirects packets to the correct SPU in the event that the NPU fails to do so.

  The central point maintains a global session table with information about the owner SPU of a particular session. It functions as a central repository and resource manager for the whole system.

- **Routing Engine (RE)**—The Routing Engine runs the control plane.

Following the Data Path for a Unicast Session

This section describes the process that the SRX Series device undertakes in establishing a session for packets belonging to a flow that transits the SRX Series devices.

To illustrate session establishment and the packet "walk" including the points at which services are applied to the packets of a flow, this example uses the simple case of a unicast session.

This packet "walk" brings together the packet-based processing and flow-based processing that the JUNOS Software performs on the packet. It includes the following sections:

- Session Lookup and Packet Match Criteria on page 12
- Understanding Session Creation: First-Packet Processing on page 12
- Understanding Fast-Path Processing on page 15

Session Lookup and Packet Match Criteria

To determine if a packet belongs to an existing flow, the SRX Series device attempts to match the packet's information to that of an existing session based on the following six match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique token from a given zone and virtual router

Understanding Session Creation: First-Packet Processing

This section explains how a session is set up to process the packets composing a flow. To illustrate the process, this section uses an example with a source "a" and a destination "b". The direction from source to destination for the packets of the flow is referred to as (a -> b). The direction from destination to source is referred to as (b-> a).

Step 1. A Packet Arrives at an Interface on the Device and the NPU Processes It.

This section describes how a packet is handled when it arrives at an SRX Series device ingress IOC.

1. The packet arrives at the device's IOC and is processed by the NPU on the card.
2. The NPU performs basic sanity checks on the packet and applies some screens configured for the interface to the packet.

3. The NPU checks its session table for an existing session for the packet. (It checks the packet's tuple against those of packets for existing sessions in its session table.)

   a. If no existent session is found, the NPU forwards the packet to the central point.

   b. If a session match is found, the session has already been created on an SPU that was assigned to it, so the NPU forwards the packet to the SPU for processing along with the session ID. (See "Understanding Fast-Path Processing" on page 15.)

Example: Packet (a -> b) arrives at NPU1. NPU1 performs sanity checks and applies DoS screens to the packet. NPU checks its session table for a tuple match and no existing session is found. NPU1 forwards the packet to the central point for assignment to an SPU.

Step 2. The Central Point (CP) Creates a Session with a "Pending" State.

The central point maintains a global session table that includes entries for all sessions that exist across all SPUs on the device. It participates in session creation and delegates and arbitrates session resources allocation.

This process entails the following parts:

1. The central point checks its session table and gate table to determine if a session or a gate exists for the packet it receives from the NPU. (An NPU has forwarded a packet to the central point because its table indicates there is no session for it. The central point verifies this information before allocating an SPU for the session.)
2. If there is no entry that matches the packet in either table, the central point creates a pending wing for the session and selects an SPU to be used for the session, based on its load-balancing algorithm.
3. The central point forwards the first packet of the flow to the selected SPU in a message telling it to set up a session locally to be used for the packet flow.

Example: The central point creates pending wing (a -> b) for the session. It selects SPU1 to be used for it. It sends SPU1 the (a-> b) packet along with a message to create a session for it.

Step 3. The SPU Sets Up the Session.

Each SPU, too, has a session table, which contains information about its sessions. When the SPU receives a message from the central point to set up a session, it checks its session table to ensure that a session does not already exist for the packet.

1. If there is no existing session for the packet, the SPU sets up the session locally.
2. The SPU sends a message to the central point telling it to install the session.

NOTE: During first-packet processing, if NAT is enabled, the SPU allocates IP address resources for NAT. In this case, the first-packet processing for the session is suspended until the NAT allocation process is completed.

The SPU adds to the queue any additional packets for the flow that it might receive until the session has been installed.

Example: SPU1 creates the session for (a -> b) and sends a message back to the central point telling it to install the pending session.

Step 4. The Central Point Installs the Session.

The central point receives the install message from the SPU.

1. It sets the state for the session's pending wing to active
2. It installs the reverse wing for the session as an active wing.
3. It sends an ACK (acknowledge) message to the SPU, indicating that the session is installed.

Example: The central point receives a message from SPU1 to install the session for (a-> b). It sets the session state for (a-> b) wing to active. It installs the reverse wing (b-> a) for the session and makes it active; this allows for delivery of packets from the reverse direction of the flow: destination (b) to be delivered to the source (a).

Step 5. The SPU Sets Up the Session on the Ingress and Egress NPUs.

NPUs maintain information about a session for packet forwarding and delivery. Session information is set up on the egress and ingress NPUs (which sometimes are the same) so that packets can be sent directly to the SPU that manages their flows and not to the central point for redirection.

Step 6. Fast-Path Processing Takes Place.

For the remainder of the steps entailed in packet processing, proceed to Step 1 in "Understanding Fast-Path Processing" on page 15.

Figure 4 on page 15 illustrates the first part of the process the first packet of a flow undergoes after it reaches the SRX Series device. At this point a session is set up to process the packet and the rest of the packets belonging to its flow. Subsequently, it and the rest of the packets of flow undergo fast-path processing.

Figure 4: Session Creation: First-Packet Processing

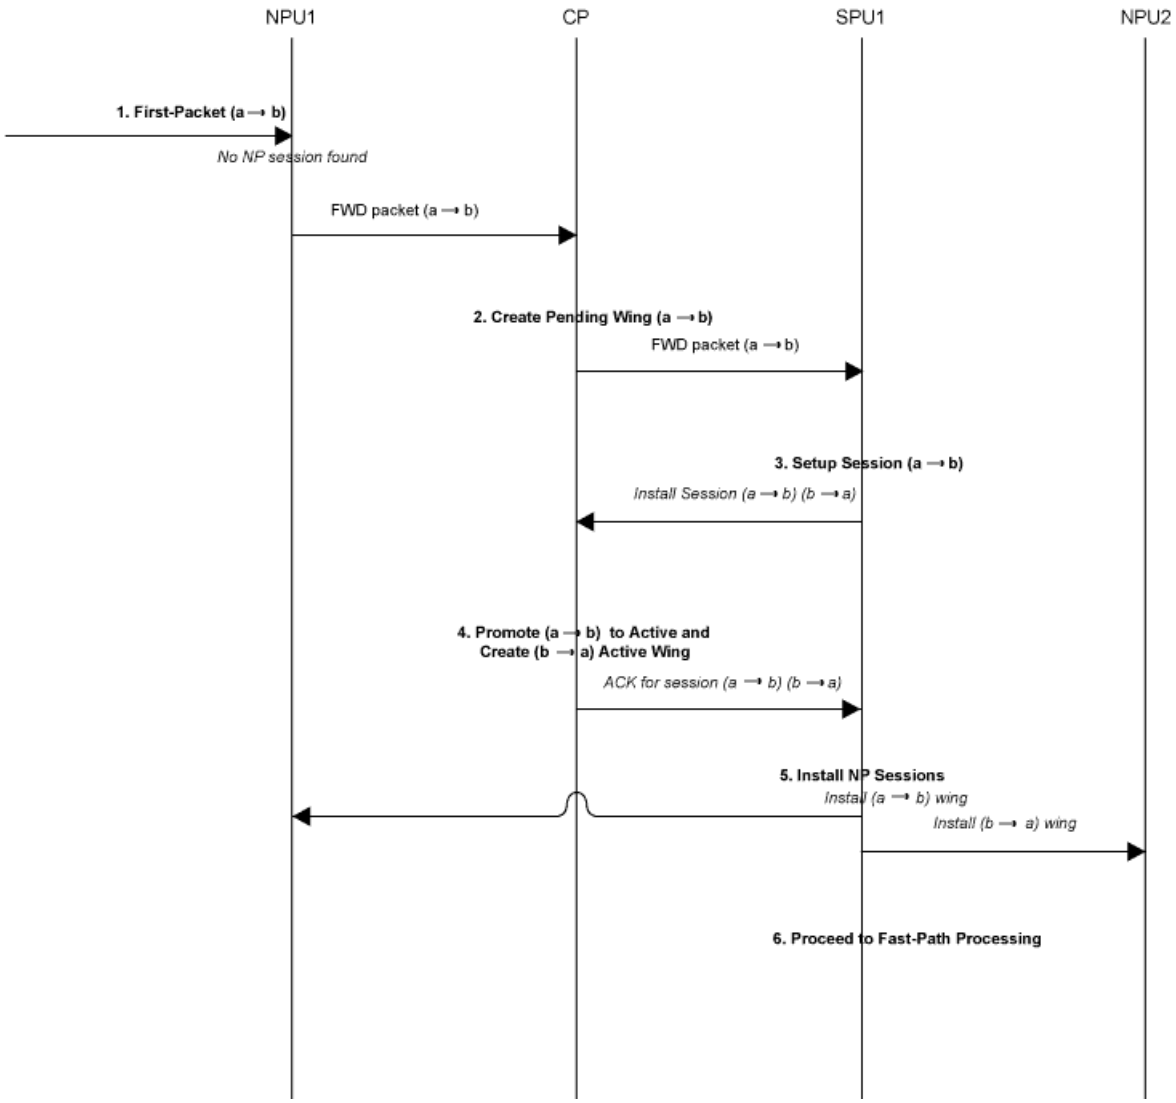

Understanding Fast-Path Processing

All packets undergo fast path-processing. However, if a session exists for a packet, the packet undergoes fast-path processing and bypasses the first-packet process. When there is already a session for the packet's flow, the packet does not transit the central point.

Here is how fast-path processing works: NPUs at the egress and ingress interfaces contain session tables that include the identification of the SPU that manages a packet's flow. Because the NPUs have this session information, all traffic for the flow, including reverse traffic, is sent directly to that SPU for processing.

To illustrate the fast-path process, this section uses an example with a source "a" and a destination "b". The direction from source to destination for the packets of

the flow is referred to as (a- > b). The direction from destination to source is referred to as (b- > a).

Step 1. A Packet Arrives at the Device and the NPU Processes It.

This section describes how a packet is handled when it arrives at a service gateway's IOC.

1. The packet arrives at the SRX Series device's IOC and is processed by the NPU on the card.

   The NPU performs sanity checks and applies some screens, such as denial-of-service (DoS) screens, to the packet.
2. The NPU identifies an entry for an existing session in its session table that the packet matches.
3. The NPU forwards the packet along with metadata from its session table, including the session ID and packet tuple information, to the SPU that manages the session for the flow, applies stateless firewall filters and CoS features to its packets, and handles the packet's flow processing and application of security and other features.

Example: Packet (a - > b) arrives at NPU1. NPU1 performs sanity checks on the packet, applies DoS screens to it, and checks its session table for a tuple match. It finds a match and that a session exists for the packet on SPU1. NPU1 forwards the packet to SPU1 for processing.

Step 2. The SPU for the Session Processes the Packet.

Most of a packet's processing occurs on the SPU to which its session is assigned. The packet is processed for packet-based features such as stateless firewall filters, traffic shapers, and classifiers, if applicable. Configured flow-based security and related services such as firewall features, NAT, ALGs, and so forth, are applied to the packet. (For information on how security services are determined for a session, see "Zones and Policies" on page 5.)

1. Before it processes the packet, the SPU checks its session table to verify that the packet belongs to one of its sessions.
2. The SPU processes the packet for applicable features and services.

Example: SPU1 receives packet (a- > b) from NPU1. It checks its session table to verify that the packet belongs to one of its sessions. Then it processes packet (a - > b) according to input filters and CoS features that apply to its input interface. The SPU applies the security features and services that are configured for the packet's flow to it, based on its zone and policies. If any are configured, it applies output filters, traffic shapers and additional screens to the packet.

Step 3. The SPU Forwards the Packet to the NPU.

1. The SPU forwards the packet to the NPU.

2. The NPU applies any applicable screens associated with the interface to the packet.

Example: SPU1 forwards packet (a -> b) to NPU2, and NPU2 applies DoS screens.

Step 4. The Interface Transmits the Packet From the Device.

Example: The interface transmits packet (a-> b) from the device.

Step 5. A Reverse Traffic Packet Arrives at the Egress Interface and the NPU Processes It.

This step mirrors Step 1 exactly in reverse. See Step 1 in this section for details.

Example: Packet (b-> a) arrives at NPU2. NPU2 checks its session table for a tuple match. It finds a match and that a session exists for the packet on SPU1. NPU2 forwards the packet to SPU1 for processing.

Step 6. The SPU for the Session Processes the Reverse Traffic Packet.

This step is the same as Step 2 except that it applies to reverse traffic. See Step 2 in this section for details.

Example: SPU1 receives packet (b-> a) from NPU2. It checks its session table to verify that the packet belongs to the session identified by NPU2. Then it applies packet-based features configured for the NPU1's interface to the packet. It processes packet (b-> a) according to the security features and other services that are configured for its flow, based on its zone and policies. (See "Zones and Policies" on page 5.)

Step 7. The SPU Forwards the Reverse Traffic Packet to the NPU.

This step is the same as Step 3 except that it applies to reverse traffic. See Step 3 in this section for details.

Example: SPU1 forwards packet (b-> a) to NPU1. NPU1 processes any screens configured for the interface.

8. The Interface Transmits the Packet From the Device.

This step is the same as Step 4 except that it applies to reverse traffic. See Step 4 in this section for details.

Example: The interface transmits packet (b-> a) from the device.

Figure 5 on page 18 illustrates the process a packet undergoes when it reaches the SRX Series device and a session exists for the flow that the packet belongs to.

Figure 5: "Packet Walk" for Fast Path Processing

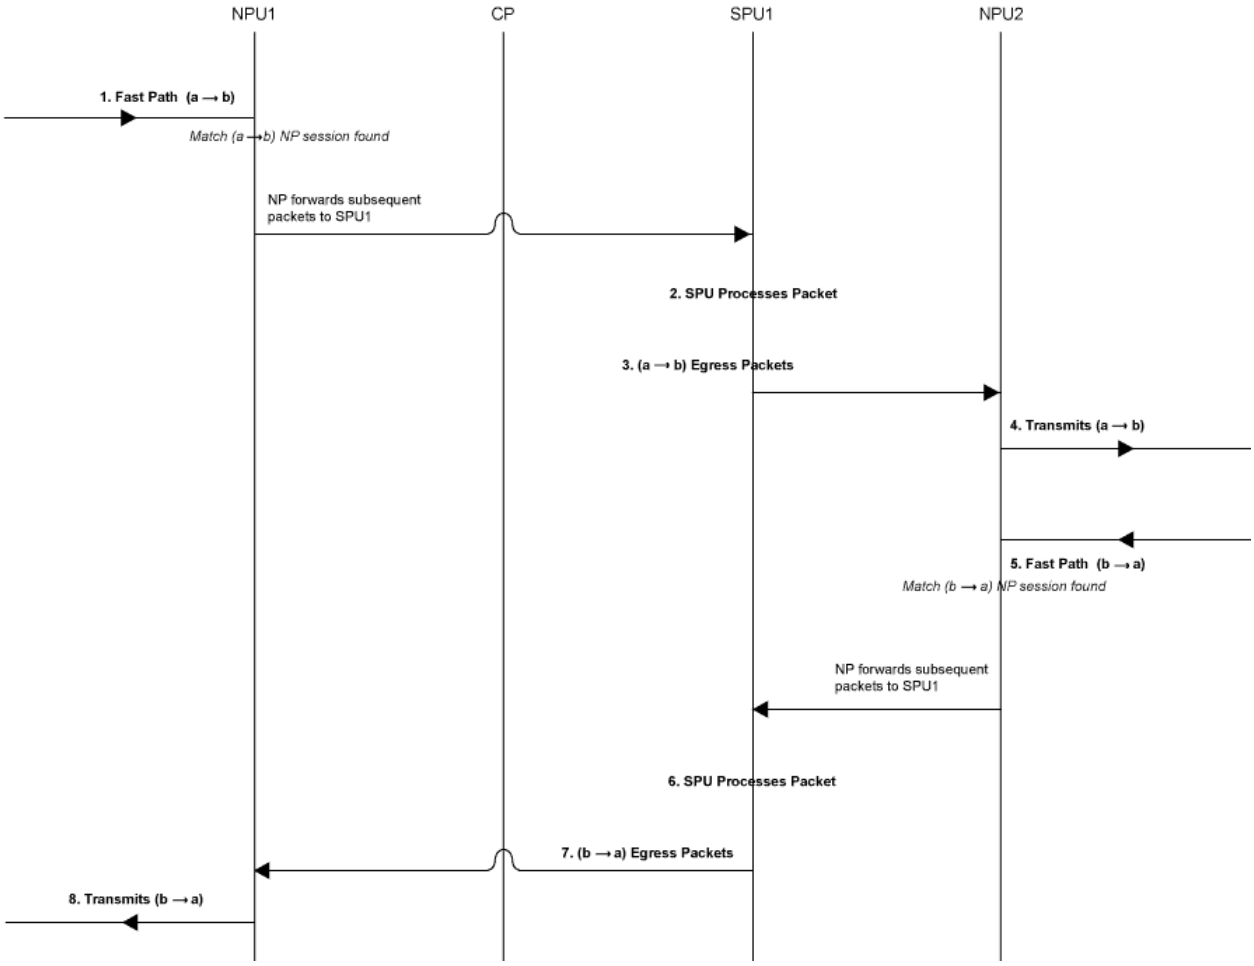

Data Path Debugging

Data path debugging provides tracing and debugging at multiple processing units along the packet-processing path. The packet filter can be executed with minimal impact to the production system.

In data path debugging, a packet goes through multiple Services Processing Units (SPUs). At the same time, several Flexible PIC Concentrator (FPC) I/O cards (IOCs) provide EZchip ingress and egress traffic management. JUNOS Software supports IOC for filter-based, per-packet counting and logging to record the processing path of a packet. Only the matched packets are traced by the IOC EZchip ingress, EZchip egress, load-balancing thread (LBT), and packet-ordering thread (POT).

The following events are defined for the SRX5600 and SRX5800 devices in the packet-processing path:

- ezchip ingress
- ezchip egress

- spu.lbt
- spu.pot

CLI Configuration

To configure the device for data path debugging:

1. Specify the following request command to set data path debugging for the multiple processing units along the packet-processing path:

   user@host# **set security datapath-debug**

2. Specify the trace options for datapath-debug using the following command:

   user@host# **set security datapath-debug traceoptions**

3. Using the request security packet-filter command, you can set the packet filter to specify the related packets to perform datapath-debug action. A maximum of four filters are supported at the same time. For example, the following command sets the first packet-filter:

   user@host# **set security datapath-debug packet-filter <name>**

4. Using the request security action-profile command, you can set the action for the packet match for a specified filter. Only the default action profile is supported, which is the trace option for network processor ezchip ingress, ezchip egress, spu.lbt, and spu.pot:

   user@host# **set security datapath-debug packet-filter <name> action-profile**

Network Processor Bundling

The network processor bundling feature is available on SRX5600 and SRX5800 Services Gateways. This feature enables distribution of data traffic from one interface to multiple network processors for packet processing. A primary network processor is assigned for an interface that receives the ingress traffic and distributes the packets to several other secondary network processors. A single network processor can act as a primary network processor or a secondary network processor to multiple interfaces. A single network processor can join only one network processor bundle.

Understanding Packet Processing

This section explains how a session is set up to process the packets composing a flow:

1. The primary network processor receives the IP packet and analyzes it to obtain several key tuples.
2. The interface is configured to enable distributed flow lookup and the packet is forwarded to a secondary network processor by referring a 5-tuple hash algorithm using the hash value of the primary network processor.

3. The secondary network processor receives the forwarded packet and performs all the tasks for a packet.

NOTE: When network processor bundling is enabled, screen settings are programmed into each network processor in the bundle, the settings chosen will affect the screen performance. For example, if network processor bundling is not enabled, and an ICMP screen is set to 100, ICMP packets will be blocked when the processing rate exceeds 100 pps. However, if network processor bundling is enabled and four network processors are bundled together (one primary network processor and three secondary network processors), the ICMP screen would not begin until the ICMP flood exceeds 300 pps, if the ICMP flooding packets are distributed evenly among the three secondary network processors.

Services Processing Units

For a given physical interface, the Services Processing Unit (SPU) receives ingress packets from all network processors of the network processor bundle associated to the physical interface. The SPU extracts network processor bundle information from the physical interface and uses the same 5-tuple hash algorithm to map a flow to a network processor index. To determine the network processor, the SPU does a lookup on the network processor index in the network processor bundle. The SPU sends egress packets to the physical interface's local PIC for the outward traffic.

NOTE: The network processor and the SPU use the same 5-tuple hash algorithm to get the hash values for the packets.

Interface Changes

For SRX5600 and SRX5800 devices, the IOC supports the following hierarchical scheduler characteristics:

- IFL – The configuration of the network processor bundle is stored in the physical interface data structure. The SRX5600 and SRX5800 devices have a maximum of 48 PICs. The physical interface can use a 48-bit bit-mask to indicate the PIC, or the network processor traffic from this physical interface is distributed in addition to the physical interface's primary network processor.
- IFD – The logical interface associated to the physical interface of a network processor bundle is passed to all the I/O cards (IOCs) that have a PIC in the network processor bundle.

Network Processor Bundling Limitations

Network processor bundling functionality has the following limitations:

- Network processor bundling allows a total of 16 PICs per bundle and eight different network processor bundles system.
- You need to reboot the device to apply the configuration changes on the bundle.
- Network processor bundling is below the reth interface in the overall architecture. You can choose one or both the interfaces from the network processor bundling to form the reth interface.
- If the IOC is removed from a network processor bundle, the packets forwarded to the PIC on that IOC is lost.
- When the network processor bundle is enabled, the ICMP, UDP and TCP sync flooding thresholds no longer apply to an interface. Packets are distributed to multiple network processors for processing. These thresholds will apply to each network processor in the network processor bundle.
- Network processor bundle is not supported in the Layer 2 mode.
- Due to memory constraints on the EZchip, the number of network processor bundled ports that are supported per PIC is limited. Within the network processor bundle, each port needs to have a global port index. The global port index is calculated using the following formula:

  Global_port_index = (global_pic * 16) + port_offset

SRX3400 and SRX3600 Services Gateways Overview

JUNOS Software for the SRX3400 and SRX3600 Services Gateways integrates the world-class network security and routing capabilities of Juniper networks. JUNOS Software for these service gateways includes the wide range of security services including policies, screens, NAT, class-of-service classifiers, and the rich, extensive set of flow-based services that are also supported on the other devices in the SRX Series Services Gateways

The distributed parallel processing architecture of the SRX3400 and SRX3600 devices includes multiple processors to manage sessions and run security and other services processing. This architecture provides greater flexibility and allows for high throughput and fast performance.

Here is an overview of the main components involved in setting up a session for a packet and processing the packets as they transit the SRX3400 and SRX3600 devices:

Services Processing Units (SPUs)—The main processors of the SRX3400 and SRX3600 devices reside on Services Processing Cards (SPCs). They establish and manage traffic flows and perform most of the packet processing on a packet as it transits the device. Each SPU maintains a hash table for fast session lookup. The SPU performs all flow-based processing for a packet, including application of security services, classifiers, and traffic shapers. All packets that belong to the same flow are processed by the same SPU.

The SPU maintains a session table with entries for all sessions that it established and whose packets it processes. When an SPU receives a packet from an NPU, it checks its session table to ensure that the packet belongs to it.

For SRX3400 and SRX3600 devices, one SPU acts in concert performing its regular session management and flow processing functions and acting as a central point in which it arbitrates sessions and allocates resources. When an SPU performs in this manner it is said to be in combo mode.

Central Point (CP)—The central point is used to allocate session management to SPUs based on load balancing criteria. It distributes sessions in an intelligent way to avoid occurrences in which multiple SPUs might wrongly handle the same flow. The central point follows load balancing criteria in allocating sessions to SPUs. If the session exists, the central point forwards packets for that flow to the SPU hosting it. It also redirects packets to the correct SPU in the event that the NPU fails to do so.

For the SRX3400 and SRX3600 devices, one SPU always runs in what is referred to as combo-mode in which it implements both the functionality of the central point and the flow and session management functionality. In combo-mode, the SPU and the central point share the same load-balancing thread (LBT) and packet-ordering thread (POT) infrastructure. For more information, see “Central Point and Combo-Mode Support” on page 22.

The central point maintains a global session table with information about the owner SPU of a particular session. It functions as a central repository and resource manager for the whole system.

Routing Engine (RE)—The routing engine runs the control plane and manages the Control Plane Processor (CPP).

Central Point and Combo-Mode Support

The central point (CP) in the architecture has two basic flow functionalities: load balancing and traffic identification (global session matching). The central point forwards a packet to its Services Processing Unit (SPU) upon session matching, or distributes traffic to an SPU for security processing if the packet does not match any existing session.

An SPU dedicated to central point functionality is called a large central point. However, when such a dedicated SPU is not affordable, you can configure the SPU to perform normal flow processing as well as the functions of the central point. When an SPU functions in such a dual manner, it is said to be in combination or combo mode. In combo mode, the central point and SPU share the same load-balancing thread (LBT) and packet-ordering thread (POT) infrastructure. SRX3400 and SRX3600 devices always support combo-mode.

Load Distribution in Combo Mode

The central point maintains SPU mapping table (for load distribution) that lists live SPUs with the logic SPU IDs mapped to the physical Trivial Network Protocol (TNP) addresses mapping. In combo mode, the SPU that hosts the central point is included

in the table. The load distribution algorithm is adjusted based on session capacity and processing power to avoid overloading of sessions.

Sharing Processing Power and Memory in Combo Mode

The CPU processing power in a combo-mode SPU is shared based on the platform and the number of SPUs in the system. Similarly, the CPU memory is also shared between the central point and SPU.

An SPU has multiple cores (CPUs) for networking processing. In "small" SPU combo mode, CPU functionality takes a small portion of the cores, whereas "medium" SPU combo-mode requires a larger portion of cores. The processing power for central point functionalities and flow processing is shared, based on the number of SPUs, as shown in Table 1.

Table 4: Combo Mode Processing:

Number of SPUs	1	2	3	4 or More than 4
SRX3400	Small	Medium	Medium	Medium
SRX3600	Small	Medium	Medium	Medium

Following the Data Path for a Unicast Session

JUNOS Software for the SRX3400 and SRX3600 Services Gateways is a distributed parallel processing high throughput and high performance system. This section describes the process that the SRX Series device undertakes in establishing a session for packets belonging to a flow that transits the device.

To illustrate session establishment and the packet “walk” including the points at which services are applied to the packets of a flow, this example uses the simple case of a unicast session.

This packet “walk” brings together the packet-based processing and flow-based processing that the JUNOS Software performs on the packet.

Session Lookup and Packet Match Criteria

To determine if a packet belongs to an existing flow, the SRX Series device attempts to match the packet's information to that of an existing session based on the following six match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique token from a given zone and virtual router

Understanding Session Creation: First Packet Processing

This section explains how a session is set up to process the packets composing a flow. To illustrate the process, this section uses an example with a source "a" and a destination "b". The direction from source to destination for the packets of the flow is referred to as (a -> b). The direction from destination to source is referred to as (b -> a).

1. A Packet Arrives at an interface on the device and the IOC processes it.

   The IOC dequeues the packet and sends it to the NPU with which it communicates.

2. The NPU receives the packet from the IOC and processes it.

   a. The NPU performs basic sanity checks on the packet and applies some screens configured for the interface to the packet.

   b. If a session match is found, the session has already been created on an SPU that was assigned to it, so the NPU forwards the packet to the SPU for processing along with the session ID.

   Example: Packet (a ->b) arrives at NPU1 from IOC1. NPU1 performs sanity checks and applies DoS screens to the packet. NPU1 checks its session table for a tuple match and no existing session is found. NPU1 forwards the packet to the central point on SPU1 for assignment to an SPU.

3. The Central Point (CP) Creates a Session with a "Pending" State.

   The central point maintains a global session table that includes entries for all sessions that exist across all SPUs on the device. It participates in session creation and delegates and arbitrates session resources allocation.

   This process entails the following parts:

   a. The central point checks its session table and gate table to determine if a session or a gate exists for the packet it receives from the NPU. (An NPU has forwarded a packet to the central point because its table indicates there is

no session for it. The central point verifies this information before allocating an SPU for the session.)

b. If there is no entry that matches the packet in either table, the central point creates a pending wing for the session and selects an SPU to be used for the session, based on its load-balancing algorithm.

c. The central point forwards the first packet of the flow to the selected SPU in a message telling it to set up a session locally to be used for the packet flow.

Example: The central point creates pending wing (a -> b) for the session. It selects SPU1 to be used for the session. It sends SPU1 the (a-> b) packet along with a message to create a session for it. (It happens to be the case that SPU1 is the SPU that runs in combo mode. Therefore, its session-management and flow-processing services are used for the session.)

4. The SPU Sets Up the Session.

Each SPU, too, has a session table, which contains information about its sessions. When the SPU receives a message from the central point to set up a session, it checks its session table to ensure that a session does not already exist for the packet.

a. If there is no existing session for the packet, the SPU sets up the session locally.

b. The SPU sends a message to the central point, telling it to install the session.

During first-packet processing, if NAT is enabled, the SPU allocates IP address resources for NAT. In this case, the first-packet processing for the session is suspended until the NAT allocation process is completed.

The SPU adds to the queue any additional packets for the flow that it might receive until the session has been installed.

Example: SPU1 creates the session for (a -> b) and sends a message back to the central point (implemented on the same SPU) telling it to install the pending session.

5. The Central Point Installs the Session.

The central point receives the install message from the SPU.

a. It sets the state for the session's pending wing to active.

b. It installs the reverse wing for the session as an active wing.

c. It sends an ACK (acknowledge) message to the SPU, indicating that the session is installed.

Example: The central point receives a message from SPU1 to install the session for (a-> b). It sets the session state for (a-> b) wing to active. It installs the reverse wing (b-> a) for the session and makes it active; this allows for delivery of packets

from the reverse direction of the flow: destination (b) to be delivered to the source (a).

6. The SPU Sets Up the Session on the Ingress and Egress NPUs.

   NPUs maintain information about a session for packet forwarding and delivery. Session information is set up on the egress and ingress NPUs (which sometimes are the same) so that packets can be sent directly to the SPU that manages their flows and not to the central point for redirection.

7. Fast-Path Processing Takes Place.

   For the remainder of the steps entailed in packet processing, proceed to Step 1 in "Understanding Fast-Path Processing".

Understanding Fast-Path Processing

All packets undergo fast path-processing. However, if a session exists for a packet, the packet undergoes fast-path processing and bypasses the first-packet process. When there is already a session for the packet's flow, the packet does not transit the central point.

Here is how fast-path processing works: NPUs at the egress and ingress interfaces contain session tables that include the identification of the SPU that manages a packet's flow. Because the NPUs have this session information, all traffic for the flow, including reverse traffic, is sent directly to that SPU for processing.

Packet Flow and Session Management in SRX210 Services Gateways

This section describes the process that the SRX210 Services Gateway undertakes in establishing a session for packets belonging to a flow that transits the device. The flow services of the SRX210 device are single-threaded and non-distributed. Although it differs from the other SRX Series devices in this respect, the same flow model is followed and the same command line interface (CLI) is implemented.

To illustrate session establishment and the packet "walk" including the points at which services are applied to the packets of a flow, this example uses the simple case of a unicast session.

Flow Processing and Session Management

This section explains how a session is set up to process the packets composing a flow. In the following sections, the SPU refers to the dataplane thread of the SRX210 Services Gateway.

At the outset, the dataplane thread fetches the packet and performs basic sanity checks on it. Then it processes the packet for stateless filters and CoS classifiers and applies some screens.

First-Packet Processing

To determine if a packet belongs to an existing flow, the SRX Series device attempts to match the packet's information to that of an existing session based on the following six match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique token from a given zone and virtual router

The SPU checks its session table for an existing session for the packet. If no existent session is found, the SPU sets up a session for the flow. If a session match is found, the session has already been created, so the SPU performs fast-path processing on the packet.

Session Creation

In setting up the session, the SPU executes the following services for the packet:

- Screens
- Route lookup
- Policy lookup
- Service lookup
- NAT, if required

After a session is set up, it is used for all packets belonging to the flow. Packets of a flow are processed according to the parameters of its session. For the remainder of the steps entailed in packet processing, proceed to Step 1 in "Fast-Path Processing". All packets undergo fast-path processing.

Fast-Path Processing

If a packet matches a session, JUNOS Software for SRX Series Services Gateways performs fast-path processing as described in the following steps. After a session has been set up for the first packet in a flow, also undergoes fast-path processing. All packets undergo fast-path processing.

1. The SPU applies flow-based security features to the packet.
   - Configured screens are applied.
   - TCP checks are performed.
   - Flow services, such as NAT, ALG, and IPsec are applied, if required.

2. The SPU prepares the packet for forwarding and transmits it.
   - Routing packet filters are applied.
   - Traffic shaping is applied.
   - Traffic prioritizing is applied.
   - Traffic scheduling is applied.
   - The packet is transmitted.

Obtaining Information About Sessions and Terminating Them

You can use the commands described in this section to display information about sessions and how to terminate them.

Obtaining Information About Sessions

You can obtain information about the sessions and packet flows active on your device, including detailed information about specific sessions. (The SRX Series device also displays information about failed sessions.) You can display this information to observe activity and for debugging purposes. For example, you can use the show security flow session command:

- To display a list of incoming and outgoing IP flows, including services
- To show the security attributes associated with a flow, for example, the policies that apply to traffic belonging to that flow
- To display the session timeout value, when the session became active, for how long it has been active, and if there is active traffic on the session

For detailed information about this command, see the *JUNOS Software CLI Reference*.

Session information can also be logged if a related policy configuration includes the logging option. See “Monitoring Policy Statistics” on page 94 for how to enable logging in a policy configuration. See “Information Provided in Session Log Entries” on page 30 for details about session information provided in system logs.

Displaying Global Session Parameters

You can use the following command to obtain information about configured parameters that apply to all flows, or sessions:

```
show security flow
```

The show security flow configuration command displays the following information:

- allow-dns-reply—Identifies if unmatched incoming Domain Name System (DNS) reply packets are allowed.
- route-change-timeout—If enabled, displays the session timeout value to be used on a route change to a nonexistent route.

- tcp-mss—Shows the current configuration for the TCP maximum segment size value to be used for all TCP packets for network traffic.
- tcp-session—Displays all configured parameters that control session parameters, examples of which are described in “Changing Session Characteristics” on page 7.
- syn-flood-protection-mode—Displays the SYN Proxy mode.

For detailed information about this command, see the *JUNOS Software CLI Reference*.

Displaying a Summary of Sessions

You can use the following show security flow command to determine the kinds of sessions on your device, how many of each kind there are—for example, the number of unicast sessions and multicast sessions—the number of failed sessions, and the maximum number of sessions that the SRX Series device supports:

```
show security flow session summary
```

Displaying Session and Flow Information About Sessions

You can use the following show security flow session command to display information about all sessions on your SRX Series device, including the session ID, the virtual system the session belongs to, the NAT source pool (if source NAT is used), the configured timeout value for the session and its standard timeout, and the session start time and how long the session has been active. The display also shows all standard flow information, including the direction of the flow, the source address and port, the destination address and port, the IP protocol, and the interface used for the session:

```
show security flow session
```

Displaying Session and Flow Information About a Specific Session

When you know the session identifier, you can use the following command to display all session and flow information for a specific session rather than for all sessions.

```
show security flow session session-identifier 40000381
```

Using Filters to Display Session and Flow Information

You can display flow and session information about one or more sessions by specifying a filter as an argument to the show security flow session command. You can use the following filters: source-prefix, destination-prefix, source-port, destination-port, protocol, interface-name, resource-manager, tunnel, and application. The SRX Series device displays the information for each session followed by a line specifying the number of sessions reported on. Here is an example of the command using the source-prefix filter:

```
show security flow session source-prefix 10/8
```

Information Provided in Session Log Entries

Session log entries are tied to policy configuration. Each main session event—create, close, and deny—will create a log entry if the controlling policy has enabled logging. See “Monitoring Policy Statistics” on page 94 for an explanation of how to enable logging in a policy configuration.

Different fields are logged for session create, session close, and session deny events as shown in Table 5 on page 30, Table 6 on page 31, and Table 7 on page 32. The same field name under each type indicates that the same information is logged, but each table is a full list of all data recorded for that type of session log.

Information displayed in session log entries includes the following:

Table 5: Session Create Log Fields

Field	Description
source-address	Source IP address of the packet that created the session.
source-port	Source port of the packet that created the session.
destination-address	Destination IP address of the packet that created the session.
destination-port	Destination port of the packet that created the session.
service-name	Application the packet traversed. (For example, “junos-telnet” for Telnet traffic during the session allowed by a policy that permits native Telnet.)
nat-source-address	The translated NAT source address if NAT was applied; otherwise, the source address as above.
nat-source-port	The translated NAT source port if NAT was applied; otherwise, the source port as above.
nat-destination-address	The translated NAT destination address if NAT was applied; otherwise, the destination address as above.
nat-destination-port	The translated NAT destination port if NAT was applied; otherwise, the destination port as above.
src-nat-rule-name	The source NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if source address translation takes place, then this field shows the static NAT rule name.*
dst-nat-rule-name	The destination NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if destination address translation takes place, then this field shows the static NAT rule name.*
protocol-id	The protocol ID of the packet that created the session.
policy-name	The name of the policy that permitted the session creation.
session-id-32	The 32–bit session ID.

** Note that some sessions may have both destination and source NAT applied and the information logged.*

Table 6: Session Close Log Fields

Field	Description
reason	The reason the session was closed.
source-address	Source IP address of the packet that created the session.
source-port	Source port of the packet that created the session.
destination-address	Destination IP address of the packet that created the session.
destination-port	Destination port of the packet that created the session.
service-name	Application the packet traversed. (For example, "junos-telnet" for Telnet traffic during the session allowed by a policy that permits native Telnet.)
nat-source-address	The translated NAT source address if NAT was applied; otherwise, the source address as above.
nat-source-port	The translated NAT source port if NAT was applied; otherwise, the source port as above.
nat-destination-address	The translated NAT destination address if NAT was applied; otherwise, the destination address as above.
nat-destination-port	The translated NAT destination port if NAT was applied; otherwise, the destination port as above.
src-nat-rule-name	The source NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if source address translation takes place, then this field shows the static NAT rule name.*
dst-nat-rule-name	The destination NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if destination address translation takes place, then this field shows the static NAT rule name.*
protocol-id	The protocol ID of the packet that created the session.
policy-name	The name of the policy that permitted the session creation.
session-id-32	The 32–bit session ID.
packets-from-client	The number of packets sent by the client related to this session.
bytes-from-client	The number of data bytes sent by the client related to this session.
packets-from-server	The number of packets sent by the server related to this session.
bytes-from-server	The number of data bytes sent by the server related to this session.
elapsed-time	The total session elapsed time from permit to close, given in seconds.

** Note that some sessions may have both destination and source NAT applied and the information logged.*

Table 7: Session Deny Log Fields

Field	Description
source-address	Source IP address of the packet that attempted to create the session.
source-port	Source port of the packet that attempted to create the session.
destination-address	Destination IP address of the packet that attempted to create the session.
destination-port	Destination port of the packet that attempted to create the session.
service-name	Application the packet attempted to traverse.
protocol-id	The protocol ID of the packet that attempted to create the session.
icmp-type	The ICMP type if the denied packet was ICMP configured; otherwise, this field will be 0.
policy-name	The name of the policy that denied the session creation.

Terminating Sessions

You can use the clear command to terminate sessions. You can clear all sessions, including sessions of a particular application type, sessions that use a specific destination port, sessions that use a specific interface or port, sessions that use a certain IP protocol, sessions that match a source prefix, and resource manager sessions.

Terminating All Sessions

You can use the following command to terminate all sessions except tunnel and resource manager sessions. The command output shows the number of sessions cleared. Be aware that this command terminates the management session through which the clear command is issued.

```
clear security flow session all
```

Terminating a Specific Session

You can use the following command to terminate the session whose session ID you specify:

```
clear security flow session session-identifier 40000381
```

Using Filters to Specify the Sessions to Be Terminated

You can terminate one or more sessions based on the filter parameter you specify for the clear command. The following example uses the protocol as a filter:

```
clear security flow session protocol 89
```

Chapter 2

Introducing JUNOS Software for J Series Services Routers

JUNOS Software on the J Series Services Router integrates and evolves the network security and routing capabilities of Juniper Networks world-class networking and security platform.

JUNOS Software uses a modular architecture as shown in Figure 6 on page 34. The operating system consists of discrete processes that implement specific features and functionality. These processes are compartmentalized with their associated data to provide functional independence, information privacy, and improved scalability.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Stateful and Stateless Data Processing on page 33
- Following the Data Path on page 39
- Understanding Secure and Router Contexts on page 42

Stateful and Stateless Data Processing

Traffic that enters and exits a Services Router running JUNOS Software is processed according to features you configure, such as security policies, packet filters, and screens. For example the software can determine:

- Whether the packet is allowed into the router
- Which class of service (CoS) to apply to the packet, if any
- Which firewall screens to apply to the packet
- Whether to send the packet through an IPsec tunnel
- Whether the packet requires an Application Layer Gateway (ALG)
- Whether to apply Network Address Translation (NAT) to translate the packet's address
- The route the packet takes to reach its destination.

Packets that enter and exit a Services Router running JUNOS Software undergo both packet-based and flow-based processing.

- Flow-based, or stateful, packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that were established for the first packet of the packet stream, which is referred to as a flow.
- Packet-based, or stateless, packet processing treats packets discretely. Each packet is assessed individually for treatment.

The software implements flow-based security and services, with packet-based application of filters, policers, traffic shapers, and other classification features.

Flow-Based Processing

A packet undergoes flow-based processing after any packet-based filters and policers have been applied to it.

Figure 6 on page 34 shows an architectural overview of traffic flow in a Services Router running JUNOS Software. See Figure 8 on page 40 to follow the path of the traffic as it traverses through the Flow services module.

Figure 6: Traffic Flow for Flow-Based Processing

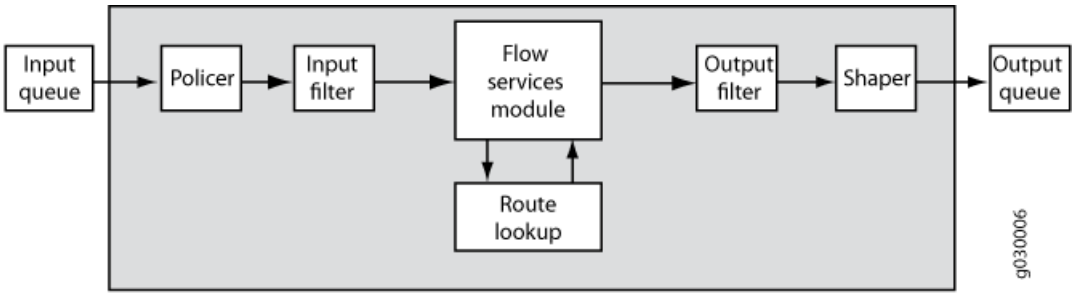

A flow is a stream of related packets that meet the same matching criteria and share the same characteristics. JUNOS Softwaretreats packets belonging to the same flow in the same manner.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, whether the packet is sent through an IPsec tunnel, if it requires an Application Layer Gateway (ALG), if Network Address Translation (NAT) is applied to translate the packet's address—are assessed for the first packet of a flow. The settings are then applied to the rest of the packets in the flow.

To determine if a packet belongs to an existing flow, the router attempts to match the packet's information to that of an existing flow based on the following six match criteria:

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

- Protocol
- Session token

If the packet matches an existing flow, processing for the packet is assessed in the context of its flow state, which is maintained by the flow's session. If it does not match the session for an existing flow, the packet is used to create a flow state and a session is allocated for it.

Zones and Policies

The security policy to be used for the first packet of a flow is cached in a flow table for use with the same flow and closely related flows. Security policies are associated with zones. A zone is a collection of interfaces that define a security boundary. A packet's incoming zone, as determined by the interface through which it arrived, and its outgoing zone, as determined by the forwarding lookup, together determine which policy is used for packets of the flow.

Flows and Sessions

Flow-based packet processing, which is stateful, requires the creation of sessions. A session is created, based on the characteristics assessed for the first packet of a flow, for the following purposes:

- To store the security measures to be applied to the packets of the flow
- To cache information about the state of the flow

  For example, logging and counting information for a flow is cached in its session. (Some stateful firewall screens rely on threshold values that pertain to individual sessions or across all sessions.)
- To allocate required resources for the flow for features such as Network Address Translation (NAT) and IPsec tunnels
- To provide a framework for features such as Application Layer Gateways (ALGs) and firewall features

Most packet processing occurs in the context of a flow. The flow engine and session bring together the following features and events that affect a packet as it undergoes flow-based processing:

- Flow-based forwarding
- Session management, including session aging and changes in routes, policy, and interfaces
- Management of virtual private networks (VPNs), ALGs, and authentication
- Management of policies, NAT, zones, and screens

Policies can be configured to log session permit, close, and deny events. To learn more about enabling logging in a policy, see “Monitoring Policy Statistics” on page 94. For information about the contents of session logs, see “Information Provided in Session Log Entries” on page 30.

Packet-Based Processing

A packet undergoes packet-based processing when it is dequeued from its input (ingress) interface and before it is enqueued on its output (egress) interface.

Packet-based processing applies stateless firewall filters and class-of-service (CoS) features to discrete packets. You can apply a firewall filter to an ingress or egress interface, or to both.

- When a packet arrives at an interface on the router, any packet-based filters and policers associated with the interface are applied to the packet before any security policies are evaluated.
- Before a packet leaves the router, any packet-based filters and traffic shapers associated with the output interface are applied to the packet after any security policies have been evaluated.

Figure 7 on page 36 shows architectural overview of traffic flow in a standard JUNOS router.

Figure 7: Traffic Flow for Packet-Based Processing

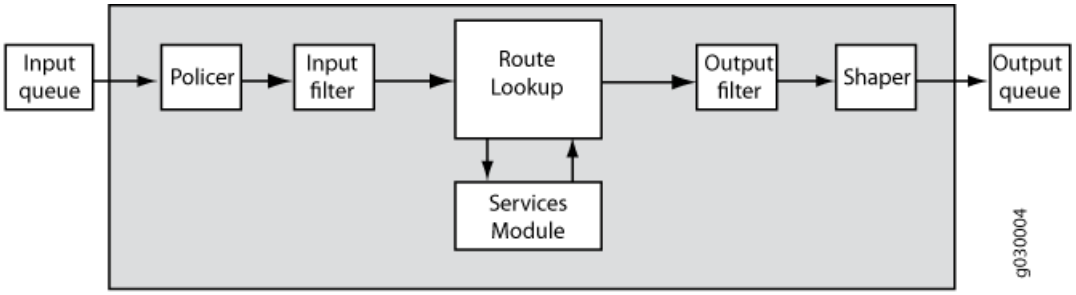

Filters and CoS features are typically associated with one or more interfaces to influence which packets are allowed to transit the system and to apply special actions to packets as necessary.

NOTE: Packet-based processing occurs only if you configure filters, CoS, IPv6, and MPLS features for an interface that handles the packet.

Here are the kinds of packet-based features that you can configure and apply to transit traffic. For details on specific stateless firewall filters and CoS features, see the *JUNOS Software Interfaces and Routing Configuration Guide* and the *JUNOS Software CLI Reference*.

- **Stateless firewall filters**—Also referred to as access control lists (ACLs), stateless firewall filters control access and limit traffic rates. They statically evaluate the contents of packets transiting the router from a source to a destination, or packets originating from or destined for the Routing Engine. A stateless firewall filter evaluates every packet, including fragmented packets.

You can apply a stateless firewall filter to an input or output interface, or to both. A filter contains one or more terms, and each terms consists of two components—match conditions and actions. By default, a packet that does not match a firewall filter is discarded.

You can plan and design stateless firewall filters to be used for various purposes—for example, to limit traffic to certain protocols, IP source or destination addresses, or data rates.

- **Class-of-service (CoS) features**—CoS features allow you to police and shape traffic.
  - **Policing traffic**—Policers allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded or assigned to a different forwarding class, a different loss priority, or both. You can use policers to limit the amount of traffic passing into or out of an interface.
  - **Traffic shaping**—You can shape traffic by assigning service levels with different delay, jitter, and packet loss characteristics to particular applications served by specific traffic flows. Traffic shaping is especially useful for real-time applications, such as voice and video transmission.

Changing Session Characteristics

Sessions are created, based on routing and other classification information, to store information and allocate resources for a flow. Sessions have characteristics, some of which you can change, such as when they terminated. For example, you might want to ensure that a session table is never entirely full to protect against an attacker's attempt to flood the table and thereby prevent legitimate users from starting sessions.

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 30 minutes. The default timeout for UDP is 1 minute. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 10 seconds.

If no traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse. You can affect the life of a session in the following ways:

- You can specify circumstances to terminate sessions using any of the following methods:
  - Aggressively age out invalid sessions aggressively based on a timeout value.
  - Age out sessions based on how full the session table is.
  - Set an explicit timeout for aging out TCP sessions.
  - Configure a TCP session to be invalidated when it receives a TCP RST (reset) message.
- You can configure sessions to accommodate other systems as follows:
  - Disable TCP packet security checks.
  - Accommodate end-to-end communication.

The following sections show you how to modify a session's characteristics. For details, see the *JUNOS Software CLI Reference*.

Controlling Session Termination

JUNOS Software terminates sessions normally in certain situations—for example, after receiving a TCP FINish Close or receiving a RST (reset) message, when encountering Internet Control Message Protocol (ICMP) errors for UDP, and when no matching traffic is received before the service timeout. When sessions are terminated, their resources are freed up for use for other sessions.

To control when sessions are terminated, you configure the router to age out sessions after a certain period of time, when the number of sessions in the session table reaches a specified percentage, or both.

- To terminate sessions based on a timeout value or the number of sessions in the session table:
  - You can use the following set security flow command to specify the number of seconds in tens of seconds after which a session is invalidated. The following command ages out sessions after 20 seconds:

    set security flow aging early-ageout 2

  - You can use the following set security flow command to specify a percentage of sessions. When the number of sessions in the session table reaches this percentage, the router begins to age sessions aggressively. When the number of sessions in the session table reaches the low-water mark, the router stops aggressively aging sessions.

    set security flow aging high-watermark 90 low-watermark 50

  - To configure an explicit timeout value, use the following command. This set security flow command removes a TCP session from the session table after 280 seconds.

    set security flow tcp-session tcp-initial-timeout 280

  - To cause any session that receives a TCP RST message to be invalidated, use the following command:

    set security flow tcp-session rst-invalidate-session

Disabling TCP Packet Security Checks

JUNOS Software provides a mechanism to disable security checks on TCP packets to ensure interoperability with hosts and routers with faulty TCP implementations. The following set security flow command disable TCP SYN checks and TCP sequence checks on all TCP sessions.

```
set security flow tcp-session no-syn-check
set security flow tcp-session no-sequence-check
```

Accommodating End-to-End TCP Communication

End-to-end TCP communication in a customer network might not work for large packets approaching 1500 bytes because of GRE or IPsec tunneling encapsulation. You can use the set security flow command to change the maximum segment size (MSS) for TCP packets to be sent or received over GRE and IPsec tunnels.

- The following set security flow commands set the TCP MSS to 1400 bytes for IPsec tunnel sessions and 1364 bytes for GRE tunnel sessions:

  ```
  set security flow tcp-mss ipsec-vpn mss 1400
  set security flow tcp-mss gre-in mss 1364
  set security flow tcp-mss gre-out mss 1364
  ```

  The following command configures the TCP MSS to 1400 bytes for all TCP sessions.

  ```
  set security flow tcp-mss all-tcp 1400
  ```

- The following command allows an unmatched incoming DNS reply packet:

  ```
  set security flow allow-dns-reply
  ```

- The following command sets the timeout value for route change to nonexistent route:

  ```
  set security flow route-change-timeout
  ```

- The following command enables TCP SYN flood protection mode:

  ```
  set security flow syn-flood-protection-mode
  ```

You can use other set security flow commands to accommodate other systems. For syntax information, see the *JUNOS Software CLI Reference*.

Following the Data Path

As a packet transits the router, it takes the following path. This packet " walk" brings together the packet-based processing and flow-based processing that JUNOS Software performs on the packet.

- Part 1—Forwarding Processing on page 40
- Part 2—Session-Based Processing on page 40
- Part 3—Forwarding Features on page 42

Figure 8 on page 40 shows the path of a data packet as it traverses through the Services Router. Refer to Figure 6 on page 34 to see how the flow module in Figure 8 on page 40 fits in with the architecture of the software.

Figure 8: Data Packet Traversing the Flow Module on the Services Router

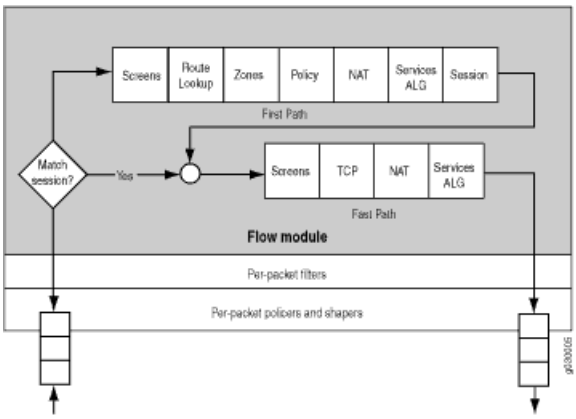

Part 1—Forwarding Processing

1. The packet enters the system and is treated on a per-packet basis.
2. The system applies stateless policing filters and class-of-service (CoS) classification to the packet.

   For details, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Part 2—Session-Based Processing

After forwarding processing, JUNOS Software performs session lookup and either first-packet processing or fast-path processing on the packet.

Session Lookup

If the packet has not already been dropped, JUNOS Software performs session lookup to determine whether the packet belongs to an existing session. The system uses six match criteria to perform the session lookup:

- Session token
- Source and destination IP addresses
- Source and destination ports
- Protocol

If the packet does not match an existing session, the system creates a new session for it. This process is called the first-packet path. (See “First-Packet Path Processing” on page 41.)

If the packet matches a session, fast-path processing is performed. (See “Fast-Path Processing” on page 42.)

First-Packet Path Processing

If a packet does not match an existing session, JUNOS Software creates a new session for it as follows:

1. For the first packet, the system creates a session based on the routing for the packet and the policy lookup so that the packet becomes the first packet of a flow.

   For policy details, see "Security Policies Overview" on page 69.

2. Depending on the protocol and whether the service is TCP or UDP, the session is programmed with a timeout value.

   - For TCP, the default timeout is 30 minutes.
   - For UDP, the default timeout is 1 minute.

   You can configure these timeouts to be more aggressive or less aggressive. If you have changed the session timeout value, it is applied here. See "Controlling Session Termination" on page 38, If no traffic uses the session during the service timeout period, the router ages out the session and releases its memory for reuse.

3. Firewall screens are applied.

   Session initialization screens are applied. For screen details, see "Attack Detection and Prevention" on page 761.

4. Route lookup is performed.

5. The destination zone is determined:

   a. The system determines a packet's *incoming* zone by the interface through which it arrives.

   b. The system determines a packet's *outgoing* zone by route lookup.

   Together they determine which policy is applied to the packet.

   For zone details, see "Security Zones and Interfaces" on page 47.

6. Policy lookup is performed.

   The system checks the packet against policies you have defined to determine how the packet is to be treated.

   For policy details, see "Security Policies" on page 69.

7. If NAT is used, the system performs address allocation.

   For NAT details, see "Network Address Translation" on page 963.

8. The system sets up the Application Layer Gateway (ALG) service vector.

   For ALG details, see the "Application Layer Gateways (ALGs)" on page 157.

9. The system creates and installs the session.

Decisions made for the first packet of a flow are cached in a flow table for use with following, related flows.

10. Fast path processing is applied to the packet.

Fast-Path Processing

If a packet matches a session, JUNOS Software performs fast-path processing as follows:

1. Configured screens are applied.
2. TCP checks are performed.
3. NAT is applied.

   For NAT details, see "Network Address Translation" on page 963.
4. Forwarding features are applied. (See the following section "Part 3—Forwarding Features" on page 42.)

Part 3—Forwarding Features

After the packet has passed through session-based processing, JUNOS Software prepares the packet and transmits it:

1. Routing packet filters are applied.
2. Traffic shaping is applied.
3. The packet is transmitted.

For information about packet filters and CoS traffic shaping, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Understanding Secure and Router Contexts

As shipped from the factory, a Services Router running JUNOS Software initially starts up and uses a configuration that places the router in secure context. You can change the context in which the Services Router is running from secure context to router context. To do so, you use a predefined template configuration file. If you plan to use the Services Router primarily as a router, change to router context with this configuration as your starting point.

CAUTION: If you plan to change contexts, do so before you configure anything else on the Services Router. If you change contexts after you have configured the Services Router, your configuration is overwritten by the default configuration for the new context.

NOTE: Both secure context and router context are flow-enabled. MPLS is not supported in either context because MPLS undergoes packet-based processing.

Secure and Router Context Support On Different Device Types

The following table lists secure and router context features, specifies whether the features are supported on various device types, and indicates where you can find more information about each feature.

Table 8: Support Information: Secure and Router Context Options

Feature	J Series Services Routers	SRX Series Services Gateways	More Information
Secure context factory configuration	Yes	No	"Secure Context" on page 43
Router context factory configuration	Yes	No	"Router Context" on page 44

Secure Context

Secure context allows a Services Router to act as a stateful firewall with only management access. To allow traffic to pass through a Services Router, you must explicitly configure a security policy for that purpose. In secure context, a Services Router forwards packets only if a security policy permits it. Certain services are also configured (in the host-inbound-traffic statement of the [edit security zones] hierarchy level) to allow host-inbound traffic for management of a Services Router. A Services Router running in secure context is a secure routing device with predefined configuration values.

When you use the router in secure context, you can configure additional security features. You can also remove security features and configure additional routing features to provide greater routing capability. The secure context configuration of the router is provided for ease of use. It is intended as a starting point that you can build on to customize the router for your environment.

To change contexts, see the *JUNOS Software Administration Guide*.

The basic configuration for secure context includes the following settings:

- A predefined interface called ge-0/0/0, which is bound to a preconfigured zone called trust. All other interfaces are bound to a preconfigured untrust zone.

  The ge-0/0/0 interface is configured to allow management access with SSH and HTTP services enabled. The following host-inbound services are configured for the ge-0/0/0 interface in the trust zone: HTTP, HTTPS, SSH, Telnet, and DHCP.

- For the trust zone, TCP reset is enabled. The default policy for the trust zone allows transmission of traffic from the trust zone to the untrust zone. All traffic within the trust zone is allowed.

- A screen is applied to a zone to protect against attacks launched from within the zone. The following screens are enabled for the untrust zone: ICMP ping-of-death, IP source route options, IP teardrop, TCP land attack, TCP SYN flood (with the alarm threshold set to 1024, attack threshold set to 200, source threshold set to 1024, destination threshold set to 2048, and a timeout value of 20 seconds).
- The default policy for the untrust zone is to deny all traffic.

For the default configuration file for secure context, see the *JUNOS Software Administration Guide*.

Router Context

Router context allows a Services Router to act as a router in which all management and transit traffic is allowed. All interfaces are bound to the trust zone, and host inbound traffic from all predefined services is allowed. In router context, the Services Router forwards all packets unless you configure a security policy that denies specific traffic.

JUNOS Software is a hardened operating system. You can use JUNOS Software with more relaxed checks for host-inbound traffic and configure the data plane with default transit policies to permit all traffic. In this scenario, the Services Router operates in a router context.

You load a predefined template configuration, jsr-series-routermode-factory.conf, to change to router context. In router context, the Services Router remains flow-enabled. All security features are available, but they are explicitly disabled.

When you use the router in router context, you can configure additional routing features. You can also configure security features selectively to provide additional protection. The router context configuration is provided for ease of use. It is intended as a starting point that you can build on to customize the router for your environment.

For the default configuration file for router context and to change contexts, see the *JUNOS Software Administration Guide*.

Part 2
Security Zones and Interfaces

■ Security Zones and Interfaces on page 47

Chapter 3

Security Zones and Interfaces

Interfaces act as a doorway through which traffic enters and exits a J Series Services Router or SRX Series Services Gateway. Many interfaces can share exactly the same security requirements; however, different interfaces can also have different security requirements for inbound and outbound data packets. Interfaces with identical security requirements can be grouped together into a single *security zone*.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding Security Zones on page 47
- Creating Security Zones on page 49
- Configuring Security Zones—Quick Configuration on page 51
- Configuring Host Inbound Traffic on page 53
- Configuring Protocols on page 58
- Configuring the TCP-Reset Parameter on page 60
- Understanding Security Zone Interfaces on page 61
- Understanding Interface Ports on page 62
- Configuring Interfaces—Quick Configuration on page 62
- Configuring a Gigabit Ethernet Interface—Quick Configuration on page 63

Understanding Security Zones

A *security zone* is a collection of one or more network segments requiring the regulation of inbound and outbound traffic through policies.

Before You Begin
For background information, read "Understanding Policies" on page 70.

Security zones are logical entities to which one or more interfaces are bound. With many types of Juniper Networks devices, you can define multiple security zones, the exact number of which you determine based on your network needs.

On a single device, you can configure multiple security zones, dividing the network into segments to which you can apply various security options to satisfy the needs of each segment. At a minimum, you must define two security zones, basically to protect one area of the network from the other. On some security platforms, you can define many security zones, bringing finer granularity to your network security design—and without deploying multiple security appliances to do so.

From the perspective of security policies, traffic enters into one security zone and goes out on another security zone. This combination of a from-zone and a to-zone is defined as a *context*. Each context contains an ordered list of policies. For more information on policies, see "Security Policies Overview" on page 69.

JUNOS Software supports functional zones and security zones.

This topic covers:

- Functional Zone on page 48
- Security Zone on page 48
- Related Topics on page 49

Functional Zone

A functional zone is used for special purposes, like management interfaces. Currently, only the management (MGT) zone is supported. Management zones have the following properties:

- Management zones host management interfaces.
- Traffic entering management zones does not match policies; therefore, traffic cannot transit out of any other interface if it was received in the management interface.
- Management zones can only be used for dedicated management interfaces.

Security Zone

Security zones are the building blocks for policies; they are logical entities to which one or more interfaces are bound. Security zones provide a means of distinguishing groups of hosts (user systems and other hosts, such as servers) and their resources from one another in order to apply different security measures to them.

Security zones have the following properties:

- Policies—Active security policies that enforce rules for the transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on the traffic as it passes through the firewall. For more information, see "Security Policies Overview" on page 69.
- Screens—A Juniper Networks stateful firewall secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another. For every security zone, and the MGT zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful. For more information, see "Reconnaissance Deterrence Overview" on page 763.

- Address books—IP addresses and address sets that make up an address book to identify its members so that you can apply policies to them. For more information, see “Configuring Address Books” on page 102.
- TCP-RST—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the SYNchronize flag set.
- Interfaces—List of interfaces in the zone.

Security zones have the following preconfigured zones:

- **junos-global zone**—Defined in the JUNOS defaults and cannot be configured by the user. The global zone serves as a storage area for static NAT addresses and can be used in policies like any other security zone.
- **Trust zone**—Available only in the factory configuration and is used for initial connection to the device. After you commit a configuration, the trust zone can be overridden.

Related Topics

- Creating Security Zones on page 49
- Configuring Security Zones—Quick Configuration on page 51
- Configuring the TCP-Reset Parameter on page 60

Creating Security Zones

For best performance, always save your changes and reboot after creating a zone.

Before You Begin

For background information, read

- Understanding Security Zones on page 47
- Configuring Host Inbound Traffic on page 53
- Configuring the TCP-Reset Parameter on page 60

When you configure a security zone, you can specify many of its parameters at the same time. For purposes of illustration, this section only shows how to configure zones and assign interfaces to them.

To configure a zone, you can use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 50
- CLI Configuration on page 50
- Related Topics on page 50

J-Web Configuration

To configure a zone using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type **ABC** and click **OK**.

To configure an interface and assign it to the created security zone:

1. Corresponding to the security zone name ABC zone, click **Edit**.
2. Next to Interfaces, click **Add new entry**.
3. In the Interface unit box, type ge-0/0/1 and click **OK**.
4. Next to Unit, click **Add new entry**.
5. In the Interface unit number box, type 1.
6. Under Family, select **Inet** and click **OK**.
7. Next to Address book, click **Configure** or **Edit**.
8. Next to Address, click **Add new entry**.
9. In the Address name box, type **10.12.12.1/24** and click **OK.**
10. If you are finished configuring the device, commit the configuration.

CLI Configuration

To configure the interface and its IP address for the ABC zone, enter the following statement in Edit mode:

```
user@host# set interfaces ge-0/0/1 unit 1 family inet address 10.12.12.1/24
```

To configure the ABC zone and assign the interface to it, enter the following statement in Edit mode:

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.1
```

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

Related Topics

- Understanding Security Zones on page 47
- Configuring Security Zones—Quick Configuration on page 51

Configuring Security Zones—Quick Configuration

You can use J-Web Quick Configuration to quickly configure security zones. See Figure 9 on page 52.

Before You Begin

For background information, read:

- Understanding Security Zones on page 47
- Creating Security Zones on page 49

Figure 9 on page 52 shows the Quick Configuration Zones page.

Figure 9: Quick Configuration Zones Page

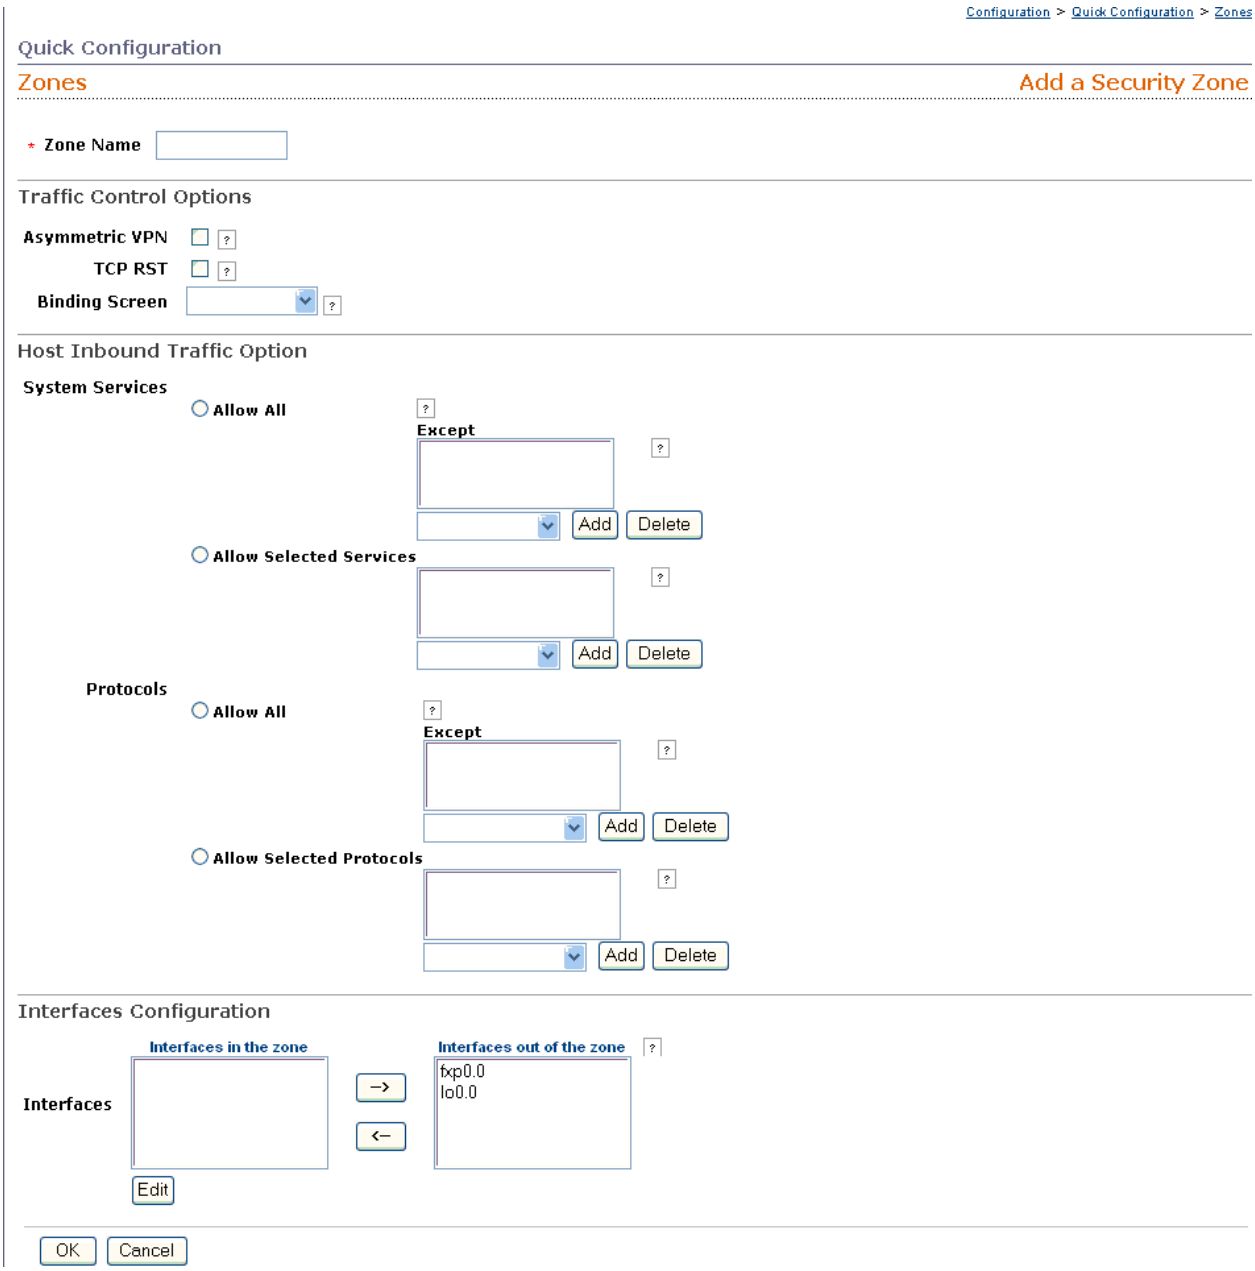

To configure security zones with Quick Configuration:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Click **Add** to create new zones.

5. Fill the form as shown in Table 9 on page 53.
6. Click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.

Table 9: Security Zone Options

Zone Name	Name of the zone for which you are enabling policies
Traffic Control Options	Asymmetric VPN—Allows any incoming VPN traffic in a zone to match any applicable VPN session, regardless of the origin for the original VPN tunnel. This feature allows free routing of VPN traffic between two or more sites when there are multiple possible paths for VPN traffic. TCP RST—Select this check box to enable the tcp-rst feature, which sends a TCP segment with the RESET flag set to 1 in response to a TCP segment with any flag set other than SYN and which does not belong to an existing session
Host Inbound Traffic Option	System Services—Configure services to permit inbound traffic of the selected type to be transmitted to hosts within the zone, provided there is a policy that permits it. You can select Allow All to permit all services, or you can select Except and Allow Selected Services to exclude selected services. Protocols—Configure protocols to permit inbound traffic of the selected type to be transmitted to hosts with the zone. You can select Allow All to permit all protocols, or use the Except and Allow Selected Protocols options to exclude selected protocols.
Interfaces Configuration	Select the interfaces that you want included in the security zone.

Configuring Host Inbound Traffic

This topic describes how to configure zones to specify the kinds of traffic that can reach the device from systems that are directly connected to its interfaces.

- You can configure these parameters at the zone level, in which case they affect all interfaces of the zone, or at the interface level. (Interface configuration overrides that of the zone.)
- You must enable all expected host-inbound traffic. Inbound traffic from devices directly connected to the device's interfaces is dropped by default.

- You can also configure a zone's interfaces to allow for use by dynamic routing protocols.

Before You Begin
For background information, read “Understanding Security Zones” on page 47.

This feature allows you to protect the device against attacks launched from systems that are directly connected to any of its interfaces. It also enables you to selectively configure the device so that administrators can manage it using certain applications on certain interfaces. You can prohibit use of other applications on the same or different interfaces of a zone. For example, most likely you would want to ensure that outsiders not use the Telnet application from the Internet to log into the device because you would not want them connecting to your system.

This topic covers:

- System Services on page 54
- J-Web Configuration on page 55
- CLI Configuration on page 56
- Related Topics on page 58

System Services

Any host-inbound traffic that corresponds to a service listed under this option is allowed. For example, suppose a user whose system was directly connected to interface 1.3.1.4 in zone ABC wanted to telnet into interface 2.1.2.4 in zone ABC. For this action to be allowed, the telnet application must be configured as an allowed inbound service on both interfaces and a policy must permit the traffic transmission.

Table 10 on page 54 lists the supported services. A value of all indicates that traffic from all of the following services is allowed inbound on the specified interfaces (of the zone, or a single specified interface)

Table 10: Supported System Services

Supported System Services			
all	http	rpm	traceroute
bootp	https	rsh	xnm-clear-text
dhcp	ike	snmp	xnm-ssl
finger	netconf	snmp-trap	
ftp	ping	ssh	
ident-reset	rlogin	telnet	

NOTE: All services listed Table 10 on page 54 (except DHCP and BOOTP) can be configured either per zone or per interface. A DHCP server is configured only per interface because the incoming interface must be known by the server to be able to send out DHCP replies.

To configure zones to allow use of supported application services as inbound services, use either J-Web or the CLI configuration editor.

J-Web Configuration

To configure the ABC zone to allow use of all the supported application services as inbound services using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type ABC and click **OK**.
6. Next to Host inbound traffic, click **Configure** or **Edit**.
7. To allow the security zone to use all of the supported application services, next to System services, click **Add new entry**.
8. From the Service name list, select **All** and click **OK**.

To enable FTP and telnet for interfaces such as ge-0/0/1.3 and ge-0/0/1:

1. Next to Interfaces, click **Add new entry**.
2. In the Interface unit box, type ge-0/0/1.3 and click **OK**.
3. Next to Host inbound traffic, click **Configure** or **Edit**.
4. Next to System services, click **Add new entry**.
5. From the Service name list, select **ftp** and click OK.
6. Next to Interfaces, click **Add new entry**.
7. In the Interface unit box, type ge-0/0/1.1 and click **OK**.
8. Next to Host inbound traffic, click **Configure** or **Edit**.
9. Next to System services, click **Add new entry**.
10. From the Service name list, select **telnet** and click **OK**.

To enable FTP and telnet for interface ge-0/0/1.3 and only SNMP for interface ge-0.0/1.1:

1. Next to Interfaces, click **Add new entry**.
2. In the Interface unit box, type ge-0/0/1.3 and click **OK**.
3. Next to Host inbound traffic, click **Configure** or **Edit**.

4. Next to System services, click **Add new entry**.
5. From the Service name list, select **ftp** and **telnet**, and click **OK**.
6. To enable SNMP for interface ge-0.0/1.1:
7. Next to Interfaces, click **Add new entry**.
8. In the Interface unit box, type ge-0/0/1.1 and click **OK**.
9. Next to Host inbound traffic, click **Configure** or **Edit**.
10. Next to System services, click **Add new entry**.
11. From the Service name list, select **snmp** and click **OK**.

To allow all configurable system services on the interface ge-0/0/1.3, except Telnet:

1. Next to Interfaces, click **Add new entry**.
2. In the Interface unit box, type ge-0/0/1.3 and click **OK**.
3. Next to Host inbound traffic, click **Configure** or **Edit**.
4. Next to System services, click **Add new entry**.
5. From the Service name list, select **al** and click **OK**.
6. Next to System services, click **Add new entry**.
7. From the Service name list, select **telnet** and click **OK**.
8. Select the **Except** check box and click **OK**.
9. If you are finished configuring the device, commit the configuration.

CLI Configuration

To configure the ABC zone to allow use of all of the supported application services as inbound services, enter the following statements in Configure mode:

```
user@host# set security zones security-zone ABC host-inbound-traffic system-services
    all
```

In the following example, FTP and telnet are enabled for interfaces ge-0/0/1.3 and ge-0/0/1. You must configure FTP and telnet at the interface level, not the zone level. For incoming FTP and telnet requests to be recognized, the interface must be known to the server.

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
    host-inbound-traffic system-services ftp

user@host# set security zones security-zone ABC interfaces ge-0/0/1.1
    host-inbound-traffic system-services telnet
```

In the following example, FTP and telnet are enabled for interface ge-0/0/1.3 and only SNMP is enabled for interface ge-0.0/1.1.

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
    host-inbound-traffic system-services ftp
```

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
  host-inbound-traffic system-services telnet

user@host# set security zones security-zone ABC interfaces ge-0/0/1.1
  host-inbound-traffic system-services snmp
```

You can use the **all** option to allow all configurable system services and use the **except** option to exclude certain services. In this example, all configurable system services are permitted on interface **ge-0/0/1.3**, except **Telnet**.

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
  host-inbound-traffic system-services all

user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
  host-inbound-traffic system-services telnet except
```

In the following example, all configurable system services are permitted on interface **ge-0/0/1.1**, except **HTTP** and **FTP**.

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.1
  host-inbound-traffic system-services all

user@host# set security zones security-zone ABC interfaces ge-0/0/1.1
  host-inbound-traffic system-services http except

user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
  host-inbound-traffic system-services ftp except
```

In the following example, **telnet** and **FTP** are enabled for security zone **ABC**/interface **ge-0/0/1.1**, but there is an interface override that takes priority and only **SNMP** is allowed on interface **ge-0/0/1.3**.

```
user@host# set security zones security-zone ABC host-inbound-traffic system-services
  telnet

user@host# set security zones security-zone ABC host-inbound-traffic system-services
  ftp

user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
  host-inbound-traffic system-services snmp
```

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

Another view of the previous configuration:

```
security zones
 security-zone ABC {
  host-inbound-traffic {
   system-services {
    telnet;
    ftp;
   }
  }
  interfaces {
   ge-0/0/1.1;
   ge-0/0/1.3 {
    host-inbound-traffic {
```

```
    system-services {
      snmp;

   }
  }
 }
}
```

For more information on host-inbound traffic parameters, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding Security Zones on page 47
- Configuring the TCP-Reset Parameter on page 60
- Creating Security Zones on page 49

Configuring Protocols

Any host-inbound traffic that corresponds to a protocol listed under this option is allowed. For example, if anywhere in the configuration, you map a protocol to a port number other than the default, you can specify the protocol in the host-inbound traffic option, and the new port number will be used.Table 11 on page 58 lists the supported protocols. A value of all indicates that traffic from all of the following protocols is allowed inbound on the specified interfaces (of the zone, or a single specified interface).

Table 11: Supported Inbound System Protocols

Supported System Services			
all	igmp	pgm	sap
bfd	ldp	pim	vrrp
bgp	msdp	rip	nhrp
router-discovery	dvmrp	ospf	rsvp

NOTE: If DVMRP or PIM is enabled for an interface, IGMP and MLD host-inbound traffic is enabled automatically. Because ISIS uses OSI addressing and should not generate any IP traffic, there is no host-inbound traffic option for the ISIS protocol.

To use supported protocols for host-inbound traffic, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 59
- CLI Configuration on page 59
- Related Topics on page 60

J-Web Configuration

To configure the ABC zone to allow use of all the supported protocols for host inbound traffic using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type ABC and click **OK**.

To configure an interface for the created security zone, corresponding to the security zone name ABC , next to ABC, click **Edit**.

1. Next to Interfaces, click **Add new entry**.
2. In the Interface unit box, type ge-0/0/1.3 and click **OK**.
3. Next to Host inbound traffic, click **Configure** or **Edit**.
4. Next to System services, click **Add new entry**.
5. From the Service name list, select **ping** and click **OK**.
6. Next to System services, click **Add new entry**.
7. From the Service name list, select **ssh** and click **OK**.
8. Next to System services, click **Add new entry**.
9. From the Service name list, select **traceroute** and click **OK**.
10. Next to Protocols, click **Add new entry**.
11. In the Protocol name box, type ospf and click **OK**.
12. If you are finished configuring the device, commit the configuration.

CLI Configuration

In the following example, ping, ssh, traceroute, and ospf host-inbound traffic is enabled for interface ge-0.0/1.1

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
    host-inbound-traffic system-services ping

user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
    host-inbound-traffic system-services ssh
```

```
user@host# set security zones security-zone ABC interfaces ge-0/0/1.3
  host-inbound-traffic system-services traceroute

user@host# set security zones security-zone ABC interfaces ge-0/0/1.1
  host-inbound-traffic protocols ospf
```

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

Another view of the previous configuration:

```
security zones security-zone ABC {
  interfaces {
  ge-0/0/0.33 {
   host-inbound-traffic {
    system-services {
     ping;
     ssh;
     traceroute;
    }
    protocols {
     ospf;
    }
   }
  }
 }
}
```

For more information on host-inbound protocols configuration, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding Security Zones on page 47
- Creating Security Zones on page 49
- Configuring Host Inbound Traffic on page 53

Configuring the TCP-Reset Parameter

When the TCP-RST feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the SYNchronize flag set.

Before You Begin
For background information, read “Understanding Security Zones” on page 47.

To configure the TCP-Reset parameter for the zone ABC, use either J-Web configuration or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 61
- CLI Configuration on page 61
- Related Topics on page 61

J-Web Configuration

To configure the TCP-Reset parameter for the zone ABC using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type ABC.
6. To set the parameter for the ABC zone, select the **Tcp rst** check box and click **OK**.
7. If you are finished configuring the device, commit the configuration.

CLI Configuration

To set the parameter for the zone ABC, enter the following statement in configure mode:

```
user@host# set security zones security-zone ABC tcp-rst
```

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

For more information on TCP-RST configuration, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding Security Zones on page 47
- Configuring Host Inbound Traffic on page 53

Understanding Security Zone Interfaces

An interface for a security zone can be thought of as a doorway through which TCP/IP traffic can pass between that zone and any other zone.

Before You Begin
For background information, read “Understanding Security Zones” on page 47.

Through the policies you define, you can permit traffic between zones to flow in one direction or in both. With the routes that you define, you specify the interfaces that traffic from one zone to another must use. Because you can bind multiple interfaces to a zone, the routes you chart are important for directing traffic to the interfaces of your choice.

Understanding Interface Ports

On J Series Services Routers, interface ports for the system are located on Physical Interface Modules (PIMs) that you can install in slots on the device. In addition, each device has four built-in Gigabit Ethernet ports in slot 0. Each physical port can have many logical interfaces configured with properties different from the port's other logical units.

Interfaces are named by type, slot number, module number (always 0), port number, and the logical unit number. Port numbering starts with 0. Interface names have the following format:

```
type-pim/0/port.logical-unit-number
```

For example, an interface on port 1 of a T1 PIM installed in slot 3 is named t1-3/0/1. Logical unit 1 on the interface is named t1-3/0/1.1. The built-in Gigabit Ethernet interfaces are named ge-0/0/0 through ge-0/0/3.

For more information about interfaces and interface names, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Related Topics

- Creating Security Zones on page 49
- Configuring Interfaces—Quick Configuration on page 62

Configuring Interfaces—Quick Configuration

You can use J-Web Quick Configuration to quickly configure interfaces. See Figure 10 on page 63.

Before You Begin

For background information, read:

- Understanding Security Zones on page 47
- Creating Security Zones on page 49
- Understanding Security Zone Interfaces on page 61

A list of the network interfaces available on the routing platform appears, as shown in Figure 10 on page 63. The third column indicates whether the interface has been configured.

Figure 10: Quick Configuration Interfaces Page

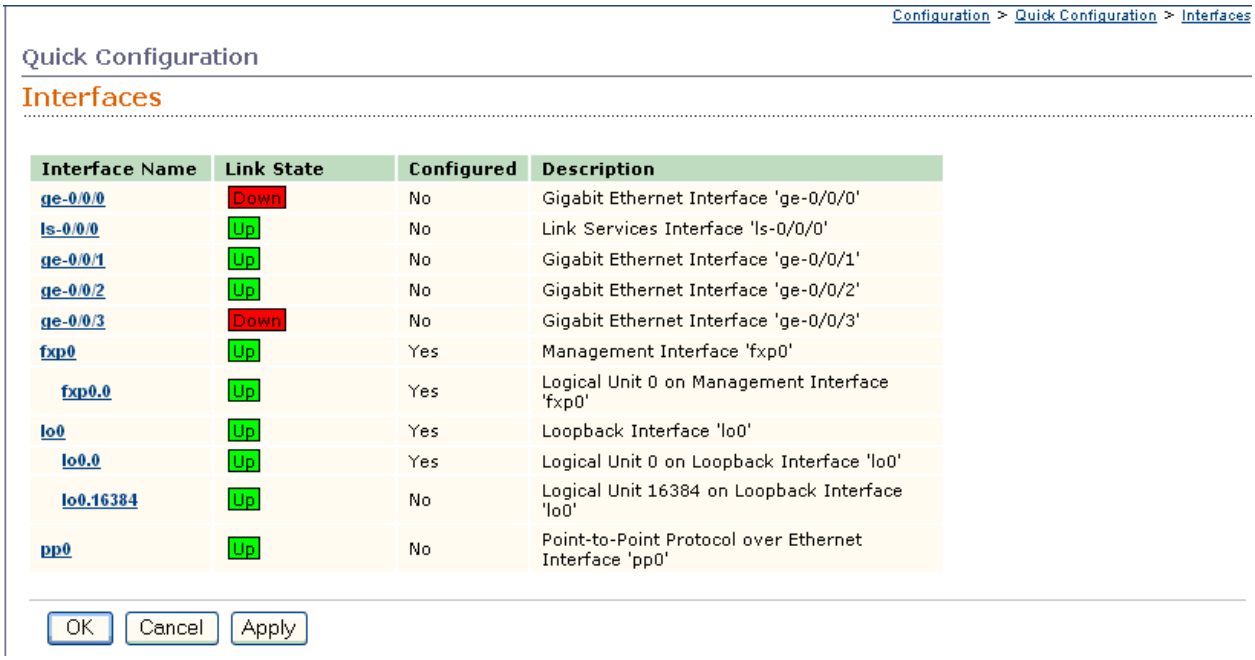

To configure a network interface with Quick Configuration:

1. Select **Configure > Interfaces**. For information about interface names, see the *JUNOS Software Interfaces and Routing Configuration Guide*.
2. Configure properties for a network interface by selecting the interface name and following the instructions in “Configuring a Gigabit Ethernet Interface—Quick Configuration” on page 63.

Configuring a Gigabit Ethernet Interface—Quick Configuration

You can use J-Web Quick Configuration to quickly configure a Gigabit Ethernet interface.

Figure 11 on page 64 shows the Gigabit Ethernet Interface Quick Configuration page.

Figure 11: Gigabit Ethernet Interface Quick Configuration

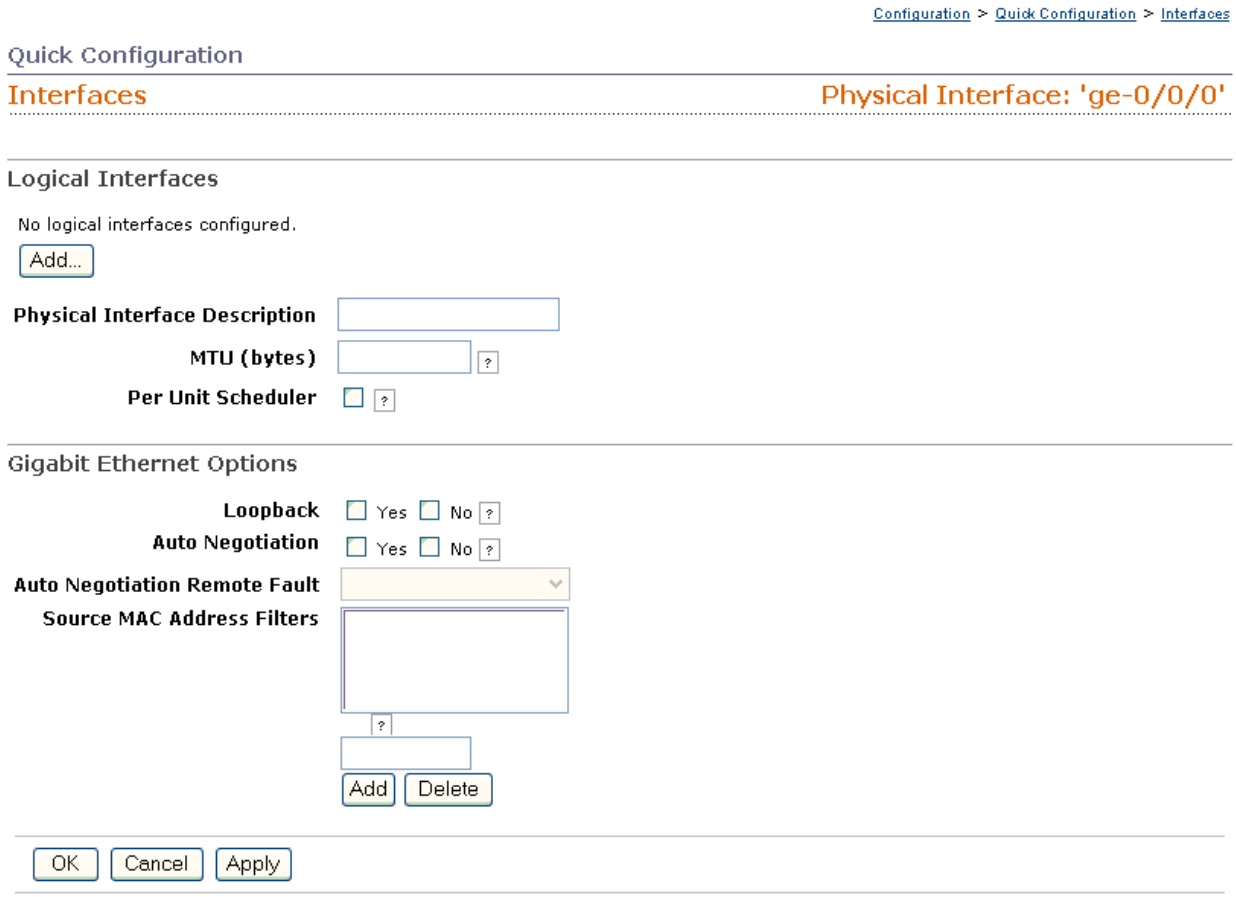

1. Select Configure>Interfaces The properties you can configure on a Gigabit Ethernet interface appear, as shown in Figure 11 on page 64.
2. Fill in the information as described in Table 12 on page 64.
3. Click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.
4. Verify that the Gigabit Ethernet interface is configured correctly, by seeing the *JUNOS Software Interfaces and Routing Configuration Guide*.

Table 12: Gigabit Ethernet Quick Configuration Page Summary

Field	Function	Actions
Logical Interfaces		

Table 12: Gigabit Ethernet Quick Configuration Page Summary *(continued)*

Field	Function	Actions
Add Logical Interfaces	Defines one or more logical units that you connect to this physical Gigabit Ethernet interface. You must define at least one logical unit for a Gigabit Ethernet interface.	Click **Add**.
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.1/24 2. Click **Add**. 3. Click **OK**. To delete an IP address and prefix, select them in the Source Addresses and Prefixes dialog box, then click **Delete**
ARP Address	Enables the device to create a static Address Resolution Protocol (ARP) entry for this interface by specifying the IP address of a node to associate with its media access control (MAC) address. The IP address must be in the same subnet as the IPv4 address or prefix of the interface you are configuring. Static ARP entries associate the IP addresses and MAC addresses of nodes on the same subnet, enabling the device to respond to ARP requests having destination addresses that are not local to the incoming interface.	Type an IPv4 address that you want to associate with the MAC address—for example, 10.10.10.1.
MAC Address	Specifies the hardware media access control (MAC) address associated with the ARP address. The MAC address uniquely identifies the system and is expressed in the following format: mm:mm:mm:ss:ss:ss. The first three octets denote the hardware manufacturer ID, and the last three are serial numbers identifying the device.	Type the MAC address to be mapped to the ARP entry—for example, 00:12:1E:A9:8A:80
Publish	Enables the device to reply to ARP requests for the specified address.	■ To enable publishing, select the check box. ■ To disable publishing, clear the check box.
Physical Interface Description	(Optional) Adds supplementary information about the physical Gigabit Ethernet interface.	Type a text description of the Gigabit Ethernet interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the Gigabit Ethernet interface.	Type a value between 256 and 9014 bytes. The default MTU for Gigabit Ethernet interfaces is 1514.

Table 12: Gigabit Ethernet Quick Configuration Page Summary *(continued)*

Field	Function	Actions
Per Unit Scheduler	Enables scheduling on logical interfaces. Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.	■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.
Gigabit Ethernet Options		
Loopback	Enables or disables the loopback option.	Select Yes to enable the loopback diagnostic option, or select No to disable the loopback option. By default, loopback is disabled.
Auto Negotiation	Enables or disables auto negotiation. By default, Gigabit Ethernet interfaces auto negotiate the link mode and speed settings. If you disable auto negotiation and do not manually configure link mode and speed, the link is negotiated at 1000 Mbps, full duplex. When you configure both the link mode and the speed, the link negotiates with the manually configured settings whether auto negotiation is enabled or disabled.	Select Yes to enable auto negotiation, or select No to disable it. By default, auto-negotiation is enabled.
Auto Negotiation Remote Fault	Indicates the auto negotiation remote fault value.	Select the auto-negotiation remote fault value from the list of options given. This field is enabled only if auto negotiation is enabled.
Source MAC Address Filters	Displays the list of media access control (MAC) addresses from which you want to receive packets on this interface.	To add MAC addresses, type them in the boxes above the Add button, then click **Add**.

Part 3
Security Policies

- Security Policies on page 69
- Security Policy Address Books and Address Sets on page 97
- Security Policy Schedulers on page 105
- Security Policy Applications on page 115

Chapter 4

Security Policies

With the advent of the Internet, the need for a secure network has become vital for businesses with an Internet connection. Before a network can be secured for a business, a network security policy has to outline all the network resources within that business and identify the required security level for those resources. The network security policy also defines the security threats and the actions taken for such threats. JUNOS Software stateful firewall policy provides a set of tools to network administrators, enabling them to implement network security for their organizations.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Security Policies Overview on page 69
- Understanding Policies on page 70
- Understanding Policy Ordering on page 74
- Configuring Policies—Quick Configuration on page 76
- Configuring Policies on page 80
- Verifying Policy Configuration on page 85
- Example: Configuring Security Policies—Detailed Configuration on page 86
- Configuring a Policy to Permit Traffic on page 86
- Configuring a Policy to Deny Traffic on page 90
- Reordering Policies After They Have Been Created on page 92
- Troubleshooting Policy Configuration on page 92
- Monitoring Policy Statistics on page 94

Security Policies Overview

In a JUNOS Software stateful firewall, the security policies enforce rules for the transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on the traffic as it passes through the firewall. From the perspective of security policies, the traffic enters one security zone and exits another security zone. This combination of a *from-zone* and *to-zone* is called a *context*. Each context contains an *ordered list* of policies.

A security policy, which can be configured from the user interface, controls the traffic flow from one zone to another zone by defining the kind(s) of traffic permitted from specified IP sources to specified IP destinations at scheduled times.

Policies allow you to deny, permit, reject (deny and send a TCP RST or ICMP port unreachable message to the source host), encrypt and decrypt, authenticate, prioritize, schedule, filter, and monitor the traffic attempting to cross from one security zone to another. You decide which users and what data can enter and exit, and when and where they can go.

NOTE: For a J Series or SRX Series device that supports virtual systems, policies set in the root system do not affect policies set in virtual systems.

Understanding Policies

A J Series or SRX Series device secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another.

By default, a device denies all traffic in all directions. Through the creation of policies, you can control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from specified sources to specified destinations at scheduled times. At the broadest level, you can allow all kinds of traffic from any source in one zone to any destination in all other zones without any scheduling restrictions. At the narrowest level, you can create a policy that allows only one kind of traffic between a specified host in one zone and another specified host in another zone during a scheduled interval of time. See Figure 12 on page 71.

Figure 12: Default Policy

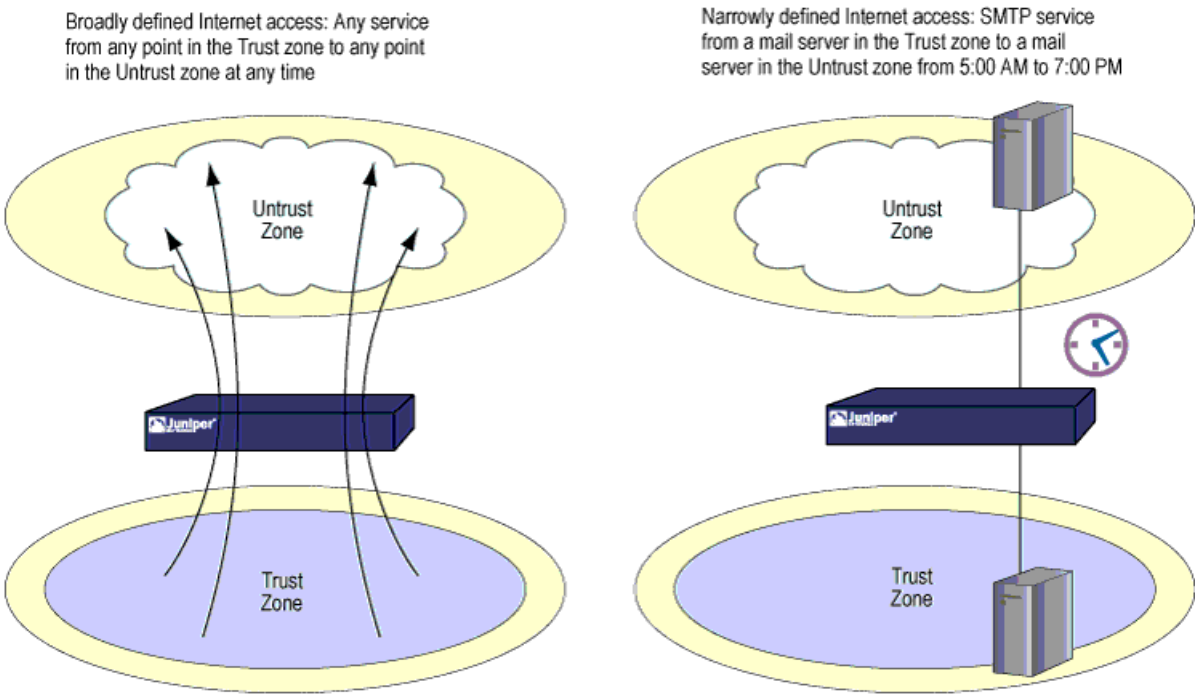

Every time a packet attempts to pass from one zone to another or between two interfaces bound to the same zone, the device checks for a policy that permits such traffic (see “Policy Application Sets Overview” on page 116). To allow traffic to pass from one security zone to another—for example, from zone A to zone B—you must configure a policy that permits zone A to send traffic to zone B. To allow traffic to flow the other way, you must configure another policy permitting traffic from zone B to zone A.

To allow data traffic to pass between zones, you must configure firewall policies.

This topic covers:

- Understanding Policy Rules on page 71
- Understanding Policy Elements on page 73
- Understanding Policy Configuration on page 74
- Related Topics on page 74

Understanding Policy Rules

The security policy applies the security rules to the transit traffic within a context (from-zone to to-zone). Each policy is uniquely identified by its name. The traffic is classified by matching its source and destination zones, the source and destination

addresses, and the application that the traffic carries in its protocol headers with the policy database in the data plane.

Before You Begin

For background information, read "Security Policies Overview" on page 69.

Each policy is associated with the following characteristics:

- A source zone
- A destination zone
- One or many source address names or address set names
- One or many destination address names or address set names
- One or many application names or application set names

These characteristics are called the *match criteria*. Each policy also has actions associated with it: permit, deny, and reject. You have to specify the match condition arguments when you configure a policy, source address, destination address, and application name. If you do not want to specify a specific application, enter any as the default application, indicating all possible applications. For example, if you do not supply an application name, then the policy is installed with the application as a wildcard (default). Therefore, any data traffic that matches the rest of the parameters in a given policy would match the policy regardless of the application type of the data traffic.

When you are creating a policy, the following policy rules apply:

- Security policies are configured in a from-zone to to-zone direction. Under a specific zone direction, each security policy contains a name, match criteria, an action, and miscellaneous options.
- The policy name, match criteria, and action are required.
- The policy name is a keyword.
- The source address in the match criteria is composed of one or more address names or address set names in the from-zone.
- The destination address of the match criteria is composed of one or more address names or address set names in the to-zone.
- The application name in the match criteria is composed of the name of one or more applications or application sets.
- One of the following actions is required: permit, deny, or reject.
- When logging is enabled, the system logs at session close time by default. You can enable logging at session creation, too.
- When the count alarm is turned on, you can, optionally, specify alarm thresholds in bytes per second and kilobytes per minute.
- You cannot specify global as either the from-zone or the to-zone except under following condition:

Any policy configured with the to-zone as a global zone must have a single destination address to indicate that either static NAT or incoming NAT has been configured in the policy.

- In SRX Series Services Gateways, policy permit option with NAT is simplified. Each policy will optionally indicate whether it allows NAT translation, no NAT translation or do not care.
- Address names cannot begin with the following reserved prefixes. These are used only for address NAT configuration:
  - static_nat_
  - incoming_nat_
  - junos_
- Application names cannot begin with the junos_ reserved prefix.

Understanding Policy Elements

A policy permits, denies, or tunnels specified types of traffic unidirectionally between two points.

To define a policy, you need:

- An incoming zone (the from-zone)
- An outgoing zone (the to-zone)
- An ordered set of policies between the from-zone and to-zone

Each policy consists of:

- A unique name for the policy.
- A set of match criteria defining the conditions that must be satisfied to apply the policy rule. The match criteria are based on a source IP address, destination IP address, and applications.
- A set of actions to be performed in case of a match—permit, deny, or reject.
- Accounting and auditing elements—counting, logging, or structured system logging.

The following example shows a policy configuration that allows traffic from the green zone (from-zone) to the red zone (to-zone).

- The red zone's address book contains the abc address.
- The green zone's address book contains the public address.

```
user@host# set security policies from-zone green to-zone red policy
abctopublic
match source-address abc
```

```
user@host# set security policies from-zone red to-zone green policy
abctopublic
match destination-address public

user@host# set security policies from-zone red to-zone green policy
abctopublic
match application ssh

user@host# set security policies from-zone red to-zone green policy
abctopublic
then permit
```

For more information on the policy configuration syntax and options, see the *JUNOS Software CLI Reference*.

Understanding Policy Configuration

You must complete the following tasks to create a security policy:

1. Configure a security zone's host addresses and subnet addresses in its address book.
2. Configure application or application sets.
3. Create the policy.
4. Create schedulers if you plan to use them for your policies.
5. Bind a policy to a schedule.

Related Topics

- Configuring Policies on page 80
- Security Policy Address Books and Address Sets on page 97
- Security Policy Schedulers on page 105
- Security Policy Applications on page 115

Understanding Policy Ordering

JUNOS Software offers a tool for verifying that the order of policies in the policy list is valid.

Before You Begin
For background information, read “Security Policies Overview” on page 69.

It is possible for one policy to eclipse, or *shadow*, another policy. Consider the following example:

```
user@host# set security policies from-zone internal to-zone external policy 1 match
  source-address any
user@host# set security policies from-zone internal to-zone external policy 1 match
  destination-address any
user@host# set security policies from-zone internal to-zone external policy 1 match
  application junos-http
user@host# set security policies from-zone internal to-zone external policy 1 then
  permit
user@host# set security policies from-zone internal to-zone external policy 2 match
  source-address any
user@host# set security policies from-zone internal to-zone external policy 2 match
  destination-address any
user@host# set security policies from-zone internal to-zone external policy 2 match
  application junos-http
user@host# set security policies from-zone internal to-zone external policy 2 then deny
```

Because JUNOS Software performs a policy lookup starting from the top of the list, when it finds a match for traffic received, it does not look any lower in the policy list. In the previous example, JUNOS Software never reaches policy 2 because the policy matches identical traffic that policy 1 matches. This matched traffic has already been acted upon with a policy action and is no longer evaluated by succeeding policies.

To correct the previous example, you can simply reverse the order of the policies, putting the more specific one first:

```
user@host# set security policies from-zone internal to-zone external policy 2 match
  source-address any
user@host# set security policies from-zone internal to-zone external policy 2 match
  destination-address any
user@host# set security policies from-zone internal to-zone external policy 2 match
  application junos-http
user@host# set security policies from-zone internal to-zone external policy 2 then deny
user@host# set security policies from-zone internal to-zone external policy 1 match
  source-address any
user@host# set security policies from-zone internal to-zone external policy 1 match
  destination-address any
user@host# set security policies from-zone internal to-zone external policy 1 match
  application junos-http
user@host# set security policies from-zone internal to-zone external policy 1 then
  permit
```

Of course, this example is purposefully simple to illustrate the basic concept. In cases where there are dozens or hundreds of policies, the eclipsing of one policy by another might not be so easy to spot. To check if there is any policy shadowing in your policy list, you can use the following CLI command:

```
show policy-options <policy-name>
```

This command reports the shadowing and shadowed policies. It is then the administrator's responsibility to correct the situation.

NOTE: The concept of policy *shadowing* refers to the situation where a policy higher in the policy list always takes effect before a subsequent policy. Because the policy lookup always uses the first policy it finds that matches the five-part tuple of source and destination zone, source and destination address, and application type, if another policy applies to the same tuple (or a subset of the tuple), the policy lookup uses the first policy in the list and never reaches the second one.

The policy verification tool cannot detect the case where a combination of policies shadows another policy. In the following example, no single policy shadows policy 3; however, policies 1 and 2 together do shadow it:

```
user@host# set security zones security-zone trust address-book address-set grp1
   address host1
user@host# set security zones security-zone trust address-book address-set grp1
   address host2
user@host# set policy id 1 from trust to untrust host1 server1 HTTP permit
user@host# set policy id 2 from trust to untrust host2 server1 HTTP permit
user@host# set policy id 3 from trust to untrust grp1 server1 HTTP deny
```

Related Topics

- Understanding Policies on page 70
- Example: Configuring Security Policies—Detailed Configuration on page 86

Configuring Policies—Quick Configuration

You can use J-Web Quick Configuration to quickly configure security policies.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Configure security zones and interfaces. See "Configuring Security Zones—Quick Configuration" on page 51.
3. Configure address books. See "Configuring Address Books" on page 102.

To configure security policies in the J-Web user interface:

1. Select **Security > Policy > FW Policies**.
2. If you want to change the default action for all new security policies, from the **Default Policy Action** list in the upper-right corner, select **permit-all** or **deny-all** and click **OK**.
3. Select one of the following policy configuration options:

- **Add**—Click **Add** to create a new security policy. The Add Policy dialog box appears.
- **Edit**—Select an existing policy and click **Edit** to change the selected policy's configuration. The Edit Policy dialog box appears.
- **Delete**—Select an existing policy and click **Delete** to delete the selected policy. (You can skip the remaining steps in this section if you select this option.)
- **Clone**—Select an existing policy and click **Clone** to create a copy of the selected policy under a different name. The Clone policy dialog box appears.
- **Deactivate/Activate**—Select an existing policy and click **Deactivate/Activate** to disable or enable processing based on the selected policy. (You can skip the remaining steps in this section if you select this option.)
- **Move**—Select an existing policy and click **Move** to move the policy up or down in the list. For information about how policy order affects how policies are applied and processed, see "Understanding Policy Ordering" on page 74. (You can skip the remaining steps in this section if you select this option.)

4. Select the **Policy** tab to specify the basic configuration requirements for the policy, including the name, from-zone, to-zone, source address, destination address, policy applications, and policy action. See Table 13 on page 77 for more information.
5. Select the **Logging/Count** tab to specify logging requirements for the policy. See Table 14 on page 79 for more information.
6. Select the **Scheduling** tab to select a scheduler whose schedule determines when the policy is active. (You must create the scheduler that you want to include in the policy before creating the policy.)
7. Select the **Permit Action** tab to specify the VPN, NAT, and authentication settings for the policy. See Table 15 on page 79 for more information. (This tab appears only if you select **Permit** as the policy action in the Policy tab.)
8. Select the **Application Services** tab to specify the IDP, UTM, and WX settings for the policy. See Table 16 on page 80 for more information. (This tab appears only if you select **Permit** as the policy action in the Policy tab.)
9. Select one of the following options:
   - **OK**–To apply the configuration and return to the main Configuration page, click **OK**.
   - **Cancel**–To cancel your entries and return to the main page, click **Cancel**.

Table 13: Policy Configuration Options

Field	Description
Policy Name	Specify a name for the security policy.
From Zone	Specify the source zone for the policy. (You must create the zones that you want to include in the policy before creating the policy.)

Table 13: Policy Configuration Options *(continued)*

Field	Description
To Zone	Specify the destination zone for the policy. (You must create the zones that you want to include in the policy before creating the policy.)
Source Address	Specify the name of the source address or address set for the policy (as entered in the source zone's address book) and move it to the **Matched** list using the arrows. If you want to add a new address to the list, select **Add New Source Address**. In the fields that appear, specify the new address and click **Add**. **NOTE:** Address names cannot begin with the following reserved prefixes. These prefixes are used only for address NAT configuration: ■ static_nat_ ■ incoming_nat_ ■ junos_
Destination Address	Specify the name of the destination address or address set for the policy (as entered in the source zone's address book) and move it to the **Matched** list using the arrows. If you want to add a new address to the list, select **Add New Destination Address**. In the fields that appear, specify the new address and click **Add**. **NOTE:** Address names cannot begin with the following reserved prefixes. These prefixes are used only for address NAT configuration: ■ static_nat_ ■ incoming_nat_ ■ junos_
Applications	Specify the name of an application or application set to which the policy applies and move it to the **Matched** list using the arrows. If you do not want to specify an application, select **any** as the default application.
Policy Action	Specify the actions that need to take place on the traffic as it passes through the firewall: ■ **Permit**—Allows the packet to pass through the firewall. ■ **Reject**—Blocks the packet from traversing the firewall. The firewall drops the packet and sends a TCP reset (RST) segment to the source host for TCP traffic and an ICMP destination unreachable, port unreachable message (type 3, code 3) for UDP traffic. For TCP and UDP traffic, the firewall drops the packet and notifies the source host as action Deny. ■ **Deny**—Blocks and drops the packet from traversing the firewall, but does not send notification back to the source.

Table 14: Logging/Count Configuration Options

Field	Description
Enable Count	
Enable Count	Select this option to enable counting. If counting is enabled, counters are collected for the number of packets, bytes, and sessions that enter the firewall for a given policy. For counts (only for packets and bytes), you can specify that alarms be generated whenever the traffic exceeds specified thresholds. **NOTE:** The alarm threshold fields are disabled if **Enable Count** is not selected.
Per Minute Alarm Threshold	Specify threshold bytes for the per-minute alarm threshold. Enter any value from 0 through 4294967295 kilobytes.
Per Second Alarm Threshold	Specify threshold bytes for the per-second alarm threshold. Enter any value from 0 through 4294967295 kilobytes.
Log Options	
Log at Session Close Time	Select this option if you want to log the events when the session closed.
Log at Session Init Time	Select this option if you want to log the events when the session is created.

Table 15: Permit Action Configuration Options

Field	Description
Tunnel — IPSec VPN	
VPN	Specify the name of the IPsec-VPN tunnel.
Pair Policy	
Pair Policy Name	Specify the name of the policy with the same IPsec-VPN in the reverse direction to create a pair policy.
NAT Translation	
Options	Select one of the following options: ■ None ■ Drop packets with translated address ■ Drop packets without translated address
Firewall Authentication	Use these options to authenticate the client before forwarding the traffic. The two types of firewall authentication are: ■ Pass-through authentication ■ Web authentication
Pass-through	Use pass-through authentication verifies traffic as it attempts to pass through the firewall.

Table 15: Permit Action Configuration Options *(continued)*

Field	Description
Access Profile	Select the access profile for the pass-through from the drop-down list.
Client name	Specify the client name for the pass-through.
Web Redirect	Select the Web Redirect option if you want to redirect the pass-through traffic for Web authentication.
Web authentication	Use Web authentication to verify client authentication.
Client name	Specify the client name for the Web authentication.

Table 16: Application Services Configuration Options

Field	Description
IDP	
Enable IDP	Select this option to enable IDP for the policy.
UTM Policy	
UTM Policy	Select the required UTM policy from the drop-down list.
Redirect	
Options	Select one of the following options: ■ None. ■ Redirect-wx ■ Reverse Redirect-wx

Configuring Policies

A security policy applies security rules to transit traffic.

Before You Begin

Before creating policies to control traffic between different security zones, you must first design the environment in which to apply those policies:

1. Create zones. See “Creating Security Zones” on page 49.
2. Configure the address book for the policy. (See “Configuring a Policy to Permit Traffic” on page 86.)
3. Create addresses for use in the policies.
4. Create an application (or application set) that indicates that the policy applies to traffic of that type.

   After completing Step 1 through Step 4, you can then create the policies necessary to permit or deny traffic in and out of your protected network. For background information, read “Security Policies Overview” on page 69.

A packet is matched against policies to determine how it is to be treated. The packet is matched against a policy's source and destination zones, source and destination addresses, and the application or application sets that the policy specifies. If the packet matches all elements of a policy, that policy's action is applied to the packet. See “Understanding Policy Rules” on page 71.

The action of the first policy that the traffic matches is applied to the packet. If there is no matching policy, the packet is dropped. Policies are searched from top to bottom, so it is a good idea to place more specific policies near the top of the list. You should also place IPsec VPN tunnel policies near the top. Place the more general policies, such as one that would allow certain users access to all Internet applications, at the bottom of the list.

Policies are applied after the packet has passed through the firewall's screens and the system has looked up its route. The packet's destination address determines its destination zone.

Depending on the policies you create, any of the actions shown in Table 13 on page 77 could be applied to the packet.

To define a policy, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 81
- CLI Configuration on page 84
- Related Topics on page 84

J-Web Configuration

To configure policies using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Interfaces, click **Edit**.
3. Next to Interface, click **Add new entry**.
4. In the Interface name box, type **ge-0/0/1**.
5. Next to Accounting profile box type **profile1**.
6. Provide a description for the interface.
7. Select the **Disable** option if the interface is disabled.
8. Enter the maximum transmit packet size. The size can be any value between 256 and 9192 MTUs.
9. Select **Traceoptions** if you want to configure traceoptions. Next to trace options, click **Configure**.
10. Next to Flag, click **Add new entry**.
11. Select the **Flag name** from the drop down list.
12. Click **OK** to return back to the traceoptions configuration page.
13. Click **OK** again to return back to the Interface configuration page.

14. Next to Unit, click **Add new entry**.
15. In the Interface unit number box, type **0**.
16. Specify the accounting profile.
17. Specify the logical bandwidth unit.
18. Select **Disable** if you want to disable the interface unit.
19. Select either of the following:
    - **Yes**, if you want to configure traps for this interface.
    - **No**, if you do not want to configure traps for this interface.
20. Under Family, select **inet** and click **Configure**. Enter the required information and click **OK** to return to the Family configuration page.
21. Under Family, select **iso** and click **Configure**. Enter the required information and click **OK** to return to the Family configuration page.
22. Click **OK** to return to the Unit configuration page.
23. Click **OK** again to return to the Interfaces configuration page.
24. To configure the other interfaces such as ge-0/0/2 and ge-0/0/3 and addresses such as 1.2.2.1/24 and 1.1.1.1/24, repeat Step 2 through Step 21 and click **OK**.
25. Click **Commit** if you are done with the configuration.

To configure interfaces:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zones, click **Add new entry**.
5. In the Name box, type **internal**.
6. Next to Interfaces, click **Add new entry**.
7. In the Interface unit box, type **ge-0/0/1.0** and click **OK**.
8. To configure security zones such as dmz and external, and interfaces such as ge-0/0/2.0 and ge-0/0/3.0, repeat Step 3 through Step 7 and click **OK**.

To configure addresses:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zones, click **Add new entry**.
5. In the Name box, type **internal**.
6. Next to the Address book, click **Configure** or **Edit**.
7. Next to Address, click **Add new entry**.

8. In the Address name box, type **corp_net 10.1.1.0/24** and click **OK**.
9. To configure more security zones such as dmz and external, and address books entries such as mail_svr 1.2.2.5/32 and r-mail_svr 2.2.2.5/32, repeat Step 3 through Step 8 and click **OK**.

To configure application sets:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Applications, click **Configure** or **Edit**.
3. Next to Application set, click **Add new entry**.
4. In the Application set name box, type **MAIL-POP3**.
5. In the Application name box, type **junos-mail** and click **OK**.

To create policies:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, select **Configure** or **Edit**.
3. Next to Policy, select the check box and click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type **internal**.
6. In the To zone name box, type **dmz corp_net**.
7. Next to Policy, click **Add new entry**.
8. In the Policy name box, type **Pol1**.
9. Select the **Match** check box.
10. Select the **Then** check box.
11. Next to Match, click **Configure**.
12. To match the policy to an application name, from the Source address choice list, select **Source address**.
13. Next to Source address, click **Add new entry**.
14. From the Value keyword list, select **Enter specific value**.
15. To specify the address of the address book, In the Address box, type **corp_net** and click **OK**.
16. To match the policy to a destination address, from the Destination address choice list, select **Destination address**.
17. Next to Destination address, click **Add new entry**.
18. In the Value keyword list, select **Enter specific value**.
19. To specify the destination address for the policy, in the Address box, type **mail_svr** and click **OK**.
20. To match the policy to an application set name, from the Application Choice list, select **Application**.

21. Next to Application, click **Add new entry**.
22. To specify the application set name to match the policy, from the Value keyword list, type **MAIL-POP3** and click **OK**.
23. Next to Then, click **Configure**.
24. From the Action list, select **Permit** and click **OK**.
25. If you are finished configuring the device, commit the configuration.
26. To check the configuration, see "Verifying Policy Configuration" on page 85.

CLI Configuration

1. Set interfaces.

```
user@host# set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.2.2.1/24
user@host# set interfaces ge-0/0/3 unit 0 family inet address 1.1.1.1/24
user@host# set security zones security-zone internal interfaces ge-0/0/1.0
user@host# set security zones security-zone dmz interfaces ge-0/0/2.0
user@host# set security zones security-zone external interfaces ge-0/0/3.0
```

2. Configure addresses.

```
user@host# set security zones security-zone internal address-book address
  corp_net 10.1.1.0/24
user@host# set security zones security-zone dmz address-book address mail_svr
  1.2.2.5/32
user@host# set security zones security-zone external address-book address
  r-mail_svr 2.2.2.5/32
```

3. Configure application sets.

```
user@host# set applications application-set MAIL-POP3 application junos-mail
user@host# set applications application-set MAIL-POP3 application junos-pop3
```

4. Create policies.

```
user@host# set security policies from-zone internal to-zone dmz corp_net mail_svr
  MAIL-POP3 permit
user@host# set security policies from-zone dmz to-zone external mail_svr r-mail_svr
  MAIL permit
user@host# set security policies from-zone external to-zone dmz r-mail_svr mail_svr
  MAIL permit
```

5. If you are finished configuring the device, commit the configuration.
6. To check the configuration, see "Verifying Policy Configuration" on page 85.

Related Topics

- Understanding Policies on page 70
- Configuring Address Books on page 102

- Policy Application Sets Overview on page 116
- Example: Configuring Applications and Application Sets on page 142
- Configuring Schedulers on page 107
- Verifying Policy Configuration on page 85

Verifying Policy Configuration

A scheduler is referred by security policies to activate or deactivate a policy according to scheduled times. Use the show schedulers command to verify that your configured policies are associated with the appropriate schedulers. For more information, see the *JUNOS Software CLI Reference*

Purpose Display information about address books and zones.

Action Use the show security policies CLI command to display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy:

```
user@host> show security policies
From zone: ZoneA, To zone: ZoneB
 Policy: p1, Sequence number: 1, State: enabled
  Source addresses: v-2-2-2-0
  Destination addresses: v-1-1-1-0
  Applications: any
  Action: permit, log, scheduled
 Policy: p2, Sequence number: 2, State: enabled
  Source addresses: v-2-2-2-0
  Destination addresses: v-1-1-1-0
  Applications: any
  Action: deny, scheduled
```

What it Means The output displays information about policies configured on the system. Verify the following information:

- From and to zones.
- Source and destination addresses.
- Match criteria.

Example: Configuring Security Policies—Detailed Configuration

The following instructions include a sample configuration and step-by-step guidelines explaining how to configure policies to protect resources and allow traffic through the firewall.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. For background information, read:
   - Security Policies Overview on page 69
   - Address Books and Address Sets Overview on page 97

To following steps show how to configure policies:

1. Configure a policy to permit traffic. For more information, see "Configuring a Policy to Permit Traffic" on page 86.
2. Configure a policy to deny traffic. For more information, see "Configuring a Policy to Deny Traffic" on page 90.
3. Configure a policy with pass-through authentication. For more information, see "Configuring for Pass-Through Authentication" on page 310.
4. Configure a policy with WebAuth authentication. For more information, see "Configuring for Web Authentication" on page 315.
5. Reorder policies after they have been created. For more information, see "Reordering Policies After They Have Been Created" on page 92.

Configuring a Policy to Permit Traffic

Configuring a policy to permit traffic is the first step in the sample configuration explaining how to configure a policy.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Create zones. See "Creating Security Zones" on page 49.
3. Configure the address book for the policy. (See "Configuring a Policy to Permit Traffic" on page 86.)
4. For background information, read "Example: Configuring Security Policies—Detailed Configuration" on page 86.

To configure a policy to permit traffic, use either J-Web or the CLI configuration editor. The following configuration commands allow traffic between the loopback addresses of both the Juniper Networks devices.

This topic covers:

- J-Web Configuration on page 87
- CLI Configuration on page 89
- Related Topics on page 89

J-Web Configuration

To configure a policy to permit traffic using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI.**
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type **green**.
6. In the To zone name box, type **red**.
7. Next to Policy, click **Add new entry**.
8. In the Policy name box, type **allowin**.
9. Select the **Match** check box and click **Configure**.
    a. From the Source address choice list, select **Source address**.
    b. Next to Source address, click **Add new entry**.
    c. From the Value keyword list, select **Enter specific value**.
    d. In the Address box, type **netTopLoopInt** and click **OK**.
    e. To match the policy to a destination address, from the Destination address choice list, select **Destination address**.
    f. Next to Destination address, click **Add new entry**.
    g. From the Value keyword list, select **Enter specific value**.
    h. In the Address box, type **netBottomLoopInt** and click **OK**.
    i. To match the policy to an application set name, from the Application Choice list, select **Application**.
    j. Next to Application, click **Add new entry**.
    k. Specify the application or application set name to match the policy and click **OK**.
    l. Expand the Advanced option and do either of the following:
        - Next to Apply groups, click **Add new entry**. Specify the group to inherit the configuration.
        - Next to Apply groups except, click **Add new entry**. Specify the group not to inherit the configuration.

10. Specify the name of the Scheduler in the Scheduler Name text box.
11. Select the **Then** check box and click **Configure**.
    - Select the **Count** check box to enable counting.
    - Next to Log, click **Configure**.
    - Do any or all of the following:
      - Select **Session Close**, if you want the log files to be created towards the end of a session.
      - Select **Session Init**, if you want the log files to be created during the beginning of a session.

      You can select both the options to enable log files to be created both during the beginning as well as towards the end of a session.
    - Expand the Advanced option and do either of the following:
      - Next to Apply groups, click **Add new entry**. Specify the group to inherit the configuration.
      - Next to Apply groups except, click **Add new entry**. Specify the group not to inherit the configuration.
    - Click **OK** to return to the main configuration page.
12. From the Action list, select **Permit** .
13. Next to Action, click **Add new entry**.
    a. Next to Application services, click **Configure**.
    b. Select any or all of the following:
       - **Idp**, if you want to permit IDP.
       - **Uac policy**, if you want to permit UAC Policy.
    c. Expand the Advanced option and do either of the following:
       - Next to Apply groups, click **Add new entry**. Specify the group to inherit the configuration.
       - Next to Apply groups except, click **Add new entry**. Specify the group not to inherit the configuration.
    d. Click **OK** to save the configuration and return to the Permit configuration page.
    e. Select the Destination address check box and click **Configure**.
    f. Select any one of the option from the **destination nat** drop down list:
       - Drop Translated.
       - Drop Untranslated.
    g. Click **OK** to return back to the Permit configuration page.

h. Next to Firewall authentication, click **Configure**.

i. Select either of the following from the Auth type drop down list:

- **Pass through**. Next to Pass through click **Configure** and specify the access profile name, client user name or group name to match. Also select the Web redirect check box. You can also specify access groups and access groups except options.
- **Web authentication**. Next to Web authentication click **Configure** and specify the client user name or group name to match. You can also specify access group and access group except options.

After specifying the firewall authentication information, click **OK** to return to the Permit configuration page.

j. Next to Tunnel, click **Configure**.

k. Specify the name of the IPsec VPN to enable it. This option is mandatory.

l. Specify the policy in the reverse direction to form a pair.

m. Expand the Advanced option and do either of the following:

- Next to Apply groups, click **Add new entry**. Specify the group to inherit the configuration.
- Next to Apply groups except, click **Add new entry**. Specify the group not to inherit the configuration.

Click **OK** to return to the Then configuration page.

14. If you are finished configuring the device, commit the configuration.
15. To check the configuration, see "Verifying Policy Configuration" on page 85.

CLI Configuration

```
user@host# set security policies from-zone RED to-zone GREEN policy allowIn match
    source-address netTopLoopInt
user@host# set security policies from-zone RED to-zone GREEN policy allowIn match
    destination-address netBottomLoopInt
user@host# set security policies from-zone RED to-zone GREEN policy allowIn match
    application any
user@host# set security policies from-zone RED to-zone GREEN policy allowIn then
    permit
```

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

To check the configuration, see "Verifying Policy Configuration" on page 85.

Related Topics

- Configuring a Policy to Deny Traffic on page 90

■ Monitoring Policy Statistics on page 94

Configuring a Policy to Deny Traffic

Configuring a policy to deny traffic is the second step in the sample configuration explaining how to configure a policy.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Create zones. See "Creating Security Zones" on page 49.
3. Configure the address book for the policy. (See "Configuring a Policy to Permit Traffic" on page 86.)
4. Configure a policy to permit traffic. (See "Configuring a Policy to Permit Traffic" on page 86.)
5. For background information, read "Example: Configuring Security Policies—Detailed Configuration" on page 86.

To configure a policy to deny traffic, use either J-Web or the CLI configuration editor. The following configuration commands deny traffic between the loopback addresses of both the Juniper Networks devices.

This topic covers:

■ J-Web Configuration on page 90
■ CLI Configuration on page 91
■ Related Topics on page 91

J-Web Configuration

To configure a policy to deny traffic using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure**.
4. In the From zone name box, type **green**.
5. In the To zone name box, type **red**.
6. In the Policy name box, type **denyin**.
7. Select the **Match** check box.
8. Select the **Then** check box.
9. Next to Match, click **Configure**.
10. To match the policy to an application name, from the Source address choice list, select **Source address**.
11. Next to Source address, click **Add new entry**.

12. From the Value keyword list, select **Enter specific value**.
13. In the Address box, type **netTopLoopInt** and click **OK**.
14. To match the policy to a destination address, from the Destination address choice list, select **Destination address**.
15. Next to Destination address, click **Add new entry**.
16. From the Value keyword list, select **Enter specific value**.
17. In the Address box, type **netBottomLoopInt** and click **OK**.
18. To match the policy to an application set name, from the Application choice list, select **Application**.
19. Next to Application, click **Add new entry**.
20. To specify the application set name to match the policy, from the Value keyword list, select **any** and click **OK**.
21. To complete matching the policy to the specific source address, destination address, and application name, click **OK**.
22. Next to Then, click **Configure**.
23. From the Action list, select **Deny** and click **OK**.
24. If you are finished configuring the device, commit the configuration.
25. To check the configuration, see "Verifying Policy Configuration" on page 85.

CLI Configuration

```
user@host# set security policies from-zone RED to-zone GREEN policy DenyIn match
    source-address netTopLoopInt
user@host# set security policies from-zone RED to-zone GREEN policy DenyIn match
    destination-address netBottomLoopInt
user@host# set security policies from-zone RED to-zone GREEN policy DenyIn match
    application any
user@host# set security policies from-zone RED to-zone GREEN policy DenyIn then
    deny
```

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

To check the configuration, see "Verifying Policy Configuration" on page 85.

Related Topics

- Configuring for Web Authentication on page 315
- Verifying Scheduled Policies on page 112

Reordering Policies After They Have Been Created

Reordering policies is the last step in the sample configuration explaining how to configure a policy.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Create zones. (See "Creating Security Zones" on page 49.)
3. Configure the address book for the policy. (See "Configuring a Policy to Permit Traffic" on page 86.)
4. Configure a policy to permit traffic. (See "Configuring a Policy to Permit Traffic" on page 86.)
5. Configure a policy to deny traffic. (See "Configuring a Policy to Deny Traffic" on page 90.)
6. Configure a policy with pass-through authentication. (See "Configuring for Pass-Through Authentication" on page 310.)
7. Configure a policy with WebAuth authentication. (See "Configuring for Web Authentication" on page 315.)
8. For background information, read "Example: Configuring Security Policies—Detailed Configuration" on page 86.

The following configuration commands show how to move policies around after they have been created.

```
user@host#edit security policies from-zone RED to-zone GREEN
insert policy passTauth after policy wAuth
```

Related Topics

- Understanding Policies on page 70
- Understanding Policy Configuration on page 74

Troubleshooting Policy Configuration

Most policy configuration failures occur during a commit or runtime.

This topic covers:

- Checking Commit Failure on page 93
- Verifying Commit on page 93
- Debugging Policy Lookup on page 93

Checking Commit Failure

Commit failures are reported directly on the CLI when you execute the CLI command commit-check in configuration mode. These errors are configuration errors and you cannot commit the configuration without fixing these errors.

To fix these errors, do the following:

1. Review your configuration data.
2. Open the file /var/log/nsd_chk_only. This file is overwritten each time you perform a commit check and contains detailed failure information.

Verifying Commit

Upon performing a policy configuration commit, if you notice that the system behavior is incorrect, use the following steps to troubleshoot this problem:

1. Operational show Commands—Execute the operational commands for security policies and verify that the information shown in the output is consistent with what you expected. If not, the configuration needs to be changed appropriately.
2. Traceoptions—Set the traceoptions command in your policy configuration. The flags under this hierarchy can be selected as per user analysis of the show command output. If you cannot determine what flag to use, the flag option all can be used to capture all trace logs.

```
user@host# set security policies traceoptions <flag all>
```

You can also configure an optional filename to capture the logs.

```
user@host# set security policies traceoptions <filename>
```

If you specified a filename in the traceoptions, you can look in the /var/log/ < *filename* > for the log file to ascertain if any errors were reported in the file. (If you did not specify a filename, the default filename is eventd.) The error messages indicate the place of failure and the appropriate reason.

After configuring the trace options, you must recommit the configuration change that caused incorrect system behavior.

Debugging Policy Lookup

When you have the correct configuration, but some traffic was incorrectly dropped or permitted, you can enable the lookup flag in the security policies traceoptions. The lookup flag logs the lookup related traces in the trace file.

```
user@host# set security policies traceoptions <flag lookup>
```

Monitoring Policy Statistics

JUNOS Software can monitor and record traffic that it permits or denies based on previously configured policies.

Before You Begin

For background information, read:

- Security Policies Overview on page 69
- Understanding Policy Rules on page 71

To monitor traffic, enable the count and log options.

Count—Can be configured in an individual policy. If count is enabled, counters are collected for the number of packets, bytes, and sessions that enter the firewall for a given policy. For counts (only for packets and bytes), you can specify that alarms be generated whenever the traffic exceeds specified thresholds.

Log—Consists of trace options and a structured system log.

- Trace options: Tracing operations for a policy.

  To trace security policies, include the traceoptions command at the [edit] hierarchy level.

  ```
  set security policies traceoptions < filename > < flag >
  ```

  filename—Name of the file in which the output of the tracing operation is saved. All files are placed in the directory /var/log < " filename" >. Enclose the name of the security-trace file within quotation marks. By default, commit script process tracing output is placed in the file. If you include the file command, you must specify a filename. To retain the default, you can specify eventd as the filename.

  The default file size is 128 KB, and 10 files are created before the first one gets overwritten.

  flag—Tracing operation to perform. To perform more than one tracing operation, include multiple *flag* commands. You can include the following flags:

  - all—All tracing operations
  - authentication—Trace authentication events
  - configuration—Trace configuration events
  - setup—Trace setup of firewall authentication service

- Structured Log: Has one directive for all policies.

For details about information collected for session logs, please see “Information Provided in Session Log Entries” on page 30.

Chapter 5

Security Policy Address Books and Address Sets

Each security zone contains an address book. Before you can set up policies between two zones, you must define the addresses for each of the zone's address books. To manage an address book with large numbers of addresses, you can create groups of addresses called address sets.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Address Books and Address Sets Overview on page 97
- Configuring Addresses and Address Sets—Quick Configuration on page 100
- Configuring Address Books on page 102
- Verifying Address Book Configuration on page 103

Address Books and Address Sets Overview

A security zone is a logical group of interfaces with identical security requirements. Each security zone contains an address book. Before you can set up policies between two zones, you must define the addresses for each of the zone's address books. A zone's address book must contain entries for the addressable networks and end hosts (and, thus, users) belonging to the zone.

Understanding Address Books

The following guidelines apply to address books:

- An address book for a security zone contains the IP address or domain names of hosts and subnets whose traffic is either allowed, blocked, encrypted, or user-authenticated.
- Address books can have address sets. Each address set has a name and a list of address names.
- Addresses and address sets in the same zone must have distinct names.
- Addresses must conform to the security requirements of the zone.

- IP addresses can be configured as IPv4 addresses with the number of prefix bits, or as Domain Name System (DNS) names.
- The predefined address any is automatically created for each security zone.
- The address book of a security zone must contain all IP addresses that are reachable within that zone.

Policies contain both source and destination zones and addresses. An address is referred to in a policy by the name you give it in its zone's address book.

- When traffic is sent to a zone, the zone and address to which the traffic is sent are used as the destination zone and address-matching criteria in policies.
- When traffic is sent from a zone, the zone and address from which it is sent are used as the matching source zone and address in policies.

For more information on the address book configuration syntax and options, see the *JUNOS Software CLI Reference*

NOTE: Specify addresses as network prefixes in the prefix/length format. For example, 1.2.3.0/24 is an acceptable address book address because it translates to a network prefix. However, 1.2.3.4/24 is not acceptable for an address book because it exceeds the subnet length of 24 bits. Everything beyond the subnet length must be entered as 0 (zero). In special scenarios, you can enter a hostname because it can use the full 32-bit address length.

Understanding Address Sets

An address book can grow to contain large numbers of addresses and become difficult to manage. To manage an address book with large numbers of addresses, you can create groups of addresses called *address sets*. You can reference an address set in a policy as you would an individual address book entry.

The following example shows addresses and address sets in the green zone:

```
user@host# set security zones security-zone green address-book address src_addr1
   64.10.4.44/32
user@host# set security zones security-zone green address-book address src_addr2
   64.10.9.28/32
user@host# set security zones security-zone green address-book address src_addr3
   10.10.10.10/24
user@host# set security zones security-zone green address-book address bbc dns-name
   www.bbc.com
user@host# set security zones security-zone green address-book address-set
   my_source_addresses address src_addr1
user@host# set security zones security-zone green address-book address-set
   my_source_addresses address src_addr2
user@host# set security zones security-zone green address-book address-set
   my_source_addresses address src_addr3
```

For more information on the address set configuration syntax and options, see the *JUNOS Software CLI Reference*

NOTE: Consider that for each address set, the system creates individual rules for its members. It creates an internal rule for each member in the group as well as for each service configured for each user. If you configure address books without taking this into account, you can exceed the number of available policy resources, especially if both the source and destination addresses are address groups and the specified service is a service group.

When you add addresses to policies, sometimes the same subset of addresses can be present in multiple policies, making it difficult to manage how policies affect each address entry. JUNOS Software allows you to create groups of addresses called *address sets*. Address sets simplify the process by allowing you to add multiple addresses within an address set and therefore manage a small number of address sets, rather than manage a large number of individual address entries. See Figure 13 on page 99.

Figure 13: Address Sets

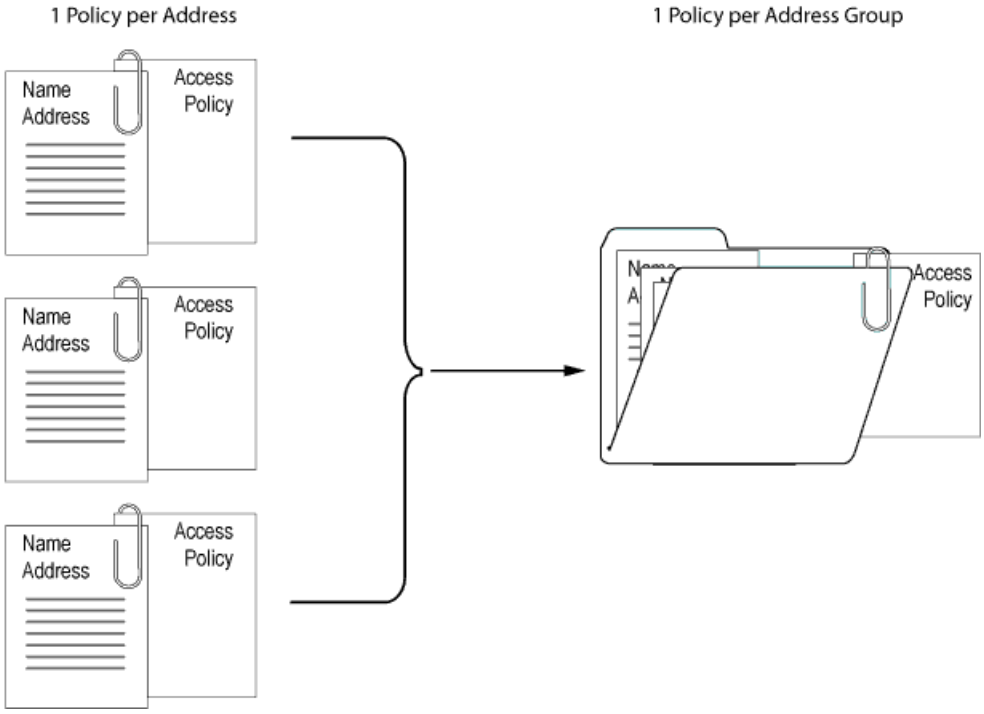

The address set option has the following features:

- You can create address sets in any zone.
- You can create address sets with existing users, or you can create empty address sets and later fill them with users.
- You can reference an address set entry in a policy like an individual address book entry.

NOTE: JUNOS Software applies policies automatically to each address set member, so you do not have to create them one by one for each address. Furthermore, JUNOS Software writes these policies to ASIC, which makes lookups run very fast.

- When you delete an individual address book entry from the address book, you must remove the address (wherever it is referred) from all the address sets.

The following constraints apply to address sets:

- To configure an address set, you need more than an address in the address book.
- Address sets can only contain address names that belong to the same security zone.
- Address names cannot be the same as address set names. For example, if the name Paris is used for an address in an individual address entry, it cannot be used for an address set name.
- If an address set is referenced in a policy, the address set cannot be removed without removing its reference in the policy. It can, however, be edited.
- You cannot add the predefined address any to an address book.

Configuring Addresses and Address Sets—Quick Configuration

You can use J-Web Quick Configuration to quickly configure address books and address sets.

Before You Begin

1. For background information, read "Understanding Security Zones" on page 47 and "Configuring Address Books" on page 102.
2. Configure zones. For more information, see "Configuring Security Zones—Quick Configuration" on page 51.

This topic contains the following instructions:

- Configuring Addresses on page 100
- Configuring Address Sets on page 101

Configuring Addresses

To configure an address in the J-Web user interface:

1. Select **Configure > Security > Policy Elements > Address Book**.
2. Select the **Addresses** tab.
3. Select one of the following options:

- **Add**–To create a new address, click **Add**. The Add Address dialog box appears.
- **Edit**–To edit an existing configuration, select the address that you want to change and click **Edit**. The Edit Address dialog box appears.
- **Delete**–To delete an existing address, select it and click **Delete**. (If you select this option, you can skip the remaining steps in this section.)

4. In the **Zone** field, specify the zone to which the address should apply.
5. In the **Address Name** field, specify a name for the address.
6. Select the **IP/Prefix** or **Domain Name** option. Then, enter the IP address or domain name of the address in the field that appears. The IP address must be an IPv4 address with the number of prefix bits. You can use domain names only if the system is configured to use DNS services.
7. (Optional) Specify an address set to which the address should be added. To create a new address set:
   - In the **Add Address Set** text box, enter an address set.
   - Click **Add**. (To remove the address after creating it, click **Undo**.)
8. Select one of the following options:
   - **OK**–To apply the configuration and return to the main Configuration page, click **OK**.
   - **Cancel**–To cancel your entries and return to the main page, click **Cancel**.

Configuring Address Sets

To configure an address set in the J-Web user interface:

1. Select **Configure > Security > Policy Elements > Address Book**.
2. Select the **Address Sets** tab.
3. Select one of the following options:
   - **Add**–To create a new address set, click **Add**. The Add Address Set dialog box appears.
   - **Edit**–To edit an existing configuration, select the address set that you want to change and click **Edit**. The Edit Address Set dialog box appears.
   - **Delete**–To delete an existing address set, select it and click **Delete**. (If you select this option, you can skip the remaining steps in this section.)
4. In the **Zone** field, specify the zone to which the address set should apply.
5. In the **Address Set Name** field, specify a name for the address set.
6. Specify which of the preexisting addresses should be included or excluded from the address set by selecting the addresses and using the arrows to move them to the **Out of This Set** and **In This Set** lists.
7. Select one of the following options:

- **OK**–To apply the configuration and return to the main Configuration page, click **OK**.
- **Cancel**–To cancel your entries and return to the main page, click **Cancel**.

Configuring Address Books

This section describes the addresses and address sets configured for the address books of the green zone. The IntranetGREEN zone contains two servers. These servers belong to the same subnet, but they are separately addressable.

This example adds individual addresses for both servers to the zone's address list to accommodate users who have access rights to one server and not the other. It also adds an address set to combine the two servers into a single addressable entity.

To create address books, use either J-Web or the CLI Configuration editor.

This topic covers:

- J-Web Configuration on page 102
- CLI Configuration on page 103
- Related Topics on page 103

J-Web Configuration

To create address books using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. In the Name box, type **IntranetGREEN**.
5. Next to Address Book, click **Configure**.
6. Next to Address, click **Add new entry**.
7. In the Address name box, type **G1 10.1.10.0/24** and click **OK**.
8. Next to Address, click **Add new entry**.

   You need to have more than one address in the Address book to configure an Address Set.
9. In the Address name box, type **G2 10.1.10.0/32** and click **OK**.
10. Next to Address Set, click **Add new entry**.
11. In the Address set name box, type **SerAll** and click **OK.**
12. Expand the Advanced option and do either of the following:
    - Next to Apply groups, click **Add new entry**. Specify the group from which address or address set should inherit the configuration.

- Next to Apply groups except, click **Add new entry**. Specify the group from which address or address set should not inherit the configuration.

13. If you are finished configuring the device, commit the configuration.
14. To check the configuration, see "Verifying Address Book Configuration" on page 103.

CLI Configuration

1. To configure address book entries for the Intranet zone, enter the following commands in edit mode.

```
user@host# set security zones security-zone IntranetGREEN address-book address
   G1 10.1.10.0/24
user@host# set security zones security-zone IntranetGREEN address-book address
   G2 10.1.10.0/32
user@host# set security zones security-zone IntranetGREEN address-book
   address-set SerAll address G1 address G2
```

2. If you are finished configuring the device, commit the configuration.
3. To check the configuration, see "Verifying Address Book Configuration" on page 103.

Related Topics

- Understanding Policy Configuration on page 74
- Policy Application Sets Overview on page 116
- Example: Configuring Applications and Application Sets on page 142
- Configuring Policies on page 80
- Configuring Schedulers on page 107

Verifying Address Book Configuration

Purpose Display information about address books and zones.

Action Use the show security zones CLI command to verify the address book and address set configuration. You get the following output:

```
user@host# show security zones
security-zone green {
    address-book {
        address src_addr3 10.10.10.10/24;
        address src_addr1 64.10.4.44/32;
        address src_addr2 64.10.9.28/32;
        address bbc dns-name www.bbc.com;
        address-set my_source_addresses {
            address src_addr1;
            address src_addr2;
            address src_addr3;
```

```
        }
    }
}
```

What it Means The output displays information about all the addresses configured in an address book in the specified. Verify the following information:

- Configured addresses belong to the correct address book.
- Configured address book belongs to the correct zone.

Chapter 6

Security Policy Schedulers

Schedulers are powerful features that allow a policy to be activated for a specified duration. You can define schedulers for a single (non-recurrent) or recurrent time slot within which a policy is active. You can create schedulers irrespective of a policy, meaning that a scheduler cannot be used by any policies. However, if you want a policy to be active within a scheduled time, then you must first create a scheduler.

When a scheduler times out, the associated policy is deactivated and all sessions associated with the policy are also timed out.

If a policy contains a reference to a scheduler, the schedule determines when the policy is active, that is, when it can be used as a possible match for traffic. Schedulers allow you to restrict access to a resource for a period of time or remove a restriction.

NOTE: A scheduler can have multiple policies associated with it; however, a policy cannot be associated with multiple schedulers.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Configuring a Scheduler—Quick Configuration on page 106
- Configuring Schedulers on page 107
- Associating a Policy to a Scheduler on page 110
- Verifying Scheduled Policies on page 112

Configuring a Scheduler—Quick Configuration

You can use J-Web Quick Configuration to quickly configure a schedule for the security policies.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Configure security zones and interfaces. See "Configuring Security Zones—Quick Configuration" on page 51.
3. Configure address books. See "Configuring Address Books" on page 102.
4. Configure security policies that are going to be associated with a schedule. See "Configuring Policies—Quick Configuration" on page 76.

To configure a scheduler in the J-Web user interface:

1. Select **Configure > Security > Policy Elements > Scheduler**.
2. Select one of the following options:
   - **Add**–To create a new scheduler, click **Add**. The Add new scheduler dialog box appears.
   - **Edit**–To edit an existing configuration, select the scheduler that you want to change and click **Edit**. The Edit scheduler dialog box appears.
   - **Delete**–To delete an existing scheduler, select it and click **Delete**. (If you select this option, you can skip the remaining steps in this section.)
3. In the **Scheduler Name** field, enter a name for the scheduler.
4. In the **Date of start** and **Date of stop** fields, use the calendar pick tool to specify the date range for the scheduler. Then, specify the exact start and stop times using the **HH.MM** fields.
5. If you want the scheduler to run at regular, recurring intervals, select the **Daily** tab or any of the day-of-the-week tabs.

   Within the specified tab, you can fine-tune when the scheduler runs by choosing one of the following options from the **Daily Options** drop-down list:
   - **all-day**—Select this option to run the scheduler throughout the specified day of the week. (This option is not available for the Daily tab.)
   - **exclude**—Select this option to configure the scheduler *not* to run on the specified day of the week. (This option is not available for the Daily tab.)

- **period**—Select this option to configure the scheduler to run only between the specified start and stop time periods. Specify the start and stop times in hh, mm, ss (hour, minute, second) format.
- **none**—Select this option to keep the scheduler "as is" with no further configuration options.

6. Select one of the following options:
   - **OK**–To apply the configuration and return to the main Configuration page, click **OK**.
   - **Cancel**–To cancel your entries and return to the main page, click **Cancel**.

Configuring Schedulers

The following guidelines apply to schedulers:

- A policy is active during the time when the scheduler it refers to is also active.
- When a scheduler is off, the policy is unavailable for policy lookup.
- A scheduler can be configured as one of the following:
  - Scheduler can be active for a single time slot, as specified by a start date and time and a stop date and time.
  - Scheduler can be active forever (recurrent), but as specified by the daily schedule. The schedule on a specific day (time slot) takes priority over the daily schedule.
  - Scheduler can be active within a time slot as specified by the weekday schedule.
  - Scheduler can have a combination of two time slots (daily and timeslot).

To configure a scheduler, you enter a meaningful name and a start and stop time for the scheduler. You can also attach comments.

To configure a scheduler, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 107
- CLI Configuration on page 109
- Related Topics on page 110

J-Web Configuration

To set a schedule that allows a policy to be used for packet match checks using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Schedulers, click **Configure** or **Edit**.

3. Next to Scheduler, click **Add new entry**.
4. In the Scheduler name box, type **Sch1**.
5. Next to the Start Date box, click **Add new entry**.
6. In the Start date box, type **2007-10-01.08:00**.
7. In the Stop date box, type **2008-04-01.21:00**.
8. Specify the days when the policy cannot be used for packet matches:
   a. To exclude Saturdays from the schedule, next to Saturday, click **Configure**.
   b. From the Daily Type list, select **Exclude** and click **OK**.
   c. To exclude Sundays from the schedule, next to Sunday, click **Configure**.
   d. From the Daily Type list, select **Exclude** and click **OK**.
9. Expand the Advanced option and do either of the following:
   - Next to Apply groups, click **Add new entry**. Specify the group to inherit the configuration.
   - Next to Apply groups except, click **Add new entry**. Specify the group not to inherit the configuration.

To associate the schedule with the policy, allowing access during regular work hours:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type **green**.
6. In the To zone name box, type **red**.
7. Next to Policy, click **Add new entry**.
8. In the Policy name box, type **abc**.
9. In the Scheduler name box, type **sch1** and click **OK**.

To set schedulers that allow associated policies to check for packet matches from noon to 6 PM on Saturdays and Sundays:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Schedulers, click **Configure** or **Edit**.
3. In the Scheduler name box, type s**cheduler SatHrs**.
4. Next to Saturday, click **Configure**.
5. From the Daily type list, select **Start time**.
6. Next to Start time, click **Add new entry**.
7. In the Start time value box, type **12:00:00**.

8. In the Stop time box, type **18:00:00** and click **OK**.
9. To set the scheduler that allow associated policies to check for packet matches from noon to 6 PM on Sundays, repeat Step 1 through Step 7 and click **OK**.

To bind these schedules with the policy, allowing access during the specified weekend hours:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. In the From zone name box, type **green**.
5. In the To zone name box, type **red**.
6. In the Policy name box, type **abc**.
7. To bind these schedules to the policy during Saturdays, in the Scheduler name box, type **SatHrs** and click **OK**.
8. To bind these schedules to the policy during Sundays, in the Scheduler name box, type **SunHrs** and click **OK**.
9. If you are finished configuring the device, commit the configuration.
10. To check the configuration, see "Verifying Scheduled Policies" on page 112.

CLI Configuration

1. Use the following commands to set a schedule that allows a policy, which refers to it, to be used for packet match checks from 8 AM to 9 PM all days of the week from October 1, 2007 to April 1, 2008 except Saturdays and Sundays. Otherwise, the policy is inactive.

```
user@host# set schedulers scheduler sch1 start-date 2007-10-01.08:00 stop-date
   2008-04-01.21:00
user@host# set schedulers scheduler sch1 saturday exclude
user@host# set schedulers scheduler sch1 sunday exclude
```

2. Use the following command to associate the schedule to the policy, allowing access during regular work hours, as specified.

```
user@host# set security policies from-zone green to-zone red policy abc
   scheduler-name sch1
```

3. Use the following commands to set schedulers that allow associated policies to check for packet matches from noon to 6 PM on Saturdays and Sundays. Otherwise the policy is inactive.

```
user@host# set schedulers scheduler SatHrs saturday start-time 12:00:00
   stop-time 18:00:00
user@host# set schedulers scheduler SunHrs sunday start-time 12:00 stop-time
   18:00
```

4. Use the following commands to bind these schedules to the policy, allowing access during the specified weekend hours.

```
user@host# set security policies from-zone green to-zone red policy abc
   scheduler-name SatHrs
user@host# set security policies from-zone green to-zone red policy abc
   scheduler-name SunHrs
```

5. If you are finished configuring the device, commit the configuration.
6. You can verify the schedules using the show schedulers CLI command. See "Verifying Scheduled Policies" on page 112.

Related Topics

- Understanding Policy Configuration on page 74
- Associating a Policy to a Scheduler on page 110
- Configuring Address Books on page 102
- Example: Configuring Applications and Application Sets on page 142
- Configuring Policies on page 80

Associating a Policy to a Scheduler

A scheduler is referred by security policies to activate or deactivate a policy according to scheduled times. You can associate a policy with a scheduler as you create the policy.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. For background information, read:
   - Security Policies Overview on page 69
   - Address Books and Address Sets Overview on page 97

To associate a policy to a scheduler, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 110
- CLI Configuration on page 112
- Related Topics on page 112

J-Web Configuration

To associate a policy to a scheduler using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. From the list of configured policies, click the policy you need to associate with a scheduler.
5. Under the Nested Configuration column, click **Policy**.
6. Under the Actions column, click **Edit**.
7. To specify the name of the schedule you wish to associate with the policy, in the **Scheduler name** box, type **SatHrs** and click **OK**.

To set schedulers that allow associated policies to check for packet matches from noon to 6 PM on Saturdays and Sundays:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Schedulers, click **Configure** or **Edit**.
3. Next to Scheduler, click **Add new entry**.
4. In the Scheduler name box, type **scheduler SatHrs**.
5. Next to Saturday, click **Configure**.
6. From the Daily type list, select **Start time**.
7. Next to Start time, click **Add new entry**.
8. In the Start time box, type **12:00:00**.
9. In the Stop time box, type **18:00:00** and click **OK**.
10. To set the scheduler that allow associated policies to check for packet matches from noon to 6 PM on Sundays, repeat Step 1 through Step 7.

To bind these schedules to the policy, allowing access during the specified weekend hours:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. In the From zone name box, type **green**.
5. In the To zone name box, type **red**.
6. In the Policy name box, type **abc**.
7. In the Scheduler name box, type **SatHrs** and click **OK**.
8. In the Scheduler name box, type **SunHrs** and click **OK**.
9. If you are finished configuring the device, commit the configuration.
10. You can verify the schedules using the **show schedulers** CLI command. See “Verifying Scheduled Policies” on page 112.

CLI Configuration

In the following example, you configure schedulers that allow associated policies to be used to check for packet matches from noon to 6 PM on Saturdays and Sundays. Otherwise the policy is inactive.

```
user@host# set schedulers scheduler SatHrs saturday start-time 12:00 stop-time
  18:00
user@host# set schedulers scheduler SunHrs sunday start-time 12:00 stop-time
  18:00
```

The following commands bind these schedules to policy abc allowing access during the specified weekend hours.

```
user@host# set security policies from-zone green to-zone red policy abc scheduler-name
  SatHrs
user@host# set security policies from-zone green to-zone red policy abc scheduler-name
  SunHrs
```

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

You can verify the schedules using the show schedulers CLI command. See "Verifying Scheduled Policies" on page 112.

Related Topics

- Configuring Policies on page 80
- Configuring Schedulers on page 107
- Verifying Scheduled Policies on page 112

Verifying Scheduled Policies

A scheduler is referred by security policies to activate or deactivate a policy according to scheduled times. Use the show schedulers command to verify that your configured policies are associated with the appropriate schedulers. For more information, see the *JUNOS Software CLI Reference*

Purpose Display information about address books and zones.

Action Use the show schedulers CLI command to display information about schedulers configured on the system. If a specific scheduler is identified, detailed information is displayed for that scheduler only.

```
user@host# #show schedulers
scheduler sche1 {
    /* This is sched1 */
    start-date 2006-11-02.12:12 stop-date 2007-11-02.12:11;
}
scheduler sche2 {
    daily {
        all-day;
```

```
    }
    sunday {
        start-time 16:00 stop-time 17:00;
    }
    friday {
        exclude;
    }
}
scheduler sche3 {
    start-date 2006-11-02.12:12 stop-date 2007-11-02.12:11;
    daily {
        start-time 10:00 stop-time 17:00
    }
    sunday {
        start-time 12:00 stop-time 14:00;
        start-time 16:00 stop-time 17:00;
    }
    monday {
        all-day;
    }
    friday {
        exclude;
    }
}
```

What it Means The output displays information about schedulers configured on the system. Verify the following information:

- Daily (recurrent) and one-time only (nonrecurrent) schedulers are configured correctly.
- Schedulers are active if policies are associated.

Chapter 7

Security Policy Applications

Applications are types of traffic for which protocol standards exist. Each application has a transport protocol and destination port number(s) associated with it, such as TCP/port 21 for FTP and TCP/port 23 for Telnet. When you create a policy, you must specify an application for it.

You can select one of the predefined applications from the application book, or a custom application or application set that you created. You can see which application you can use in a policy by using the show application CLI command.

NOTE: Each predefined application has a source port range of 1-65535, which includes the entire set of valid port numbers. This prevents potential attackers from gaining access by using a source port outside of the range. If you need to use a different source port range for any predefined application, create a custom application. For information, see "Understanding Custom Policy Applications" on page 135.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Policy Application Sets Overview on page 116
- Understanding the ICMP Predefined Policy Application on page 117
- Understanding Internet-Related Predefined Policy Applications on page 121
- Understanding Microsoft Predefined Policy Applications on page 123
- Understanding Dynamic Routing Protocols Predefined Policy Applications on page 126
- Understanding Streaming Video Predefined Policy Applications on page 127
- Understanding Sun RPC Predefined Policy Applications on page 128
- Understanding Security and Tunnel Predefined Policy Applications on page 129
- Understanding IP-Related Predefined Policy Applications on page 130
- Understanding Instant Messaging Predefined Policy Applications on page 130
- Understanding Management Predefined Policy Applications on page 131
- Understanding Mail Predefined Policy Applications on page 133
- Understanding UNIX Predefined Policy Applications on page 133

- Understanding Miscellaneous Predefined Policy Applications on page 134
- Understanding Custom Policy Applications on page 135
- Configuring Applications and Application Sets—Quick Configuration on page 136
- Example: Configuring Applications and Application Sets on page 142
- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146
- Example: Defining a Custom Internet Control Message Protocol Application on page 147
- Understanding Policy Application Timeouts on page 149
- Setting a Policy Application Timeout on page 152

Policy Application Sets Overview

When you create a policy, you must specify an application, or service, for it to indicate that the policy applies to traffic of that type. Sometimes the same applications or a subset of them can be present in multiple policies, making it difficult to manage. JUNOS Software allows you to create groups of applications called *application sets*. Application sets simplify the process by allowing you to manage a small number of application sets, rather than a large number of individual application entries.

The application (or application set) is referred to by security policies as match criteria for packets initiating sessions. If the packet matches the application type specified by the policy and all other criteria match, then the policy action is applied to the packet.

Before You Begin

For background information, read:

- Security Policies Overview on page 69
- Understanding Policies on page 70

You can specify the name of an application set in a policy. In this case, if all of the other criteria match, any one of the applications in the application set serves as valid matching criteria; any is the default application name that indicates all possible applications.

Applications are created in the .../applications/application/<application-name> directory. You do not need to configure an application for any of the services that are predefined by the system.

In addition to predefined services, you can configure a custom service. After you create a custom service, you can refer to it in a policy.

Related Topics

- Security Policy Applications on page 115

- Example: Configuring Applications and Application Sets on page 142

Understanding the ICMP Predefined Policy Application

When you create a policy, you can specify the ICMP predefined application for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Internet Control Message Protocol (ICMP) is a part of IP and provides a way to query a network (ICMP query messages) and to receive feedback from the network for error patterns (ICMP error messages). ICMP does not, however, guarantee error message delivery or report all lost datagrams; and it is not a reliable protocol. ICMP codes and type codes describe ICMP query messages and ICMP error messages.

You can choose to permit or deny any or specific types of ICMP messages to improve network security. Some types of ICMP messages can be exploited to gain information about your network that might compromise security. For example, ICMP, TCP, or UDP packets can be constructed to return ICMP error messages that contain information about a network, such as its topology, and access list filtering characteristics. Table 17 on page 117 lists ICMP message names, the corresponding code, type, and description.

Table 17: ICMP Messages

ICMP Message Name	Type	Code	Description
ICMP-ANY	all	all	ICMP-ANY affects any protocol using ICMP. Denying ICMP-ANY impairs any attempt to ping or monitor a network using ICMP. Permitting ICMP-ANY allows all ICMP messages.
ICMP-ADDRESS-MASK ■ Request ■ Reply	17 18	0 0	ICMP address mask query is used for systems that need the local subnet mask from a bootstrap server. Denying ICMP address mask request messages can adversely affect diskless systems. Permitting ICMP address mask request messages might allow others to fingerprint the operating system of a host in your network.

Table 17: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
ICMP-DEST-UNREACH	3	0	ICMP destination unreachable error message indicates that the destination host is configured to reject the packets. Codes 0, 1, 4, or 5 can be from a gateway. Codes 2 or 3 can be from a host (RFC 792). Denying ICMP destination unreachable error messages can remove the assumption that a host is up and running behind a J Series or SRX Series device. Permitting ICMP destination unreachable error messages can allow some assumptions, such as security filtering, to be made about the network.
ICMP Fragment Needed	3	4	ICMP fragmentation error message indicates that fragmentation is needed but the don't fragment flag is set. We recommend denying these messages from the Internet to an internal network.
ICMP FragmentReassembly	11	1	ICMP fragment reassembly time exceeded error indicates that a host reassembling a fragmented message ran out of time and dropped the packet. This message is sometimes sent. We recommend denying these messages from the Internet (external) to the trusted (internal) network.
ICMP-HOST-UNREACH	3	1	ICMP host unreachable error messages indicate that routing table entries do not list or list as infinity a particular host. Sometimes this error is sent by gateways that cannot fragment when a packet requiring fragmentation is received. We recommend denying these messages from the Internet to a trusted network. Permitting these messages allows others to be able to determine your internal hosts IP addresses by a process of elimination or make assumptions about gateways and fragmentation.
ICMP-INFO ■ Request ■ Reply	15 16	0 0	ICMP-INFO query messages allow diskless host systems to query the network and self-configure. Denying ICMP address mask request messages can adversely affect diskless systems. Permitting ICMP address mask request messages might allow others to broadcast information queries to a network segment to determine computer type.

Table 17: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
ICMP-PARAMETER-PROBLEM	12	0	ICMP parameter problem error messages notify you when incorrect header parameters are present and have caused a packet to be discarded We recommend denying these messages from the Internet to a trusted network. Permitting ICMP parameter problem error messages allows others to make assumptions about your network.
ICMP-PORT-UNREACH	3	3	ICMP port unreachable error messages indicate that gateways processing datagrams requesting certain ports are unavailable or unsupported in the network. We recommend denying these messages from the Internet to a trusted network. Permitting ICMP port unreachable error messages can allow others to determine which ports you use for certain protocols.
ICMP-PROTOCOL-UNREACH	3	2	ICMP protocol unreachable error messages indicate that gateways processing datagrams requesting certain protocols are unavailable or unsupported in the network. We recommend denying these messages from the Internet to a trusted network. Permitting ICMP protocol unreachable Error messages can allow others to determine what protocols your network is running.
ICMP-REDIRECT	5	0	ICMP redirect network error messages are sent by a J Series or SRX Series device. We recommend denying these messages from the Internet to a trusted network.
ICMP-REDIRECT-HOST	5	1	ICMP redirect messages indicate datagrams destined for the specified host to be sent along another path.
ICMP-REDIRECT-TOS-HOST	5	3	ICMP redirect type of service (TOS) and host error is a type of message.
ICMP-REDIRECT-TOS-NET	5	2	ICMP redirect TOS and network error is a type of message.

Table 17: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
ICMP-SOURCE-QUENCH	4	0	ICMP source quench error message indicates that a device does not have the buffer space available to accept, queue, and send the packets on to the next hop. Denying these messages will not help or impair internal network performance. Permitting these messages can allow others to know that a device is congested, making it a viable attack target.
ICMP-SOURCE-ROUTE-FAIL	3	5	ICMP source route failed error message We recommend denying these messages from the Internet (external).
ICMP-TIME-EXCEEDED	11	0	ICMP time-to-live (TTL) exceeded error message indicates that a packet's TTL setting reached zero before the packet reached its destination. This ensures that older packets are discarded before resent ones are processed. We recommend denying these messages from a trusted network out to the Internet.
ICMP-TIMESTAMP ■ Request ■ Reply	13 14	0 0	ICMP-TIMESTAMP query messages provide the mechanism to synchronize time and coordinate time distribution in a large, diverse network.
Ping (ICMP ECHO)	8	0	Ping is a utility to determine whether a specific host is accessible by its IP address. Denying ping functionality removes your ability to check to see if a host is active. Permitting ping can allow others to execute a denial-of-service (DoS) or Smurf attack.
ICMP-ECHO-FRAGMENT-ASSEMBLY-EXPIRE	11	1	ICMP fragment echo reassembly time expired error message indicates that the reassembly time was exceeded. We recommend denying these messages.
Traceroute ■ Forward ■ Discard	30 30	0 1	Traceroute is a utility to indicate the path to access a specific host. We recommend denying this utility from the Internet (external) to your trusted network (internal).

Handling ICMP Unreachable Errors

For different levels of security, the default behavior for ICMP unreachable errors from downstream Juniper Networks device is handled as follows:

- Sessions do not close for ICMP type-3 code-4 messages.

  ICMP messages pass through without dropping sessions. Packets are, however, dropped per session.

- Sessions do not close on receiving any kind of ICMP unreachable messages.
- Sessions store ICMP unreachable message, thereby restricting the number of messages flowing through to 1.

  One ICMP unreachable message is generated globally per router. The remaining ICMP unreachable errors are dropped.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Internet-Related Predefined Policy Applications

When you create a policy, you can specify predefined Internet-related applications for the policy.

Before You Begin
For background information, read “Security Policy Applications” on page 115.

Table 18 on page 121 lists Internet-related predefined applications. Depending on your network requirements, you can choose to permit or deny any or all of these applications. Each entry lists the application name, default receiving port, and application description.

Table 18: Predefined Applications

Application Name	Port(s)	Application Description
AOL	5190-5193	America Online Internet service provider (ISP) provides Internet, chat, and instant messaging applications.
DHCP relay	67 (default)	Dynamic Host Configuration Protocol.
DHCP	68 client 67 server	Dynamic Host Configuration Protocol allocates network addresses and delivers configuration parameters from server to hosts.

Table 18: Predefined Applications *(continued)*

Application Name	Port(s)	Application Description
DNS	53	Domain Name System translates domain names into IP addresses.
FTP ■ FTP-Get ■ FTP-Put	20 data 21 control	File Transfer Protocol (FTP) allows the sending and receiving of files between machines. You can choose to deny or permit ANY (GET or PUT) or to selectively permit or deny either GET or PUT. GET receives files from another machine and PUT sends files to another machine. We recommend denying FTP applications from untrusted sources (Internet).
Gopher	70	Gopher organizes and displays Internet servers' contents as a hierarchically structured list of files. We recommend denying Gopher access to avoid exposing your network structure.
HTTP	8080	HyperText Transfer Protocol is the underlying protocol used by the World Wide Web (WWW). Denying HTTP application disables your users from viewing the Internet. Permitting HTTP application allows your trusted hosts to view the Internet.
HTTP-EXT	—	Hypertext Transfer Protocol with extended non-standard ports
HTTPS	443	Hypertext Transfer Protocol with Secure Sockets Layer (SSL) is a protocol for transmitting private documents through the Internet. Denying HTTPS disables your users from shopping on the Internet and from accessing certain online resources that require secure password exchange. Permitting HTTPS allows your trusted hosts to participate in password exchange, shop online, and visit various protected online resources that require user login.
Internet Locator Service	—	Internet Locator Service includes LDAP, User Locator Service, and LDAP over TSL/SSL.
IRC	6665-6669	Internet Relay Chat (IRC) allows people connected to the Internet to join live discussions.
LDAP	389	Lightweight Directory Access Protocol is a set of protocols used to access information directories.
PC-Anywhere	—	PC-Anywhere is a remote control and file transfer software.
TFTP	69	Trivial File transfer Protocol (TFTP) is a protocol for simple file transfer.
WAIS	—	Wide Area Information Server is a program that finds documents on the Internet.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Microsoft Predefined Policy Applications

When you create a policy, you can specify predefined Microsoft applications for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Table 19 on page 123 lists predefined Microsoft applications, parameters associated with each application, and a brief description of each application. Parameters include universal unique identifiers (UUIDs) and TCP/UDP source and destination ports. A UUID is a 128-bit unique number generated from a hardware address, a timestamp, and seed values.

Table 19: Predefined Microsoft Applications

Application	Parameter/UUID	Description
MS-RPC-EPM	135 e1af8308-5d1f-11c9-91a4-08002b14a0fa	Microsoft remote procedure call (RPC) Endpoint Mapper (EPM) Protocol
MS-RPC-ANY	—	Any Microsoft remote procedure call (RPC) applications
MS-AD	4 members	Microsoft Active Directory application Group includes: ■ MS-AD-BR ■ MS-AD-DRSUAPI ■ MS-AD-DSROLE ■ MS-AD-DSSETUP
MS-EXCHANGE	6 members	Microsoft Exchange application group includes: ■ MS-EXCHANGE-DATABASE ■ MS-EXCHANGE-DIRECTORY ■ MS-EXCHANGE-INFO-STORE ■ MS-EXCHANGE-MTA ■ MS-EXCHANGE-STORE ■ MS-EXCHANGE-SYSATD

Table 19: Predefined Microsoft Applications *(continued)*

Application	Parameter/UUID	Description
MS-IIS	6 members	Microsoft IIS Server application group includes: ■ MS-IIS-COM ■ MS-IIS-IMAP4 ■ MS-IIS-INETINFO ■ MS-IIS-NNTP ■ MS-IIS-POP3 ■ MS-IIS-SMTP
MS-AD-BR	ecec0d70-a603-11d0-96b1-00a0c91ece30 16e0cf3a-a604-11d0-96b1-00a0c91ece30	Microsoft Active Directory backup and restore application
MS-AD-DRSUAPI	e3514235-4b06-11d1-ab04-00c04fc2dcd2	Microsoft Active Directory replication
MS-AD-DSROLE	1cbcad78-df0b-4934-b558-87839ea501c9	Microsoft Active Directory DSROLE application
MS-AD-DSSETUP	3919286a-b10c-11d0-9ba8-00c04fd92ef5	Microsoft Active Directory setup application
MS-DTC	906b0ce0-c70b-1067-b317-00dd010662da	Microsoft Distributed Transaction Coordinator application
MS-EXCHANGE-DATABASE	1a190310-bb9c-11cd-90f8-00aa00466520	Microsoft Exchange Database application
MS-EXCHANGE-DIRECTORY	f5cc5a18-4264-101a-8c59-08002b2f8426 f5cc5a7c-4264-101a-8c59-08002b2f8426 f5cc59b4-4264-101a-8c59-08002b2f8426	Microsoft Exchange Directory application
MS-EXCHANGE-INFO-STORE	0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde 1453c42c-0fa6-11d2-a910-00c04f990f3b 10f24e8e-0fa6-11d2-a910-00c04f990f3b 1544f5e0-613c-11d1-93df-00c04fd7bd09	Microsoft Exchange Information Store application
MS-EXCHANGE-MTA	9e8ee830-4459-11ce-979b-00aa005ffebe 38a94e72-a9bc-11d2-8faf-00c04fa378ff	Microsoft Exchange MTA application
MS-EXCHANGE-STORE	99e66040-b032-11d0-97a4-00c04fd6551d 89742ace-a9ed-11cf-9c0c-08002be7ae86 a4f1db00-ca47-1067-b31e-00dd010662da a4f1db00-ca47-1067-b31f-00dd010662da	Microsoft Exchange Store application

Table 19: Predefined Microsoft Applications *(continued)*

Application	Parameter/UUID	Description
MS-EXCHANGE-SYSATD	67df7c70-0f04-11ce-b13f-00aa003bac6c f930c514-1215-11d3-99a5-00a0c9b61b04 83d72bf0-0d89-11ce-b13f-00aa003bac6c 469d6ec0-0d87-11ce-b13f-00aa003bac6c 06ed1d30-d3d3-11cd-b80e-00aa004b9c30	Microsoft Exchange System Attendant application
MS-FRS	f5cc59b4-4264-101a-8c59-08002b2f8426 d049b186-814f-11d1-9a3c-00c04fc9b232 a00c021c-2be2-11d2-b678-0000f87a8f8e	Microsoft File Replication application
MS-IIS-COM	70b51430-b6ca-11d0-b9b9-00a0c922e750 a9e69612-b80d-11d0-b9b9-00a0c922e70	Microsoft Internet Information Server COM GUID/UUID application
MS-IIS-IMAP4	2465e9e0-a873-11d0-930b-00a0c90ab17c	Microsoft Internet Information Server IMAP4 application
MS-IIS-INETINFO	82ad4280-036b-11cf-972c-00aa006887b0	Microsoft Internet Information Server Administration application
MS-IIS-NNTP	4f82f460-0e21-11cf-909e-00805f48a135	Microsoft Internet Information Server NNTP application
MS-IIS-POP3	1be617c0-31a5-11cf-a7d8-00805f48a135	Microsoft Internet Information Server POP3 application
MS-IIS-SMTP	8cfb5d70-31a4-11cf-a7d8-00805f48a135	Microsoft Internet Information Server STMP application
MS-ISMSERV	68dcd486-669e-11d1-ab0c-00c04fc2dcd2 130ceefb-e466-11d1-b78b-00c04fa32883	Microsoft Inter-site Messaging application
MS-MESSENGER	17fdd703-1827-4e34-79d4-24a55c53bb37 5a7b91f8-ff00-11d0-a9b2-00c04fb6e6fc	Microsoft Messenger application
MS-MQQM	fdb3a030-065f-11d1-bb9b-00a024ea5525 76d12b80-3467-11d3-91ff-0090272f9ea3 1088a980-eae5-11d0-8d9b-00a02453c33 5b5b3580-b0e0-11d1-b92d-0060081e87f0 41208ee0-e970-11d1-9b9e-00e02c064c39	Microsoft Windows Message Queue Management application
MS-NETLOGON	12345678-1234-abcd-ef00-01234567cffb	Microsoft Netlogon application

Table 19: Predefined Microsoft Applications *(continued)*

Application	Parameter/UUID	Description
MS-SCHEDULER	1ff70682-0a51-30e8-076d-740be8cee98b 378e52b0-c0a9-11cf-822d-00aa0051e40f 0a74ef1c-41a4-4e06-83ae-dc74fb1cdd53	Microsoft Scheduler application
MS-WIN-DNS	50abc2a4-574d-40b3-9d66-ee4fd5fba076	Microsoft Windows DNS server
MS-WINS	45f52c28-7f9f-101a-b52b-08002b2efabe 811109bf-a4e1-11d1-ab54-00a0c91e9b45	Microsoft WINS application

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Dynamic Routing Protocols Predefined Policy Applications

When you create a policy, you can specify predefined dynamic routing protocol applications for the policy.

Before You Begin
For background information, read “Security Policy Applications” on page 115.

Depending on your network requirements, you can choose to permit or deny messages generated from and packets of these dynamic routing protocols. Table 20 on page 126 lists each supported dynamic routing protocol by name, port, and description.

Table 20: Dynamic Routing Protocols

Dynamic Routing Protocol	Port	Description
RIP	520	RIP is a common distance-vector routing protocol.
OSPF	89	OSPF is a common link-state routing protocol.
BGP	179	BGP is an exterior/interdomain routing protocol.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Streaming Video Predefined Policy Applications

When you create a policy, you can specify predefined streaming video applications for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Table 21 on page 127 lists each supported streaming video application by name and includes the default port and description. Depending on your network requirements, you can choose to permit or deny any or all of these applications.

Table 21: Supported Streaming Video Applications

Application	Port	Description
H.323	TCP source 1-65535; TCP destination 1720, 1503, 389, 522, 1731 UDP source 1-65535; UDP source 1719	H.323 is a standard approved by the International Telecommunication Union (ITU) that defines how audiovisual conference data is transmitted across networks.
NetMeeting	TCP source 1-65535; TCP destination 1720, 1503, 389, 522 UDP source 1719	Microsoft NetMeeting uses TCP to provide teleconferencing (video and audio) applications over the Internet.
Real media	TCP source 1-65535; TCP destination 7070	Real Media is streaming video and audio technology.
RTSP	554	Real-Time Streaming Protocol (RTSP) is for streaming media applications
SIP	5056	Session Initiation Protocol (SIP) is an Application-Layer control protocol for creating, modifying, and terminating sessions.
VDO Live	TCP source 1-65535; TCP destination 7000-7010	VDOLive is a scalable, video streaming technology.

Related Topics

- Example: Adding a Custom Policy Application on page 144

- Example: Modifying a Custom Policy Application on page 146

Understanding Sun RPC Predefined Policy Applications

When you create a policy, you can specify predefined Sun RPC applications for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Table 22 on page 128 lists each Sun remote procedure call Application Layer Gateway (RPC ALG) application name, parameters, and full name.

Table 22: RPC ALG Applications

Application	Program Numbers	Full Name
SUN-RPC-PORTMAPPER	111100000	Sun RPC Portmapper protocol
SUN-RPC-ANY	ANY	Any Sun RPC applications
SUN-RPC-PROGRAM-MOUNTD	100005	Sun RPC Mount Daemon
SUN-RPC-PROGRAM-NFS	100003 100227	Sun RPC Network File System
SUN-RPC-PROGRAM-NLOCKMGR	100021	Sun RPC Network Lock Manager
SUN-RPC-PROGRAM-RQUOTAD	100011	Sun RPC Remote Quota Daemon
SUN-RPC-PROGRAM-RSTATD	100001	Sun RPC Remote Status Daemon
SUN-RPC-PROGRAM-RUSERD	100002	Sun RPC Remote User Daemon
SUN-RPC-PROGRAM-SADMIND	100232	Sun RPC System Administration Daemon
SUN-RPC-PROGRAM-SPRAYD	100012	Sun RPC Spray Daemon
SUN-RPC-PROGRAM-STATUS	100024	Sun RPC Status
SUN-RPC-PROGRAM-WALLD	100008	Sun RPC Wall Daemon
SUN-RPC-PROGRAM-YPBIND	100007	SUN RPC Yellow Page Bind application

Related Topics

- Example: Adding a Custom Policy Application on page 144

- Example: Modifying a Custom Policy Application on page 146

Understanding Security and Tunnel Predefined Policy Applications

When you create a policy, you can specify predefined security and tunnel applications for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Table 23 on page 129 lists each supported application and gives the default port(s) and a description of each entry.

Table 23: Supported Applications

Application	Port	Description
IKE	UDP source 1-65535; UDP destination 500 4500 (used for NAT traversal)	Internet Key protocol (IKE) is a protocol to obtain authenticated keying material for use with ISAKMP. When configuring auto IKE, you can choose from three predefined Phase 1 or Phase 2 proposals: ■ Standard: AES and 3DES ■ Basic: DES and two different types of authentication algorithms ■ Compatible: Four commonly used authentication and encryption algorithms
L2TP	1723	L2TP combines PPTP with Layer 2 Forwarding (L2F) for remote access.
PPTP	1723	Point-to-Point Tunneling Protocol allows corporations to extend their own private network through private *tunnels* over the public Internet.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding IP-Related Predefined Policy Applications

When you create a policy, you can specify predefined IP-related applications for the policy.

Before You Begin
For background information, read “Security Policy Applications” on page 115.

Table 24 on page 130 lists the predefined IP-related applications. Each entry includes the default port and a description of the application.

Table 24: Predefined IP-related Applications

Application	Port	Description
Any	—	Any application
TCP-ANY	1-65535	Any protocol using the TCP TCPMUX port 1
UDP-ANY	137	Any protocol using the UDP

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Instant Messaging Predefined Policy Applications

When you create a policy, you can specify predefined instant messaging applications for the policy.

Before You Begin
For background information, read “Security Policy Applications” on page 115.

Table 25 on page 131 lists predefined Internet-messaging applications. Each entry includes the name of the application, the default or assigned port, and a description of the application.

Table 25: Predefined Internet-messaging Applications

Application	Port	Description
Gnutella	6346 (default)	Gnutella is a public domain file sharing protocol that operates over a distributed network. You can assign any port, but the default is 6346.
MSN	1863	Microsoft Network Messenger is a utility that allows you to send instant messages and talk online.
NNTP	119	Network News Transport Protocol is a protocol used to post, distribute, and retrieve USENET messages.
SMB	445	Server Message Block (SMB) over IP is a protocol that allows you to read and write files to a server on a network.
YMSG	5010	Yahoo! Messenger is a utility that allows you to check when others are online, send instant messages, and talk online.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Management Predefined Policy Applications

When you create a policy, you can specify predefined management applications for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Table 26 on page 131 lists the predefined management applications. Each entry includes the name of the application, the default or assigned port, and a description of the application.

Table 26: Predefined Management Applications

Application	Port	Description
NBNAME	137	NetBIOS Name application displays all NetBIOS name packets sent on UDP port 137.
NDBDS	138	NetBIOS Datagram application, published by IBM, provides connectionless (datagram) applications to PCs connected with a broadcast medium to locate resources, initiate sessions, and terminate sessions. It is unreliable and the packets are not sequenced.

Table 26: Predefined Management Applications *(continued)*

Application	Port	Description
NFS	—	Network File System uses UDP to allow network users to access shared files stored on computers of different types. SUN RPC is a building block of NFS.
NS Global	—	NS-Global is the central management protocol for Juniper Networks Firewall/VPN devices.
NS Global PRO	—	NS Global-PRO is the scalable monitoring system for the Juniper Networks Firewall/VPN device family.
NSM	—	NetScreen-Security Manager
NTP	123	Network Time Protocol provides a way for computers to synchronize to a time reference.
RLOGIN	513	RLOGIN starts a terminal session on a remote host.
RSH	514	RSH executes a shell command on a remote host.
SNMP	161	Simple Network Management Protocol is a set of protocols for managing complex networks.
SQL*Net V1	66	SQL*Net Version 1 is a database language that allows for the creation, access, modification, and protection of data.
SQL*Net V2	66	SQL*Net Version 2 is a database language that allows for the creation, access, modification, and protection of data.
MSSQL	1433 (default instance)	Microsoft SQL is a proprietary database server tool that allows for the creation, access, modification, and protection of data.
SSH	22	SSH is a program to log into another computer over a network through strong authentication and secure communications on an unsecure channel.
SYSLOG	514	Syslog is a UNIX program that sends messages to the system logger.
Talk	517-518	Talk is a visual communication program that copies lines from your terminal to that of another user.
Telnet	23	Telnet is a UNIX program that provides a standard method of interfacing terminal devices and terminal-oriented processes to each other.
WinFrame	—	WinFrame is a technology that allows users on non-Windows machines to run Windows applications.
X-Windows	—	X-Windows is the windowing and graphics system that Motif and OpenLook are based on.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Mail Predefined Policy Applications

When you create a policy, you can specify predefined mail applications for the policy.

Before You Begin
For background information, read “Security Policy Applications” on page 115.

Table 27 on page 133 lists the predefined mail applications. Each includes the name of the application, the default or assigned port number, and a description of the application.

Table 27: Predefined Mail Applications

Application	Port	Description
IMAP	143	Internet Message Access Protocol is used for retrieving messages.
Mail (SMTP)	25	Simple Mail Transfer Protocol is used to send messages between servers.
POP3	110	Post Office Protocol is used for retrieving e-mail.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding UNIX Predefined Policy Applications

When you create a policy, you can specify predefined UNIX applications for the policy.

Before You Begin
For background information, read “Security Policy Applications” on page 115.

Table 28 on page 134 lists the predefined UNIX applications. Each entry includes the name of the application, the default or assigned port, and a description of the application.

Table 28: Predefined UNIX Applications

Application	Port	Description
FINGER	79	Finger is a UNIX program that provides information about the users.
UUCP	117	UNIX-to-UNIX Copy Protocol (UUCP) is a UNIX utility that enables file transfers between two computers over a direct serial or modem connection.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Miscellaneous Predefined Policy Applications

When you create a policy, you can specify miscellaneous predefined applications for the policy.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

Table 29 on page 134 lists predefined miscellaneous applications. Each entry includes the application name, default or assigned port, and a description of the application.

Table 29: Predefined Miscellaneous Applications

Application	Port	Description
CHARGEN	19	Character Generator Protocol is a UDP- or TCP-based debugging and measurement tool.
DISCARD	9	Discard protocol is an Application Layer protocol that describes a process for discarding TCP or UDP data sent to port 9.
IDENT	113	Identification protocol is a TCP/IP Application Layer protocol used for TCP client authentication.
LPR	515 listen; 721-731 source range (inclusive)	Line Printer Daemon protocol is a TCP-based protocol used for printing applications.
RADIUS	1812	Remote Authentication Dial-In User Service application is a server program used for authentication and accounting purposes.
SQLMON	1434 (SQL Monitor Port)	SQL monitor (Microsoft)

Table 29: Predefined Miscellaneous Applications *(continued)*

Application	Port	Description
VNC	5800	Virtual Network Computing facilitates viewing and interacting with another computer or mobile Juniper Networks device connected to the Internet.
WHOIS	43	Network Directory Application Protocol is a way to look up domain names.
IPsec-NAT	—	IPSEC-NAT allows Network Address Translation for ISAKMP and ESP packets.
SCCP	2000	Cisco Station Call Control Protocol (SCCP) uses the signaling connection control port to provide high availability and flow control.
VoIP	—	Voice over IP application group provides voice applications over the Internet and includes H.323 and Session Initiation Protocol (SIP).

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146

Understanding Custom Policy Applications

If you do not want to use predefined applications in your policy, you can easily create custom applications.

Before You Begin
For background information, read "Security Policy Applications" on page 115.

You can assign each custom application the following attributes:

- Name
- Transport protocol
- Source and destination port numbers for applications using TCP or UDP
- Type and code values for applications using ICMP
- Timeout value

This topic covers:

- Custom Application Mappings on page 136
- Related Topics on page 136

Custom Application Mappings

The application option specifies the Layer 7 application that maps to the Layer 4 application that you reference in a policy. A predefined application already has a mapping to a Layer 7 application. However, for custom applications, you must link the application to an application explicitly, especially if you want the policy to apply an Application Layer Gateway (ALG) or deep inspection to the custom application.

NOTE: JUNOS Software supports ALGs for numerous applications, including DNS, FTP, H.323, HTTP, RSH, SIP, Telnet, and TFTP.

Applying an ALG to a custom application involves the following two steps:

- Define a custom application with a name, timeout value, transport protocol, and source and destination ports.
- When configuring a policy, reference that application and the application type for the ALG that you want to apply.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146
- Example: Defining a Custom Internet Control Message Protocol Application on page 147

Configuring Applications and Application Sets—Quick Configuration

You can use J-Web Quick Configuration to quickly configure application and application sets for a security policy. When configuring custom applications, you can configure global settings or individual terms.

Before You Begin

1. For background information, read “Security Policies Overview” on page 69 and “Understanding Policies” on page 70.
2. Configure zones. For more information, see “Configuring Security Zones—Quick Configuration” on page 51.
3. Configure address books. For more information, see “Configuring Addresses and Address Sets—Quick Configuration” on page 100.

This topic includes the following instructions:

- Configuring Global Custom Application Settings on page 137
- Configuring Custom Application Terms on page 141

Configuring Global Custom Application Settings

To configure global application settings in the J-Web user interface:

1. Click **Configure > Security > Policy Elements > Applications**.
2. Select the **Custom-Applications** tab.
3. Select one of the following options:
   - **Add**–To create a new application, click **Add**. The Add an Application dialog box appears.
   - **Edit**–To edit an existing configuration, select the application that you want to change and click **Edit**. The Edit an Application dialog box appears.
   - **Delete**–To delete an existing configuration, select it and click **Delete**. (If you select this option, you can skip the remaining steps in this section.)
4. Select the **Global** tab.
5. In the **Application name** field, specify a name for the custom application.
6. In the **Application-protocol** field, specify the application protocol. For a complete list of options, see Table 30 on page 138.
7. In the **Match IP protocol** field, specify the network protocol. For a complete list of options, see Table 31 on page 138.
8. In the **Destination Port** field, specify the destination port identifier. You can use a numeric value or one of the text synonyms listed in Table 32 on page 139.
9. In the **Source Port** field, specify the source port identifier. You can use a numeric value or one of the text synonyms listed in Table 32 on page 139.
10. In the **Inactivity-timeout** field, specify the length of time (in seconds) that the application is inactive before it times out. For more information about timeouts, see "Understanding Policy Application Timeouts" on page 149.
11. In the **RPC-program-number** field, specify the remote procedure call (RPC) value. Valid values range from 0 to 65535.
12. In the **Match ICMP message code** field, specify the Internet Control Message Protocol (ICMP) code value, such as host-unreachable or host-unreachable-for-tos.
13. In the **Match ICMP message type** field, specify the ICMP packet type value, such as echo or echo-reply.
14. In the **UUID** field, specify a universal unique identifier (UUID). A UUID is a 128-bit unique number generated from a hardware address, a timestamp, and seed values.
15. In the **Application-Set** field, select the application set to which this application should belong.
16. Select one of the following options:
    - **OK**–To apply the configuration and return to the main Configuration page, click **OK**.
    - **Cancel**–To cancel your entries and return to the main page, click **Cancel**.

Table 30: Custom Application Protocols

Protocol	Description
dns	Domain Name Service
ftp	File Transfer Protocol
ignore	Ignore application type
mgcp-ca	Media Gateway Control Protocol with Call Agent
mgcp-ua	MGCP with User Agent
ms-rpc	Microsoft RPC
pptp	Point-to-Point Tunneling Protocol
q931	ISDN connection control protocol (Q.931)
ras	Remote Access Service
realaudio	RealAudio
rsh	UNIX remote shell services
rtsp	Real-Time Streaming Protocol
sccp	Skinny Client Control Protocol
sip	Session Initiation Protocol
sqlnet-v2	Oracle SQLNET v2
sun-rpc	Sun Microsystems RPC
talk	TALK program
tftp	Trivial File Transfer Protocol

Table 31: Match IP Protocols

Protocol	Description
ah	IP Security Authentication Header
egp	Exterior gateway protocol
esp	IPsec Encapsulating Security Payload
gre	Generic routing encapsulation
icmp	Internet Control Message Protocol
igmp	Internet Group Management Protocol

Table 31: Match IP Protocols *(continued)*

Protocol	Description
ipip	IP over IP
ospf	Open Shortest Path First
pim	Protocol Independent Multicast
rsvp	Resource Reservation Protocol
sctp	Stream Control Transmission Protocol
tcp	Transmission Control Protocol
udp	User Datagram Protocol

Table 32: Port Names

Port Name	Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
exec	512
finger	79
ftp	21
ftp-data	20
http	80
https	443

Table 32: Port Names *(continued)*

Port Name	Port Number
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760
kshell	544
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520

Table 32: Port Names *(continued)*

Port Name	Port Number
rkinit	2108
smtp	25
snmp	161
snmp-trap	162
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69
timed	525
who	513
xdmcp	177

Configuring Custom Application Terms

To configure individual application terms in the J-Web user interface:

1. Click **Configure > Security > Policy Elements > Applications**.
2. Select the **Custom-Applications** tab.
3. Select one of the following options:
   - **Add**–To create a new application, click **Add**. The Add an Application dialog box appears.
   - **Edit**–To edit an existing configuration, select the application that you want to change and click **Edit**. The Edit an Application dialog box appears.
   - **Delete**–To delete an existing configuration, select it and click **Delete**. (If you select this option, you can skip the remaining steps in this section.)

4. Select the **Terms** tab.
5. Select one of the following options:
   - **Add**–To create a new application term, click **Add**. The Add new term dialog box appears.
   - **Edit**–To edit an existing configuration term, select the application that you want to change and click **Edit**. The Edit term dialog box appears.
6. In the **Term Name** field, enter a name for the application term.
7. In the **ALG** field, specify the application protocol. For a complete list of options, see Table 30 on page 138.
8. In the **Match IP protocol** field, specify the network protocol. For a complete list of options, see Table 31 on page 138.
9. In the **Destination Port** field, specify the destination port identifier. You can use a numeric value or one of the text synonyms listed in Table 32 on page 139.
10. In the **Source Port** field, specify the source port identifier. You can use a numeric value or one of the text synonyms listed in Table 32 on page 139.
11. In the **Inactivity-timeout** field, specify the length of time (in seconds) that the application is inactive before it times out. For more information about timeouts, see "Understanding Policy Application Timeouts" on page 149.
12. In the **RPC-program-number** field, specify the remote procedure call (RPC) value. Valid values range from 0 to 65535.
13. In the **Match ICMP message code** field, specify the Internet Control Message Protocol (ICMP) code value, such as host-unreachable or host-unreachable-for-tos.
14. In the **Match ICMP message type** field, specify the ICMP packet type value, such as echo or echo-reply.
15. In the **UUID** field, specify a universal unique identifier (UUID). A UUID is a 128-bit unique number generated from a hardware address, a timestamp, and seed values.
16. Select one of the following options:
    - **OK**–To apply the configuration and return to the Add Application or Edit Application dialog box, click **OK**.
    - **Cancel**–To cancel your entries and return to the Add Application or Edit Application dialog box, click **Cancel**.
17. Select one of the following options:
    - **OK**–To apply the configuration and return to the main Configuration page, click **OK**.
    - **Cancel**–To cancel your entries and return to the main page, click **Cancel**.

Example: Configuring Applications and Application Sets

Rather than create or add multiple individual application names to a policy, you can create an application set and refer to the name of the set in a policy. For example,

for a group of employees, you can create an application set that contains all the approved applications.

The following example shows how to create an application set for the three applications that are used to log into the servers in the ABC (intranet) zone, to use the database, and to transfer files.

To configure applications and application sets, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 143
- CLI Configuration on page 144
- Related Topics on page 144

J-Web Configuration

To specify applications and application sets using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. To create an application set for the three applications, next to Applications, click **Configure** or **Edit**.
3. Next to Application set, click **Add new entry**.
4. In the Application set name box, type MgrAppSet.
5. Next to Application, click **Add new entry**.
6. In the Application name box, type ssh and click **OK**.
7. To specify the second application name:
   a. Next to Application, click **Add new entry**.
   b. In the Application name box, type telnet and click **OK**.
8. To specify the third application name:
   a. Next to Application, click **Add new entry**.
   b. In the Application name box, type Cust App and click **OK**.

To create an application set for the applications for e-mail and Web-based applications delivered by the two servers in the external zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Applications, click **Configure** or **Edit**.
3. Next to Application set box, click **Add new entry**.
4. In the Application set name box, type WebMailApps.
5. Use the left and right arrows to add the following applications to the Applications in this set column: **smtp**, **http**, **https**, and **pops**.

6. Click **OK**
7. If you are finished configuring the device, commit the configuration.

CLI Configuration

Managers in zoneA and managers in zoneB use these services. Therefore, the application set is given a generic name such as MgrAppSet

```
user@host# set applications application-set MgrAppSet application ssh
user@host# set applications application-set MgrAppSet application telnet
user@host# set applications application-set MgrAppSet application custApp
```

The following example shows how to create an application set for the applications that are used for e-mail and Web-based applications delivered by the two servers in the external zone.

```
user@host# set applications application-set WebMailApps application smtp
user@host# set applications application-set WebMailApps application http
user@host# set applications application-set WebMailApps application https
user@host# set applications application-set WebMailApps application POPS
```

Related Topics

- Security Policies Overview on page 69
- Policy Application Sets Overview on page 116

Example: Adding a Custom Policy Application

Applications are types of traffic for which protocol standards exist. When you create a policy, you must specify a application for it.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Create addresses. See “Configuring a Policy to Permit Traffic” on page 86.
3. Create zones. See “Creating Security Zones” on page 49.
4. For background information, read:
   - Understanding Custom Policy Applications on page 135
   - Security Policy Applications on page 115

To add a custom application to the application book, you need the following information:

- A name for the application: in this example, cust-telnet.
- A range of source port numbers: 1—65535.

- A range of destination port numbers to receive the application request: for example: 23000—23000.
- Whether the application uses TCP or UDP protocol, or some other protocol as defined by the Internet specifications.

To add a custom application to the application book, use either J-Web or the CLI configuration editor. In this example, the protocol is TCP.

This topic covers:

- J-Web Configuration on page 145
- CLI Configuration on page 145
- Related Topics on page 146

J-Web Configuration

To add a custom application to application book using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Applications, click **Configure** or **Edit**.
3. Next to Application, click **Add new entry**.
4. In the Application name box, type cust-telnet.
5. From the Protocol list, select **tcp**.
6. From the Source port list, select **Enter specific value**.
7. In the Range box, type 1—65535.
8. From the Destination port list, select **Enter specific value**.
9. In the Range box, type 23000—23000 and click **OK**.
10. In the Inactivity-timeout box, type 30 and click **OK**.
11. If you are finished configuring the device, commit the configuration.

CLI Configuration

In this example, the protocol is TCP.

```
user@host# set applications application cust-telnet protocol tcp source-port 1-65535
    destination-port 23000-23000
user@host# set applications application cust-telnet inactivity-timeout 30
```

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

NOTE: The timeout value is in minutes. If you do not set it, the timeout value of a custom application is 180 minutes. If you do not want an application to time out, type never.

Related Topics

- Example: Modifying a Custom Policy Application on page 146
- Example: Defining a Custom Internet Control Message Protocol Application on page 147

Example: Modifying a Custom Policy Application

Applications are types of traffic for which protocol standards exist. When you create a policy, you must specify an application for it.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Create addresses. See "Configuring a Policy to Permit Traffic" on page 86.
3. Create zones. See "Creating Security Zones" on page 49.
4. For background information, read:
   - Understanding Custom Policy Applications on page 135
   - Security Policy Applications on page 115

Use the delete applications *application_name* CLI command to remove the definition of a custom application without removing the application from the application book.

```
user@host# delete applications application cust-telnet
```

To modify the custom policy application, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 146
- CLI Configuration on page 147
- Related Topics on page 147

J-Web Configuration

To modify the custom policy application using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Applications, click **Configure** or **Edit**.
3. To modify the custom policy application, click **Edit**.
4. In the Application name box, type cust-telnet and click **OK**.
5. From the Application protocol list, select **cust-telnet** and click **OK**.
6. From the Protocol list, select **ftp**.

7. From the Source port list, select **Enter specific value**.
8. In the Range box, type 1—65535.
9. From the Destination port list, select **Enter specific value**.
10. In the Range box, type 23230—23230 and click **OK**.
11. If you are finished modifying the custom policy application, commit the configuration.

CLI Configuration

```
user@host# set applications application cust-telnet application-protocol ftp source-port
    1-65535 destination-port 23230-23230
```

If you are finished modifying the custom policy application, commit the configuration.

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Defining a Custom Internet Control Message Protocol Application on page 147

Example: Defining a Custom Internet Control Message Protocol Application

JUNOS Software supports Internet Control Message Protocol (ICMP)—as well as several ICMP messages—as predefined or custom applications. When configuring a custom ICMP application, you must define a type and code.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. For background information, read:
   - Understanding the ICMP Predefined Policy Application on page 117
   - Understanding Custom Policy Applications on page 135
   - Security Policy Applications on page 115

There are different message types within ICMP. For example:

- type 0 = Echo Request message

- type 3 = Destination Unreachable message

NOTE: For more information about ICMP types and codes, refer to RFC 792, *Internet Control Message Protocol*.

An ICMP message type can also have a message code. The code provides more specific information about the message, as shown in Table 33 on page 148.

Table 33: Message Descriptions

Message Type	Message Code
5 = Redirect	0 = Redirect datagram for the network (or subnet)
	1 = Redirect datagram for the host
	2 = Redirect datagram for the type of application and network
	3 = Redirect datagram for the type of application and host
11 = Time Exceeded Codes	0 = Time to live exceeded in transit
	1 = Fragment reassembly time exceeded

JUNOS Software supports any type or code within the 0—255 range.

To define a Custom Internet Control Message Protocol application (ICMP) as the transport protocol, use either J-Web or the CLI configuration editor.

In this example, you define a custom application named " host-unreachable" using ICMP as the transport protocol. The type is 3 (for destination unreachable) and the code is 1 (for host unreachable). You set the timeout value at 4 minutes.

This topic covers:

- J-Web Configuration on page 148
- CLI Configuration on page 149
- Related Topics on page 149

J-Web Configuration

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Applications, click **Configure** or **Edit**.
3. Next to Application, click **Configure** or **Edit**.
4. In the Application name box, type host-unreachable.

5. From the ICMP type list, select **Enter specific value**.
6. In the Number box, type 5.
7. From the ICMP code list, select **Enter specific value**.
8. In the Number box, type 0 and click **OK**.
9. In the Inactivity timeout box, type 4 and click **OK**.
10. If you are finished configuring the device, commit the configuration.

CLI Configuration

```
user@host# set applications application host-unreachable icmp-type 5 icmp-code 0
user@host# set applications application host-unreachable inactivity-timeout 4
```

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

Related Topics

- Example: Adding a Custom Policy Application on page 144
- Example: Modifying a Custom Policy Application on page 146
- Example: Defining a Custom Internet Control Message Protocol Application on page 147

Understanding Policy Application Timeouts

The application timeout value you set for an application determines the session timeout. You can set the timeout threshold for a predefined or custom application; you can use the application default timeout, specify a custom timeout, or use no timeout at all. Application timeout behavior is the same in virtual systems (vsys) security domains as at the root level.

Before You Begin
For background information, read: ■ Security Policy Applications on page 115 ■ Understanding Policies on page 70

This topic covers:

- Application Timeout Configuration and Lookup on page 150
- Contingencies on page 150
- Related Topics on page 152

Application Timeout Configuration and Lookup

Application timeout values are stored in the application entry database and in the corresponding vsys TCP and UDP port-based timeout tables. When you set a application timeout value, JUNOS Software updates these tables with the new value. There are also default timeout values in the applications entry database, which are taken from predefined applications. You can set a timeout, but you cannot alter the default values.

Applications with multiple rule entries share the same timeout value. If multiple applications share the same protocol and destination port range, all applications share the last timeout value configured.

For single application entries, an application timeout lookup proceeds as follows:

1. The specified timeout in the application entry database, if set.
2. The default timeout in the application entry database, if specified in the predefined application.
3. The protocol-based default timeout table. See Table 34 on page 150.

Table 34: Protocol-Based Default Timeout

Protocol	Default Timeout (minutes)
TCP	30
UDP	1
ICMP	1
OSPF	1
Other	30

For application groups, including hidden groups created in multi-cell policy configurations, and for the predefined application ANY (if timeout is not set), application timeout lookup proceeds as follows:

1. The vsys TCP and UDP port-based timeout table, if a timeout is set.
2. The protocol-based default timeout table.

Contingencies

When setting timeouts, be aware of the following contingencies:

- If an application contains several application rule entries, all rule entries share the same timeout. The timeout table is updated for each rule entry that matches the protocol (for UDP and TCP—other protocols use the default). You need to define the application timeout only once. For example, if you create an application with two rules, the following commands will set the timeout to 20 minutes for both rules:

```
user@host# set applications application test protocol tcp destination-port
1035-1035 inactivity-timeout 20

user@host# set applications application test term test protocol udp

user@host# set applications application test term test source-port 1-65535

user@host# set applications application test term test destination-port
1111-1111
```

- If multiple applications are configured with the same protocol and overlapping destination ports, the latest application timeout configured overrides the others in the port-based table. For example:

```
user@host# set applications application ftp-1 protocol tcp source-port
0-65535 destination-port 2121-2121 inactivity-timeout 10

user@host# set applications application telnet-1 protocol tcp source-port
0-65535 designating-port 2100-2148 inactivity-timeout 20
```

  With this configuration, JUNOS Software applies the 20-minute timeout for destination port 2121 in an application group, because the destination port numbers for telnet-1 (2100-2148) overlap those for ftp-1 (2121), and you defined telnet-1 after you defined ftp-1.

  To modify an application timeout when multiple applications use the same protocol and an overlapping destination port range, you must unset the application and reset it with the new timeout value. This is because, during reboot, applications are loaded according to creation time, not modification time.

  To avoid the unintended application of the wrong timeout to an application, do not create applications with overlapping destination port numbers.

- If you unset an application timeout, the default protocol-based timeout in the application entry database is used, and the timeout values in both the application entry and port-based timeout tables are updated with the default value.

  If the modified application has overlapping destination ports with other applications, the default protocol-based timeout might not be the desired value. In that case, reboot JUNOS Software, or set the application timeout again for the desired timeout to take effect.

- When you modify a predefined application and reboot, the modified application might not be the last one in the configuration. This is because predefined applications are loaded before custom applications, and any change made to a custom application, even if made earlier, will show as later than the predefined application change when you reboot.

For example, if suppose you create the following application:

```
user@host# set applications application my-application protocol tcp
destination-port 179-179 inactivity-timeout 20
```

Later you modify the timeout of the predefined application BGP as follows:

```
user@host# set applications application bgp inactivity-timeout 75
```

The BGP application will use the 75-minute timeout value, because it is now written to the application entry database. But the timeout for port 179, the port BGP uses, is also changed to 75 in the TCP port-based timeout table. After you reboot, the BGP application will continue to use the 75-minute timeout which, as a single application, it gets from the application entry database. But the timeout in the TCP port-based table for port 179 will now be 60. You can verify this by entering the show applications application bgp command.

The BGP application has no effect on single applications. But if you add BGP or my_application to an application group, the 60-minute timeout value will be used for destination port 179. This is because application group timeout is taken from the port-based timeout table, if one is set.

To ensure predictability when you modify a predefined application timeout, therefore, you can create a similar application, for example:

```
user@host# set applications application my-bgp protocol tcp
destination-port 179-179 inactivity-timeout 75
```

Related Topics

- Setting a Policy Application Timeout on page 152
- Understanding Policies on page 70
- Policy Application Sets Overview on page 116

Setting a Policy Application Timeout

Application timeout values are stored in the application entry database and in the corresponding vsys TCP and UDP port-based timeout tables. When you set an application timeout value, JUNOS Software updates these tables with the new value.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. For background information, read:
   - Understanding Policy Application Timeouts on page 149
   - Address Books and Address Sets Overview on page 97

In the following example, you change the timeout threshold for the FTP predefined application to 75 minutes.

```
user@host# set applications application ftp inactivity-timeout 75
```

Related Topics

- Understanding Policy Application Timeouts on page 149

Part 4
Application Layer Gateways

- Application Layer Gateways (ALGs) on page 157

Chapter 8

Application Layer Gateways (ALGs)

An Application Layer Gateway (ALG) is a software component that is designed to manage specific protocols such as Session Initiation Protocol (SIP) or File Transfer Protocol (FTP) on J Series Services Routers and SRX Series Services Gateways running JUNOS Software. The ALG intercepts and analyzes the specified traffic, allocates resources, and defines dynamic policies to permit the traffic to pass securely through the Juniper Networks device.

A security policy contains many elements including services and applications. Services are objects that identify application protocols using Layer 4 information, such as standard and accepted TCP and UDP port numbers for application services like Telnet, FTP, SMTP, and HTTP.

The application option specifies the Layer 7 application that maps to a Layer 4 service. A predefined service already has a mapping to a Layer 7 application. However, for custom services, you must link the service to an application explicitly, especially if you want the policy to apply an Application Layer Gateway (ALG).

This chapter describes voice-over-IP (VoIP) ALGs and basic data ALGs. VoIP ALGs provide stateful application layer inspection and Network Address Translation (NAT) capabilities to VoIP signaling and media traffic. The ALG inspects the state of transactions, or calls, and forwards or drops packets based on the those states. The RPC ALG is a data ALG.

ALGs are supported on chassis clusters. For information about chassis clusters, see “Chassis Cluster” on page 867.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding Application Layer Gateways on page 159
- Configuring Application Layer Gateways—Quick Configuration on page 159
- Understanding the H.323 ALG on page 162
- Configuring the H.323 ALG—Quick Configuration on page 164
- Setting H.323 Endpoint Registration Timeout on page 165
- Setting H.323 Media Source Port Range on page 167
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 168
- Allowing Unknown H.323 Message Types on page 169

- Verifying the H.323 Configuration on page 171
- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 172
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 176
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 181
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 183
- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- Configuring the SIP ALG—Quick Configuration on page 193
- Understanding SIP ALG Call Duration and Timeouts on page 196
- Setting SIP Call Duration and Inactive Media Timeout on page 197
- Configuring SIP Denial of Service (DoS) Attack Protection on page 199
- Allowing Unknown SIP Message Types on page 200
- Disabling SIP Call ID Hiding on page 202
- Retaining SIP Hold Resources on page 203
- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213
- Configuring Interface Source NAT for Incoming SIP Calls on page 216
- Configuring a Source NAT Pool for Incoming SIP Calls on page 218
- Configuring Static NAT for Incoming SIP Calls on page 223
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring the SIP Proxy in the Public Zone on page 230
- Configuring a Three-Zone SIP Scenario on page 235
- Verifying the SIP Configuration on page 243
- Understanding the SCCP ALG on page 249
- Configuring the SCCP ALG—Quick Configuration on page 255
- Setting SCCP Inactive Media Timeout on page 257
- Allowing Unknown SCCP Message Types on page 258
- Configuring SCCP Denial of Service (DoS) Attack Protection on page 260
- Configuring Call Manager/TFTP Server in the Private Zone on page 261
- Verifying the SCCP Configuration on page 263
- Understanding the MGCP ALG on page 266
- Configuring the MGCP ALG—Quick Configuration on page 272
- Understanding MGCP ALG Call Duration and Timeouts on page 274
- Setting MGCP Call Duration on page 275
- Setting MGCP Inactive Media Timeout on page 276
- Setting the MGCP Transaction Timeout on page 278
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 279

- Allowing Unknown MGCP Message Types on page 280
- Configuring a Media Gateway in Subscribers' Homes on page 282
- Configuring Three-Zone ISP-Hosted Service Using Source and Static NAT on page 289
- Verifying the MGCP Configuration on page 293
- Understanding the RPC ALG on page 296
- Disabling and Enabling RPC ALG on page 300
- Verifying the RPC ALG Tables on page 301

Understanding Application Layer Gateways

The ALG module is responsible for application-layer aware packet processing. ALG functionality can be triggered either by a service or application configured in the policy. ALGs for packets destined to well-known ports are triggered by service type. When a packet arrives at the Juniper Networks device, the flow module forwards the packet according to the security rule set in the policy. If a policy is found to permit the packet, the associated service type or application type is assigned and a session is created for this type of traffic. If a session is found for the packet, no policy rule match is needed. ALG module is triggered if that particular service or application type requires the supported ALG processing.

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

Related Topics

- Reconnaissance Deterrence Overview on page 763
- Suspicious Packet Attributes Overview on page 791
- Denial-of-Service Attack Overview on page 807

Configuring Application Layer Gateways—Quick Configuration

You can use J-Web Quick Configuration to quickly enable or disable JUNOS Software Application Layer Gateways (ALGs). All ALGs are enabled by default.

Before You Begin
For background information, read “Understanding Application Layer Gateways” on page 159

Figure 14: Quick Configuration Page for General ALGs

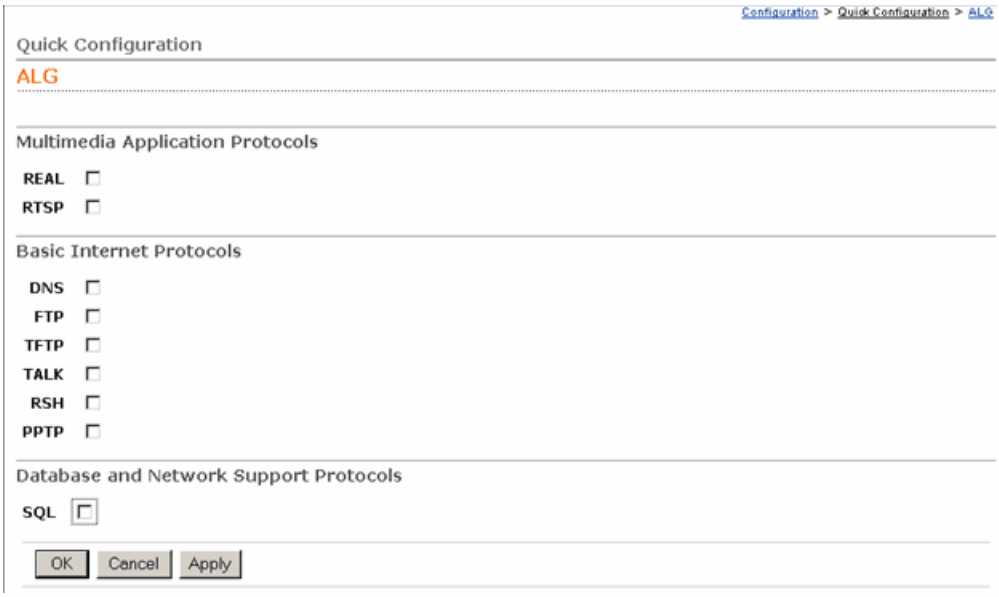

To enable or disable an ALG with Quick Configuration:

1. In the J-Web user interface, select **Configure > Security > ALGs > General ALG**.

   Figure 14 on page 160 shows the General ALG page.

2. Select the check box next to an ALG, described in Table 35 on page 160, then click one of the following buttons:

- To apply the configuration and stay on the Quick Configuration page, click **Apply**.
- To apply the configuration and return to the main Configuration page, click **OK**.
- To cancel your entries and return to the main page, click **Cancel**.

Table 35: General Configuration Options

Field	Function	Action
Multimedia Application Protocols		
RTSP	Provides an ALG for the Real-Time Streaming Protocol.	Select the check box to enable the ALG.
Basic Internet Protocols		
DNS	Provides an ALG for the Domain Name System. The DNS ALG monitors DNS query and reply packets and closes session if the DNS flag indicates the packet is a reply message.	Select the check box to enable the ALG.

Table 35: General Configuration Options *(continued)*

Field	Function	Action
Multimedia Application Protocols		
FTP	Provides an ALG for the File Transfer Protocol. The FTP ALG monitors PORT, PASV and 227 commands. It performs NAT of IP/port in the message and gate opening on the device as necessary. The FTP ALG supports FTP put and FTP get command blocking. When the FTP_NO_PUT or FTP_NO_GET is set in the policy, the FTP ALG sends back a blocking command and closes the associated opened gate when FTP STOR or FTP RETR command is observed.	Select the check box to enable the ALG.
TFTP	Provides an ALG for the Trivial File Transfer Protocol. The TFTP ALG processes TFTP packet that initiate the request and opens a gate to allow return packets from the reverse direction to the port that sends the request.	Select the check box to enable the ALG.
TALK	Provides an ALG for the TALK Protocol. The TALK protocol uses UDP port 517 and port 518 for control channel connections. The talk program consists of a server and a client. The server handles client notifications and helps to establish talk sessions. There are two types of talk servers: ntalk and talkd. The TALK ALG processes packets of both ntalk and talkd formats. It also performs NAT and gate opening as necessary.	Select the check box to enable the ALG.
RSH	Provides an ALG for the Provides an ALG for the Remote Shell. The RSH ALG handles TCP packets destined for port 514 and process the RSH port command. The RSH ALG performs NAT on the port in the port command and opens gates as necessary.	Select the check box to enable the ALG.
PPTP	Provides an ALG for the Point-to-Point Tunneling Protocol. The PPTP is a layer 2 protocol that tunnels PPP data across TCP/IP networks. The PPTP client is freely available on Windows systems and is widely deployed for building Virtual Private Networks (VPNs).	Select the check box to enable the ALG.
Database and Network Support Protocols		

Table 35: General Configuration Options *(continued)*

Field	Function	Action
Multimedia Application Protocols		
SQL	Provides an ALG for the Structured Query Language. The SQLNET ALG processes SQL TNS response frame from the server side. It parses the packet and looks for (HOST = ipaddress), (PORT = port) pattern and performs NAT and gate opening on the client side for the TCP data channel.	Select the check box to enable the ALG.

Understanding the H.323 ALG

The H.323 standard is a legacy VoIP protocol defined by the International Telecommunication Union (ITU-T). H.323 consists of a suite of protocols (such as H.225.0 and H.245) that are used for call signaling and call control for VoIP.

Before You Begin
For background information, read “Application Layer Gateways (ALGs)” on page 157.

H.323 uses the ASN.1 coding format. It sets up the dynamic links for data, video, and audio streams, following the protocols Q.931 (with port number 1720) and H.245. There are three major processes in H.323:

- Gatekeeper Discovery—An endpoint finds its gatekeeper through the gatekeeper discovery process, through broadcast or unicast (to a known IP and the well-known UDP port 1719). (JUNOS Software supports unicast only.)
- Endpoint Registration, Admission, and Status—An endpoint registers to a gatekeeper and asks for its management. Before making a call, an endpoint asks its gatekeeper for permission to place the call. In both registration and admission phases, the remote access server (RAS) channel is used. The Transport Service Access Point (TSAP) may be either the well-known UDP port (1719), or a dynamically assigned port from the discovery or registration phase.
- Call Control and Call Setup—Calls can be established within a zone or across two zones, or even across multiple zones (multipoint conference). The call setup and teardown is performed through the call signaling channel whose TSAP is the well-known TCP port (1720). The call control, including opening/closing media channels between two endpoints, is performed through the call control channel whose TSAP is dynamically assigned from the previous call signaling process. H.245 messages are used in the call control channel, and are encoded using ASN.1.

NOTE: Detailed information on H.323 can be found in ITU-T RECOMMENDATION H.323.

The H.323 ALG lets you secure VoIP communication between terminal hosts, such as IP phones and multimedia devices. In such a telephony system, gatekeeper J series device manage call registration, admission, and call status for VoIP calls. Gatekeepers can reside in the two different zones or in the same zone. (See Figure 15 on page 163.)

Figure 15: H.323 Protocol

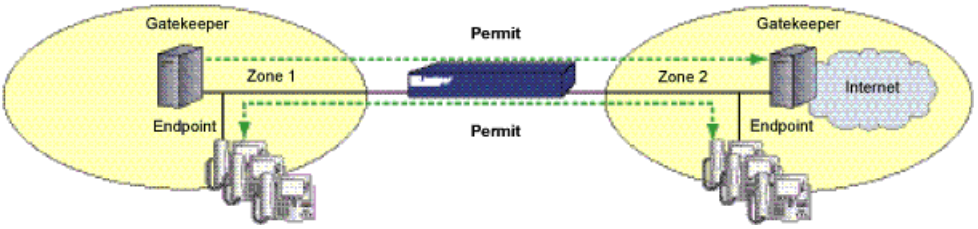

NOTE: The illustrations use IP phones for illustrative purposes, although it is possible to make configurations for other hosts that use VoIP, such as Microsoft NetMeeting multimedia devices.

Related Topics

- Understanding the SIP ALG on page 187
- Understanding the SCCP ALG on page 249
- Understanding the MGCP ALG on page 266
- Setting H.323 Endpoint Registration Timeout on page 165
- Setting H.323 Media Source Port Range on page 167
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 168
- Allowing Unknown H.323 Message Types on page 169
- Verifying the H.323 Configuration on page 171
- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 172
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 176
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 183
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 181

Configuring the H.323 ALG—Quick Configuration

You can use J-Web Quick Configuration to quickly configure H.323 ALG Parameters

Before You Begin
For background information, read “Understanding the H.323 ALG” on page 162

Figure 16: H.323 ALG Configurationy

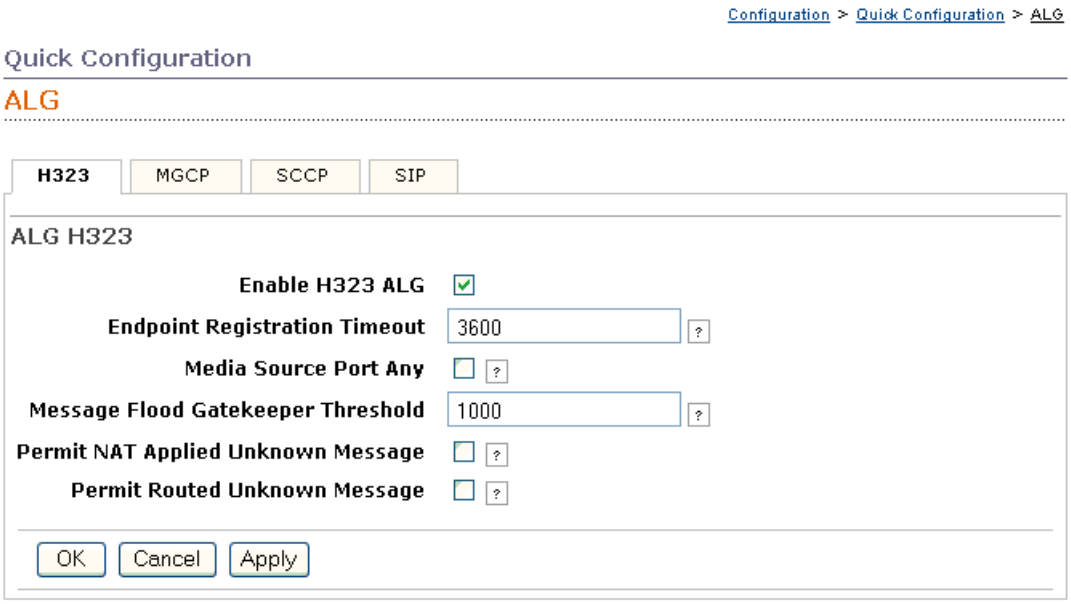

To configure the H.323 ALG with Quick Configuration:

1. In the J-Web user interface, select **Configure > Security > ALGs > VoIP ALG**.

   Figure 16 on page 164 shows the H.323 ALG page.

2. Fill in the parameter settings as described in Table 36 on page 164 and click one of the following buttons:

   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.

Table 36: H.323 Configuration Options

Enable H323 ALG	Enable or disable the H.323 ALG	Click the check box.

Table 36: H.323 Configuration Options *(continued)*

Endpoint Registration Timeout	Controls how long entries remain in the NAT table.	Enter a value between 10 and 50,000 seconds.
Media Source Port Any	Allows media traffic from any port number. By default, this feature is disabled. When disabled, the device allows a temporary opening, or pinhole, in the firewall as needed for media traffic.	Enter a value between 1 and 50,000 seconds.
Message Flood Gatekeeper Threshold	Limits the rate per second at which remote access server (RAS) requests to the gatekeeper are processed. Messages exceeding the threshold are dropped. This feature is disabled by default.	Enter a value
Permit NAT Applied Unknown Message	Specifies how unidentified H.323 messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown H.323 (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped. This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.	Click the check box.
Permit Routed Unknown Message	Specifies that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)	Click the check box.

Setting H.323 Endpoint Registration Timeout

You set an endpoint registration timeout to specify how long an endpoint registration entry remains in the NAT table

Before You Begin

For background information, read “Understanding the H.323 ALG” on page 162.

In NAT mode, when endpoints in the protected network behind the J Series device register with the H.323 gatekeeper, the device adds an entry to the NAT table containing a mapping of the public-to-private address for each endpoint. These entries make it possible for endpoints in the protected network to receive incoming calls. To ensure uninterrupted incoming call service, set the endpoint registration timeout to a value equal to or greater than the keepalive value the administrator configures

on the gatekeeper. The range is 10 to 50000 seconds, the default value is 3600 seconds.

To set the H.323 endpoint registration, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 166
- CLI Configuration on page 166
- Related Topics on page 166

J-Web Configuration

To specify endpoint registration timeout using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Alg, click **Configure** or **Edit**.
4. Select the H.323 check box if it is not selected and click **Configure** or **Edit**.
5. In the Endpoint-registration-timeout box, type 5000 and click **OK**.
6. If you are finished configuring the device, commit the configuration.
7. To check the configuration, see "Verifying the H.323 Configuration" on page 171.

CLI Configuration

In this example, you set the endpoint registration timeout to 5,000 seconds.

```
user@host# set security alg h323 endpoint-registration-timeout 5000
```

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

To check the configuration, see "Verifying the H.323 Configuration" on page 171.

Related Topics

- Allowing Unknown H.323 Message Types on page 169
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 168
- Setting H.323 Media Source Port Range on page 167
- Verifying the H.323 Configuration on page 171

Setting H.323 Media Source Port Range

The media source port feature enables you to configure the device to allow media traffic on a narrow or wide range of ports.

Before You Begin
For background information, read "Understanding the H.323 ALG" on page 162.

By default, the device listens for H.323 traffic on a wide range of ports. If your endpoint equipment allows you to specify a sending port and a listening port, you might want to narrow the range of ports the device allows media traffic on. This enhances security by opening a smaller pinhole for H.323 traffic.

This topic covers:

- J-Web Configuration on page 167
- CLI Configuration on page 167
- Related Topics on page 168

J-Web Configuration

To configure the device to open a narrow gate for media traffic by disabling the media source port feature using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Alg, click **Configure** or **Edit**.
4. Select the H.323 check box if it is not selected and click **Configure** or **Edit**.
5. Select the **Media source port any** check box and click **OK**.
6. If you are finished configuring the device, commit the configuration.
7. To check the configuration, see "Verifying the H.323 Configuration" on page 171.

CLI Configuration

In this example, you configure the device to open a narrow gate for media traffic by disabling the media source port feature.

```
user@host# set security alg h323 media-source-port-any disable
```

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

To check the configuration, see "Verifying the H.323 Configuration" on page 171.

Related Topics

- Setting H.323 Endpoint Registration Timeout on page 165
- Allowing Unknown H.323 Message Types on page 169
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 168
- Verifying the H.323 Configuration on page 171

Configuring H.323 Denial of Service (DoS) Attack Protection

You can protect the H.323 gatekeeper from flood attacks by limiting the number of remote access service (RAS) messages per second it will attempt to process.

Before You Begin
For background information, read "Understanding the H.323 ALG" on page 162.

Incoming RAS request messages exceeding the threshold you specify are dropped by H.323 ALG. The range is 1 to 50000 messages per second, the default value is 1000.

To configure the H.323 DoS attack protection feature, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 168
- CLI Configuration on page 169
- Related Topics on page 169

J-Web Configuration

To limit the number of incoming RAS request messages to the H.323 gatekeeper to 5,000 messages per second using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Alg, click **Configure** or **Edit**.
4. Select the **H.323** check box if it is not selected and click **Configure** or **Edit**.
5. Next to Application Screen, click **Configure** or **Edit**.
6. Next to Message flood, click **Configure** or **Edit**.
7. To specify the gatekeeper threshold, in the Threshold box, type 5000 and click **OK**.

8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see "Verifying the H.323 Configuration" on page 171.

CLI Configuration

In this example, you limit the number of incoming RAS request messages to the H.323 gatekeeper to 5,000 messages per second.

```
user@host# set security alg h323 application-screen message-flood message-flood
  gatekeeper threshold 5000
```

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

To check the configuration, see "Verifying the H.323 Configuration" on page 171.

Related Topics

- Setting H.323 Endpoint Registration Timeout on page 165
- Setting H.323 Media Source Port Range on page 167
- Allowing Unknown H.323 Message Types on page 169
- Verifying the H.323 Configuration on page 171

Allowing Unknown H.323 Message Types

To accommodate on-going development of the H.323 protocol, you might want to allow traffic containing new H.323 message types. The unknown H.323 message type feature enables you to configure the device to accept H.323 traffic containing unknown message types in both NAT and route modes.

Before You Begin
For background information, read "Understanding the H.323 ALG" on page 162.

This feature enables you to specify how unidentified H.323 messages are handled by the device. The default is to drop unknown (unsupported) messages. We do not recommend permitting unknown messages because they can compromise. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown H.323 messages can help you get your network operational, so that you can analyze your VoIP traffic to determine why some messages were being dropped.

Note that this command applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

- permit-nat-applied—Specifies that unknown messages be allowed to pass if the session is in NAT mode.
- permit-routed—Specifies that unknown messages be allowed to pass if the session is in route mode.

To configure the allow unknown messages feature, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 170
- CLI Configuration on page 170
- Related Topics on page 171

J-Web Configuration

To configure the device to allow unknown H.323 message types in both route and NAT modes using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Alg, click **Configure** or **Edit**.
4. Select the **H.323** check box if it is not selected and click **Configure** or **Edit**.
5. Next to Application screen, click **Configure** or **Edit**.
6. Next to Unknown message, click **Configure** or **Edit**.
7. Select one of the following boxes:
   - To allow unknown message types in NAT mode, select the Permit nat applied check box and click **OK**.
   - To allow unknown message types in route mode, select the Permit route check box and click **OK**.
8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see "Verifying the H.323 Configuration" on page 171.

CLI Configuration

In this example, you configure the device to allow unknown H.323 message types in both route and NAT modes.

```
user@host# set security alg h323 application-screen unknown-message
  permit-nat-applied permit-routed
```

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

To check the configuration, see "Verifying the H.323 Configuration" on page 171.

Related Topics

- Setting H.323 Endpoint Registration Timeout on page 165
- Setting H.323 Media Source Port Range on page 167
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 168
- Verifying the H.323 Configuration on page 171

Verifying the H.323 Configuration

To verify the H.323 configuration, perform this task.

Verifying H.323 Counters

Purpose Display information about active calls.

Action From the J-Web interface, select **Monitor > ALGs > H323.** Alternatively, from the CLI, enter the show security alg sip calls command.

```
user@host> show security alg h323 counters
H.323 counters summary:
  Packets received         : 0
  Packets dropped          : 0
  RAS message received     : 0
  Q.931 message received   : 0
  H.245 message received   : 0
  Number of calls          : 0
  Number of active calls   : 0
H.323 error counters:
  Decoding errors          : 0
  Message flood dropped    : 0
  NAT errors               : 0
  Resource manager errors  : 0
H.323 message counters:
  RRQ       : 0
  RCF       : 0
  ARQ       : 0
  ACF       : 0
  URQ       : 0
  UCF       : 0
  DRQ       : 0
  DCF       : 0
  Oth RAS   : 0
  Setup     : 0
  Alert     : 0
  Connect   : 0
  CallProd  : 0
  Info      : 0
  RelCmpl   : 0
  Facility  : 0
  Empty     : 0
  OLC       : 0
```

```
  OLC-ACK   : 0
  Oth H245  : 0
```

What it Means The output provides counts at the packet, message, and call levels. Verify the following information:

- Number of packets received and dropped
- Number of remote access service (RAS), Q.931, and H.245 messages received
- Number of calls and active calls
- Number of all types of H.323 messages

Note that H.323 counters for calls and active calls in the output to this show security command do not apply to the proprietary Avaya implementation of H.323. This is because Q.931 setup and connect messages are exchanged right after the phone is powered up and call creation and tear down is done by Facility messages.

Note also that counters for calls and active calls are increased when the resources allocated for calls are increased—that is, messages belonging to the same call and that pass the firewall multiple times increment the counters. This applies when resources for a call need to be allocated multiple times. For example, in a two-zone scenario the setup and connect message pair allocates one call resource, and the active call counter is increased by one. But in a three-zone scenario the setup and connect message pair passes the firewall twice, each time allocating different call resources. In this case, the counter is incremented.

Counters for H.245 messages received also will not be accurate in the case of H.245 tunneling. Because H.245 messages are encapsulated in Q.931 packets, the counter for H.245 messages received will remain zero even when there are H.245 messages. The Other H245 counter will, however, reflect these packet transmissions.

Related Topics

- Understanding the H.323 ALG on page 162

Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone

In the following example, you set up two policies that allow H.323 traffic to pass between IP phone hosts and a gatekeeper in the private zone, and an IP phone host (2.2.2.5) in the public zone.

Before You Begin
For background information, read “Understanding the H.323 ALG” on page 162.

In this example, the device can be in either route or NAT mode. See Figure 17 on page 173.

Figure 17: H.323 Gatekeeper in Zone1

To configure a gatekeeper in the internal zone, use either J-Web or the CLI Configuration editor.

This topic covers:

- J-Web Configuration on page 173
- CLI Configuration on page 175
- Related Topics on page 175

J-Web Configuration

To configure an address book, to configure a policy from the internal zone to the external zone, and to configure policies from the external zone to the internal zone using the J-Web configuration editor, follow the sequence of steps listed below:

To configure an address book:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone.
4. Next to Security zones, click **Add new entry**.
5. In the Name box, type public.
6. Next to Address book, click **Configure** or **Edit**.
7. Next to Address, click **Add new entry**.
8. In the Address name box, type ip_phone 2.2.2.5/32 and click **OK**.

To configure a policy from the internal zone to the external zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry.**
5. In the From-zone name box, type **private.**
6. In the To-zone name box, type public and click **OK**.
7. Under the From zone name column, click **private**.
8. Next to Policy, click **Add new entry**.

9. In the Policy name box, type p1.
10. Select the **Match** check box.
11. Select the **Then** check box.
12. Next to Match, click **Configure**.
13. From the Source address list, select **Source address.**
14. Next to Source address, click **Add new entry**.
15. From the Value keyword list, select **any** and click **OK**.
16. From the Destination address choice list, select **Destination address**.
17. Next to destination address, click **Add new entry**.
18. From the Value keyword list, select **Enter Specific Value.**
19. In the Address box, type ip_phone and click **OK**.
20. From the Application choice list, select **Application**.
21. Next to Application, click **Add new entry**.
22. In the Value keyword box, type junos-h323 and click **OK**.
23. Next to Then, click **Configure**.
24. Next to Action, select **permit** and click **OK.**

To configure policies from the external zone to the internal zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Under From zone name column, click **private**.
5. Next to Policy, click **Add new entry**.
6. In the From zone name box, type private.
7. In the To zone name box, type public and click **OK**.
8. Under the From zone name column, click **private**.
9. Next to Policy, click **Add new entry**.
10. In the Policy name box, type p2.
11. Select the **Match** check box.
12. Select the **Then** check box.
13. Next to Match, click **Configure**.
14. From the Source address list, select **Source address.**
15. Next to Source address, click **Add new entry**.
16. From the Value keyword list, select **any** and click **OK**.
17. Next to Destination address, click **Add new entry**.

18. From the Value keyword list, select **Enter Specific Value.**
19. In the Address box, type ip_phone and click **OK**.
20. Next to Application, click **Add new entry**.
21. In the Value keyword box, type junos-h323 and click **OK**.
22. Next to Then, click **Configure**.
23. From the Action list, select **permit** and click **OK.**
24. If you are finished configuring the device, commit the configuration.

CLI Configuration

To configure an address book, to configure a policy from the internal zone to the external zone, and to configure policies from the external zone to the internal zone, follow the sequence of steps listed below:

1. Configure an address book.

```
user@host# set security zones security-zone public address-book address ip_phone
    2.2.2.5/32
```

2. Configure a policy from the internal zone to the external zone.

```
user@host# set security policies from-zone private to-zone public policy p1 match
    source-address any
user@host# set security policies from-zone private to-zone public policy p1 match
    destination-address ip_phone
user@host# set security policies from-zone private to-zone public policy p1 match
    application junos-h323
user@host# set security policies from-zone private to-zone public policy p1 then
    permit
```

3. Configure policies from the external zone to the internal zone.

```
user@host# set security policies from-zone public to-zone private policy p2 match
    source-address any
user@host# set security policies from-zone public to-zone private policy p2 match
    destination-address ip_phone
user@host# set security policies from-zone public to-zone private policy p2 match
    application junos-h323
user@host# set security policies from-zone public to-zone private policy p2 then
    permit
```

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

Related Topics

- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 176
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 181
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 183
- Understanding the SIP ALG on page 187

- Setting SCCP Inactive Media Timeout on page 257
- Configuring a Media Gateway in Subscribers' Homes on page 282

Passing H.323 ALG Traffic to a Gatekeeper in the External Zone

Because route mode does not require address mapping of any kind, a device configuration for a gatekeeper in the external, or public, zone is usually identical to the configuration for a gatekeeper in an internal, or private, zone.

Before You Begin
For background information, read "Understanding the H.323 ALG" on page 162.

In the following example, you set up two policies to allow H.323 traffic to pass between IP phone hosts in the internal zone, and the IP phone at IP address 2.2.2.5 (and the gatekeeper) in the external zone. The device can be in transparent or route mode. See Figure 18 on page 176.

Figure 18: H.323 Gatekeeper in Zone2

To configure a gatekeeper in the external zone, use either J-Web or the CLI Configuration editor.

This topic covers:

- J-Web Configuration on page 176
- CLI Configuration on page 180
- Related Topics on page 181

J-Web Configuration

To configure addresses, to configure a policy from the internal zone to the external zone, and to configure a policy to allow traffic between the internal zone and the gatekeeper in the external zone using the J-Web configuration editor, follow the sequence of steps listed below:

To configure addresses:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.

4. Next to Security zone, click **Add new entry**.
5. In the Name box, type external.
6. Next to Address book, click **Configure** or **Edit**.
7. Next to Address, click **Add new entry**.
8. In the Address name box, type ip_phone 2.2.2.5/32 and click **OK**.
9. To configure another security zone internal, repeat Step 2 through Step 4 and click **OK**.
10. Next to Address book, click **Configure**.
11. Next to Address, click **Add new entry**.
12. In the Address name box, type gatekeeper 2.2.2.10/32 and click **OK**.

To configure a policy from the internal zone to the external zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type internal.
6. In the To zone name box, type external and click **OK**.
7. Under the From zone name column, click **internal**.
8. Next to Policy, click **Add new entry**.
9. In the Name box, type p_1.
10. Select the **Match** check box.
11. Select the **Then** check box.
12. Next to Match, click **Configure** or **Edit**.
13. Next to Source address, select **Enter Source address** and click **Add new entry**.
14. From the Value keyword list, select **any** and click **OK**.
15. Next to Destination address, click **Add new entry**.
16. From the Value keyword list, select **Enter Specific Value.**
17. In the Address box, type ip_phone and click **OK**.
18. Next to Application, click **Add new entry**.
19. In the Value keyword box, type junos-h323 and click **OK**.
20. Next to Then, click **Configure**.
21. Next to Action, enter permit and click **OK.**

To configure a policy to allow traffic between the internal zone and the gatekeeper in the external zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type internal.
6. In the To zone name box, type external and click **OK**.
7. Next to Policy, click **Add new entry**.
8. In the Policy name box, type p_2.
9. Select the **Match** check box.
10. Select the **Then** check box.
11. Next to Match, click **Configure**.
12. Next to Source address, select **Enter Source address** and click **Add new entry**.
13. From the Value keyword list, select **any** and click **OK**.
14. Next to Destination address, click **Add new entry**.
15. From the Value keyword list, select **Enter Specific Value.**
16. Next to Address box, type gatekeeper and click **OK**.
17. Next to Application, click **Add new entry**.
18. In the Value keyword box, type junos-h323 and click **OK**
19. Next to Then, click **Configure**.
20. Next to Action, enter permit and click **OK.**

To configure a policy to allow traffic between phones in the internal zone and the external zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, **Add new entry**.
5. In the From zone name box, type external.
6. In the To zone name box, type internal and click **OK**.
7. Under the From zone name column, click **external**.
8. Next to Policy, click **Add new entry**.
9. In the Policy name box, type p_3.
10. Select the **Match** check box.
11. Select the **Then** check box.
12. Next to Match, click **Configure**.
13. Next to Source address, select **Enter Source address** and click **Add new entry**.

14. From the Value keyword list, select **ip_phone** and click **OK**.
15. Next to Destination address, click **Add new entry**.
16. From the Value keyword list, select **Enter Specific Value.**
17. In the Address box, type any and click **OK**.
18. Next to Application, click **Add new entry**.
19. In the Value keyword box, type junos-h323 and click **OK**
20. Next to Then, click **Configure**.
21. Next to Action, enter permit and click **OK.**

To configure a policy to allow traffic between phones in the internal zone and the gatekeeper in the external zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, **Add new entry**.
5. In the From zone name box, type external.
6. In the To zone name box, type internal and click **OK**.
7. Next to Policy, click **Add new entry**.
8. Under the From zone name column, click **external**.
9. In the Policy name box, type p_4.
10. Select the **Match** check box.
11. Select the **Then** check box.
12. Next to Match, click **Configure**.
13. Next to Source address, select **Enter Source address** and click **Add new entry**.
14. From the Value keyword list, select **gatekeeper** and click **OK**.
15. Next to Destination address, click **Add new entry**.
16. From the Value keyword list, select **Enter Specific Value.**
17. In the Address box, type any and click **OK**.
18. Next to Application, click **Add new entry**.
19. In the Value keyword box, type junos-h323 and click **OK**
20. Next to Then, click **Configure**.
21. Next to Action, enter permit and click **OK.**
22. If you are finished configuring the device, commit the configuration.

CLI Configuration

1. Configure addresses.

```
user@host# set security zones security-zone external address-book address
  IP_Phone 2.2.2.5/32
user@host# set security zones security-zone internal address-book address
  gatekeeper 2.2.2.10/32
```

2. Configure a policy from the internal zone to the external zone.

```
user@host# set security policies from-zone internal to-zone external policy p_1
  match source-address any
user@host# set security policies from-zone internal to-zone external policy p_1
  match destination-address IP_Phone
user@host# set security policies from-zone internal to-zone external policy p_1
  match application junos-h323
user@host# set security policies from-zone internal to-zone external policy p_1
  then permit
```

3. Configure a policy to allow traffic between the internal zone and the gatekeeper in the external zone.

```
user@host# set security policies from-zone internal to-zone external p_2 match
  source-address any
user@host# set security policies from-zone internal to-zone external policy p_2
  match destination-address gatekeeper
user@host# set security policies from-zone internal to-zone external policy p_2
  match application junos-h323
user@host# set security policies from-zone internal to-zone external policy p_2
  then permit
```

4. Configure a policy to allow traffic between phones in the internal zone and the external zone.

```
user@host# set security policies from-zone external to-zone internal policy p_3
  match source-address IP_Phone
user@host# set security policies from-zone external to-zone internal policy p_3
  match destination-address any
user@host# set security policies from-zone external to-zone internal policy p_3
  match application junos-h323
user@host# set security policies from-zone external to-zone internal policy p_3
  then permit
```

5. Configure a policy to allow traffic between phones in the internal zone and the gatekeeper in the external zone.

```
user@host# set security policies from-zone external to-zone internal policy id_4
  match source-address gatekeeper
user@host# set security policies from-zone external to-zone internal policy p_4
  match destination-address any
user@host# set security policies from-zone external to-zone internal policy p_4
  match application junos-h323
```

```
user@host# set security policies from-zone external to-zone internal policy p_4
    then permit
```

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

Related Topics

- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 172
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 183
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 181
- Understanding the SIP ALG on page 187
- Setting SCCP Inactive Media Timeout on page 257
- Configuring a Media Gateway in Subscribers' Homes on page 282

Using NAT and the H.323 ALG to Enable Outgoing Calls

In this example, the devices in the external zone include the endpoint host (10.1.1.5) and the gatekeeper (10.1.1.25). IP_Phone2 (2.2.2.5) is in Zone2. You configure the device to allow traffic between the endpoint host IP_Phone1 and the gatekeeper in the external zone and the endpoint host IP_Phone2 in the internal zone.

Before You Begin
For background information, read “Understanding the H.323 ALG” on page 162.

When the J Series device uses NAT, a gatekeeper or endpoint device in the external zone has a private address, and when it is in the internal zone, it has a public address. See Figure 19 on page 181.

Figure 19: Network Address Translation—Outgoing Calls

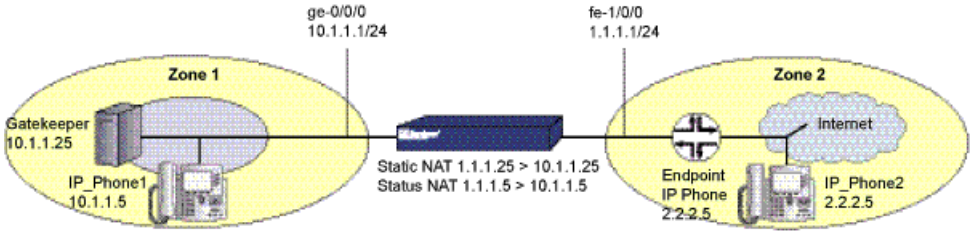

You can use either J-Web or the CLI configuration editor to enable outgoing calls.

This topic covers:

- CLI Configuration on page 182
- Related Topics on page 183

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces fe-1/0/0 unit 0 family inet address 1.1.1.1/24
user@host# set security zones security-zone zone1 interfaces ge-0/0/0.0
user@host# set security zones security-zone zone2 interfaces fe-1/0/0.0
```

2. Configure zones.

```
user@host# set security zones security-zone zone1 interfaces ge-0/0/0.0
user@host# set security zones security-zone zone1 address-book address
  IP_Phone1 10.1.1.5/32
user@host# set security zones security-zone zone1 address-book address
  gatekeeper 10.1.1.25/32
user@host# set security zones security-zone zone2 interfaces fe-1/0/0.0
user@host# set security zones security-zone zone2 address-book address
  IP_Phone2 2.2.2.5/32
user@host# set security zones Global
```

3. Configure interface NAT.

```
user@host# set security nat interface fe-1/0/0.0 static-nat 1.1.1.5/32 host
  10.1.1.5/32
user@host# set security nat interface fe-1/0/0.0 static-nat 1.1.1.25/32 host
  10.1.1.25/32
```

4. Configure policies.

```
user@host# set security policy from-zone zone1 to-zone zone2 policy
  zone1_to_zone2 match source-address IP_Phone1
user@host# set security policy from-zone zone1 to-zone zone2 policy
  zone1_to_zone2 match source-address gatekeeper
user@host# set security policy from-zone zone1 to-zone zone2 policy
  zone1_to_zone2 match destination-address IP_Phone2
user@host# set security policy from-zone zone1 to-zone zone2 policy
  zone1_to_zone2 match application junos-h323
user@host# set security policy from-zone zone1 to-zone zone2 policy
  zone1_to_zone2 then permit
user@host# set security policy from-zone zone2 to-zone Global policy
  zone2_to_Global match source-address IP_Phone2
user@host# set security policy from-zone zone2 to-zone Global policy
  zone2_to_Global match destination-address static_nat_1.1.1.5_32
user@host# set security policy from-zone zone2 to-zone Global policy
  zone2_to_Global match destination-address static_nat_1.1.1.25_32
user@host# set security policy from-zone zone2 to-zone Global policy
  zone2_to_Global match application junos-h323
user@host# set security policy from-zone zone2 to-zone Global policy
  zone2_to_Global then permit
```

Related Topics

- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 172
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 176
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 181
- Understanding the SIP ALG on page 187
- Setting SCCP Inactive Media Timeout on page 257
- Configuring a Media Gateway in Subscribers' Homes on page 282

Using NAT and the H.323 ALG to Enable Incoming Calls

In this example, you configure the device to accept incoming calls over a NAT boundary. To do this, you can create an interface NAT address pool for dynamically allocating destination addresses. This differs from most configurations, where a source pool provides source addresses only. See Figure 20 on page 183.

Before You Begin
For background information, read “Understanding the H.323 ALG” on page 162.

Figure 20: Network Address Translation—Incoming Calls

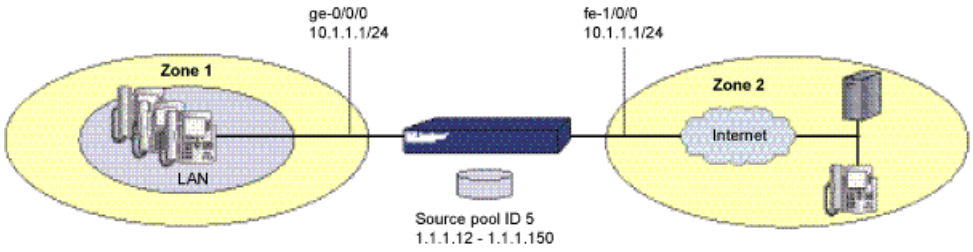

Interface NAT is when the source pool uses the same address as an interface IP address. You can use such address entries as destination addresses in policies, together with H.323, SIP, or other VoIP protocols, to support incoming calls.

To configure incoming calls using NAT, use either the J-Web or CLI Configuration editor.

This topic covers:

- CLI Configuration on page 184
- Related Topics on page 185
- Understanding the Avaya H.323 ALG on page 185

CLI Configuration

In the following example, you configure interfaces, a NAT address pool, zones, and security policies for incoming and outgoing traffic:

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces fe-1/0/0 unit 0 family inet address 1.1.1.1/24
```

2. Configure interface NAT.

```
user@host# set security nat source-nat address-persistent
user@host# set security nat interface fe-1/0/0.0 source-nat pool p1 address-range
   low 1.1.1.12 high 1.1.1.150
```

3. Configure zones.

```
user@host# set security zones security-zone zone1 interfaces ge-0/0/0.0
user@host# set security zones security-zone zone1 address-book address
   IP_Phone1 10.1.1.5/32
user@host# set security zones security-zone zone1 address-book address
   gatekeeper 10.1.1.25/32
user@host# set security zones security-zone zone2 interfaces fe-1/0/0.0
user@host# set security zones security-zone zone2 address-book address
   IP_Phone2 2.2.2.5/32
user@host# set security zones Global
```

4. Configure policies for outgoing traffic.

```
user@host# set security policy from-zone zone1 to-zone zone2 policy
   zone1_to_zone2 match source-address IP_Phone1
user@host# set security policy from-zone zone1 to-zone zone2 policy
   zone1_to_zone2 match source-address gatekeeper
user@host# set security policy from-zone zone1 to-zone zone2 policy
   zone1_to_zone2 match destination-address IP_Phone2
user@host# set security policy from-zone zone1 to-zone zone2 policy
   zone1_to_zone2 match application junos-h323
user@host# set security policy from-zone zone1 to-zone zone2 policy
   zone1_to_zone2 then permit source-nat pool p1
```

5. Configure policies for incoming traffic.

```
user@host# set security policy from-zone zone2 to-zone Global policy
   zone2_to_Global match source-address IP_Phone2
user@host# set security policy from-zone zone2 to-zone Global policy
   zone2_to_Global match destination-address incoming_nat_p1
user@host# set security policy from-zone zone2 to-zone Global policy
   zone2_to_Global match application junos-h323
user@host# set security policy from-zone zone2 to-zone Global policy
   zone2_to_Global then permit
```

Related Topics

- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 172
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 176
- Understanding the SIP ALG on page 187
- Setting SCCP Inactive Media Timeout on page 257
- Configuring a Media Gateway in Subscribers' Homes on page 282

Understanding the Avaya H.323 ALG

The H.323 standard is a legacy VoIP protocol defined by the International Telecommunication Union (ITU-T). H.323 consists of a suite of protocols (such as H.225.0 and H.245) that are used for call signaling and call control for VoIP. The processes for configuring the H.323 standard ALG and the proprietary Avaya H.323 ALG are the same. For more information on the H.323 ALG and how to configure it, see the section "Understanding the H.323 ALG" on page 162.

However, Avaya H.323 ALG has some special features. To understand and configure the Avaya H.323-specific features listed below, see the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.

Before You Begin
For background information, read "Application Layer Gateways (ALGs)" on page 157.

The topics in this section are:

- Avaya H.323-Specific Features on page 185
- Call Flow Details in Avaya H.323 on page 186
- Related Topics on page 187

Avaya H.323-Specific Features

Avaya H.323-specific features are as follows:

- H.323 Fast Connect
- H.323 asymmetric media
- Call waiting
- Call forwarding
- Voice mail
- Call identification
- Conference calling

Call Flow Details in Avaya H.323

- Connecting the Phone in to the Network—Avaya performs the Q.931 Setup/Connect negotiation when the phone is wired into the network rather when a call is being initiated.
- Making a call—When a call is made, since the PBX has already stored the capabilities for each phone when the phone is connected to the network, no further Q.931 and PBX negotiations are required to set up the call. It no longer exchanges Q.931 Setup and Connect messages with the PBX. The phone and the PBX exchange RAS and H.323 Facility messages to setup the call.
- Registering with a CM—When a call has been made, Avaya H.323 registers with the Avaya Communication Manager (CM). The registration process is similar to a generic H.323 standard registration process.

NOTE: The direct mode and tunnel mode are not defined by Avaya H.323 ALG.

For a call to work, the Avaya Communication Manager must be deployed with Avaya Endpoints. During the call, RAS and Q.931 messages are exchanged between the Avaya Communication Manager and the Avaya Endpoints.

NOTE: For Avaya H.323 with a source NAT pool, the registration process allows only one IP address in the pool.

- Setting up RTP/RTCP ports—The Q.931 Setup, Facility and Information messages are used to set up RTP/RTCP ports. The hierarchy for an Avaya H.323 session is Q.931, RTP/RTCP, Parent, and then Child.

NOTE: H.245 ports are not used in an Avaya call flow process.

- Using Avaya H.323 counters—The counters for calls and active calls are not applicable to the Avaya H.323 ALG. The call creation and tearing down is done by Facility messages afterward. When resources are allocated for a call, all counters for calls and active calls increment. If resources are allocated for a call multiple times, messages that belonging to the same call that pass the firewall multiple times will trigger multiple increments of the counters. In other words, messages belong to the same call and pass firewall multiple times might trigger multiple increments of the counters if the resource for a call needs to be allocated multiple times.

  For example, in the two-zone case, the setup, connect message pair allocates one call resource. The active call counter is increased once. Each time the setup-connect message pair passes the firewall, a different call resource with unique interfaces and NAT is allocated. Therefore, the counter increments twice in a three-zone scenario.

Related Topics

- Setting H.323 Endpoint Registration Timeout on page 165
- Setting H.323 Media Source Port Range on page 167
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 168
- Allowing Unknown H.323 Message Types on page 169
- Verifying the H.323 Configuration on page 171
- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 172
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 176
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 183
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 181

Understanding the SIP ALG

Session Initiation Protocol (SIP) is an Internet Engineering Task Force (IETF)-standard protocol for initiating, modifying, and terminating multimedia sessions over the Internet. Such sessions might include conferencing, telephony, or multimedia, with features such as instant messaging and application-level mobility in network environments.

Before You Begin
For background information, read "Application Layer Gateways (ALGs)" on page 157.

J Series devices support SIP as a service and screen SIP traffic, allowing and denying it based on a policy that you configure. SIP is a predefined service in JUNOS Software and uses port 5060 as the destination port.

SIP's primary function is to distribute session-description information and, during the session, to negotiate and modify the parameters of the session. SIP is also used to terminate a multmedia session.

Session-description information is included in INVITE and ACK messages and indicates the multimedia type of the session, for example, whether it is voice or video. Although SIP can use different description protocols to describe the session, the Juniper Networks SIP ALG supports only the Session Description Protocol (SDP).

SDP provides information that a system can use to join a multimedia session. SDP might include information such as IP addresses, port numbers, times, and dates. Note that the IP address and port number in the SDP header (the c= and m= fields, respectively) are the address and port where the client wants to receive the media streams and not the IP address and port number from which the SIP request originates (although they can be the same).

SIP messages consist of requests from a client to a server and responses to the requests from a server to a client with the purpose of establishing a session (or a call). A User Agent (UA) is an application that runs at the endpoints of the call and consists of two parts:

- User Agent Client (UAC), which sends SIP requests on behalf of the user
- User Agent Server (UAS), which listens to the responses and notifies the user when they arrive

Examples of UAs are SIP proxy servers and phones.

This topic covers:

- SIP ALG Operation on page 188
- SDP Session Descriptions on page 189
- Pinhole Creation on page 190

SIP ALG Operation

There are two types of SIP traffic, the signaling and the media stream. SIP signaling traffic consists of request and response messages between client and server and uses transport protocols such as UDP or TCP. The media stream carries the data (audio data, for example) and uses Application Layer protocols such as Real-Time Protocol (RTP) over UDP.

Before You Begin
For background information, read "Understanding the SIP ALG" on page 187.

J Series devices support SIP signaling messages on port 5060. You can simply create a policy that permits SIP service, and the J Series device filters SIP signaling traffic like any other type of traffic, permitting or denying it. The media stream, however, uses dynamically assigned port numbers that can change several times during the course of a call. Without fixed ports, it is impossible to create a static policy to control media traffic. In this case, the device invokes the SIP ALG. The SIP ALG reads SIP messages and their SDP content and extracts the port-number information it needs to dynamically open pinholes to let the media stream traverse the device.

NOTE: We refer to a pinhole as the limited opening of a port to allow exclusive traffic.

The SIP ALG monitors SIP transactions and dynamically creates and manages pinholes based on the information it extracts from these transactions. The Juniper Networks SIP ALG supports all SIP methods and responses (see "SIP ALG Request Methods Overview" on page 191 and "Classes of SIP Responses" on page 212). You can allow SIP transactions to traverse the Juniper Networks firewall by creating a static policy that permits SIP service. This policy enables the device to intercept SIP traffic and do one of the following actions: permit or deny the traffic or enable the SIP ALG to open pinholes to pass the media stream. The SIP ALG needs to open pinholes only

for the SIP requests and responses that contain media information (SDP). For SIP messages that do not contain SDP, the device simply lets them through.

The SIP ALG intercepts SIP messages that contain SDP and, using a parser, extracts the information it requires to create pinholes. The SIP ALG examines the SDP portion of the packet, and a parser extracts information such as IP addresses and port numbers, which the SIP ALG records in a pinhole table. The SIP ALG uses the IP addresses and port numbers recorded in the pinhole table to open pinholes and allow media streams to traverse the device.

NOTE: J Series devices do not support encrypted SDP. If a J Series device receives a SIP message in which SDP is encrypted, the SIP ALG permits it through the firewall but generates a log message informing the user that it cannot process the packet. If SDP is encrypted, the SIP ALG cannot extract the information it needs from SDP to open pinholes. As a result, the media content that SDP describes cannot traverse the device.

SDP Session Descriptions

A Session Description Protocol (SDP) session description is text-based and consists of a set of lines. It can contain session-level and media-level information. The session-level information applies to the whole session, while the media-level information applies to a particular media stream. An SDP session description always contains session-level information, which appears at the beginning of the description, and might contain media-level information, which comes after.

NOTE: In the SDP session description, the media-level information begins with the m = field.

Of the many fields in the SDP description, two are particularly useful to the SIP ALG because they contain Transport Layer information.

- c= for connection information

  This field can appear at the session or media level. It displays in this format:

  c = < network type > < address type > < connection address >

  Currently, the J Series device supports only " IN" (for Internet) as the network type, " IP4" as the address type, and a unicast IP address or domain name as the destination (connection) IP address.

  If the destination IP address is a unicast IP address, the SIP ALG creates pinholes using the IP address and port numbers specified in the media description field m = .

- m= for media announcement

  This field appears at the media level and contains the description of the media. It displays in this format:

m = < media > < port > < transport > < fmt list >

Currently, the J Series device supports only " audio" as the media and " RTP" as the Application Layer transport protocol. The port number indicates the destination (not the origin) of the media stream. The format list (fmt list) provides information on the Application Layer protocol that the media uses.

The J Series device opens ports only for RTP and RTCP. Every RTP session has a corresponding Real Time Control Protocol (RTCP) session. Therefore, whenever a media stream uses RTP, the SIP ALG must reserve ports (create pinholes) for both RTP and RTCP traffic. By default, the port number for RTCP is one higher than the RTP port number.

Pinhole Creation

Both pinholes for the RTP and RTCP traffic share the same destination IP address. The IP address comes from the c = field in the SDP session description. Because the c = field can appear in either the session-level or media-level portion of the SDP session description, the parser determines the IP address based on the following rules (in accordance with SDP conventions):

- First, the SIP ALG parser verifies if there is a c = field containing an IP address in the media level. If there is one, the parser extracts that IP address, and the SIP ALG uses it to create a pinhole for the media.
- If there is no c = field in the media level, the SIP ALG parser extracts the IP address from the c = field in the session level, and the SIP ALG uses it to create a pinhole for the media. If the session description does not contain a c = field in either level, this indicates an error in the protocol stack, and the J Series device drops the packet and logs the event.

The SIP ALG needs the following information to create a pinhole. This information comes from the SDP session description and parameters on the device:

- Protocol—UDP.
- Source IP—Unknown.
- Source port—Unknown.
- Destination IP—The parser extracts the destination IP address from the c = field in the media or session level.
- Destination port—The parser extracts the destination port number for RTP from the m = field in the media level and calculates the destination port number for RTCP using the following formula:

  RTP port number + one

- Lifetime—This value indicates the length of time (in seconds) during which a pinhole is open to allow a packet through. A packet must go through the pinhole before the lifetime expires. When the lifetime expires, the SIP ALG removes the pinhole.

When a packet goes through the pinhole within the lifetime period, immediately afterwards the SIP ALG removes the pinhole for the direction from which the packet came.

Figure 21 on page 191 describes a call setup between two SIP clients and how the SIP ALG creates pinholes to allow RTP and RTCP traffic. The illustration assumes that the J Series device has a policy that permits SIP, thus opening port 5060 for SIP signaling messages.

Figure 21: SIP ALG Call Setup

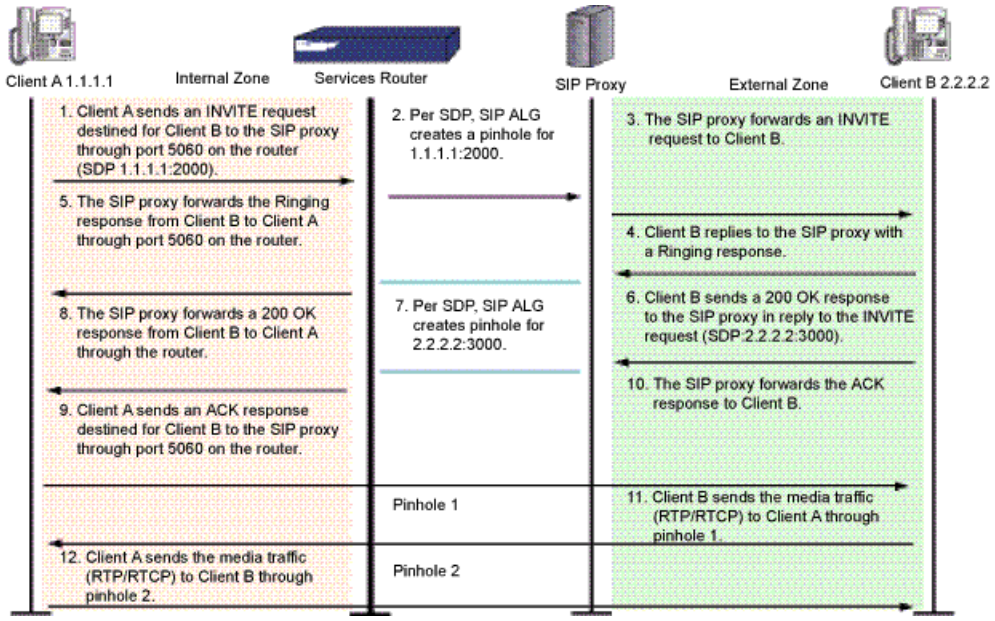

NOTE: The SIP ALG does not create pinholes for RTP and RTCP traffic when the destination IP address is 0.0.0.0, which indicates that the session is on hold. To put a session on hold during a telephone communication, for example, user A sends user B a SIP message in which the destination IP address is 0.0.0.0. Doing so indicates to user B that it should not send any media until further notice. If user B sends media anyway, the device drops the packets.

SIP ALG Request Methods Overview

The SIP transaction model includes a number of request and response messages, each of which contains a *method* field that denotes the purpose of the message.

Before You Begin
For background information, read "Understanding the SIP ALG" on page 187.

JUNOS Software supports the following method types and response codes:

- INVITE—A user sends an INVITE request to invite another user to participate in a session. The body of an INVITE request may contain the description of the session.
- ACK—The user from whom the INVITE originated sends an ACK request to confirm reception of the final response to the INVITE request. If the original INVITE request did not contain the session description, the ACK request must include it.
- OPTIONS—Used by the User Agent (UA) to obtain information about the capabilities of the SIP proxy. A server responds with information about what methods, session description protocols, and message encoding it supports.
- BYE—A user sends a BYE request to abandon a session. A BYE request from either user automatically terminates the session.
- CANCEL—A user can send a CANCEL request to cancel a pending INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE had sent a final response for the INVITE before it received the CANCEL.
- REGISTER—A user sends a REGISTER request to a SIP registrar server to inform it of the current location of the user. A SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server attempting to locate a user.
- Info—Used to communicate mid-session signaling information along the signaling path for the call.
- Subscribe—Used to request current state and state updates from a remote node.
- Notify—Sent to inform subscribers of changes in state to which the subscriber has a subscription.
- Refer—Used to refer the recipient (identified by the Request-URI) to a third party by the contact information provided in the request.

  For example, if user A in a private network refers user B, in a public network, to user C, who is also in the private network, the SIP ALG allocates a new IP address and port number for user C so that user C can be contacted by user B. If user C is registered with a registrar, however, its port mapping is stored in the ALG NAT table and is reused to perform the translation.
- Update—Used to open pinhole for new or updated SDP information. The Via:, From:, To:, Call-ID:, Contact:, Route:, and Record-Route: header fields are modified.
- 1*xx*, 202, 2*xx*, 3*xx*, 4*xx*, 5*xx*, 6*xx* Response Codes—Used to indicate the status of a transaction. Header fields are modified.

Related Topics

- Understanding SIP ALG Call Duration and Timeouts on page 196

Configuring the SIP ALG—Quick Configuration

You can use J-Web Quick Configuration to quickly configure SIP ALG Parameters

Before You Begin
For background information, read “Understanding the SIP ALG” on page 187

Figure 22 on page 194 shows the SIP ALG configuration page.

Figure 22: SIP ALG Configuration

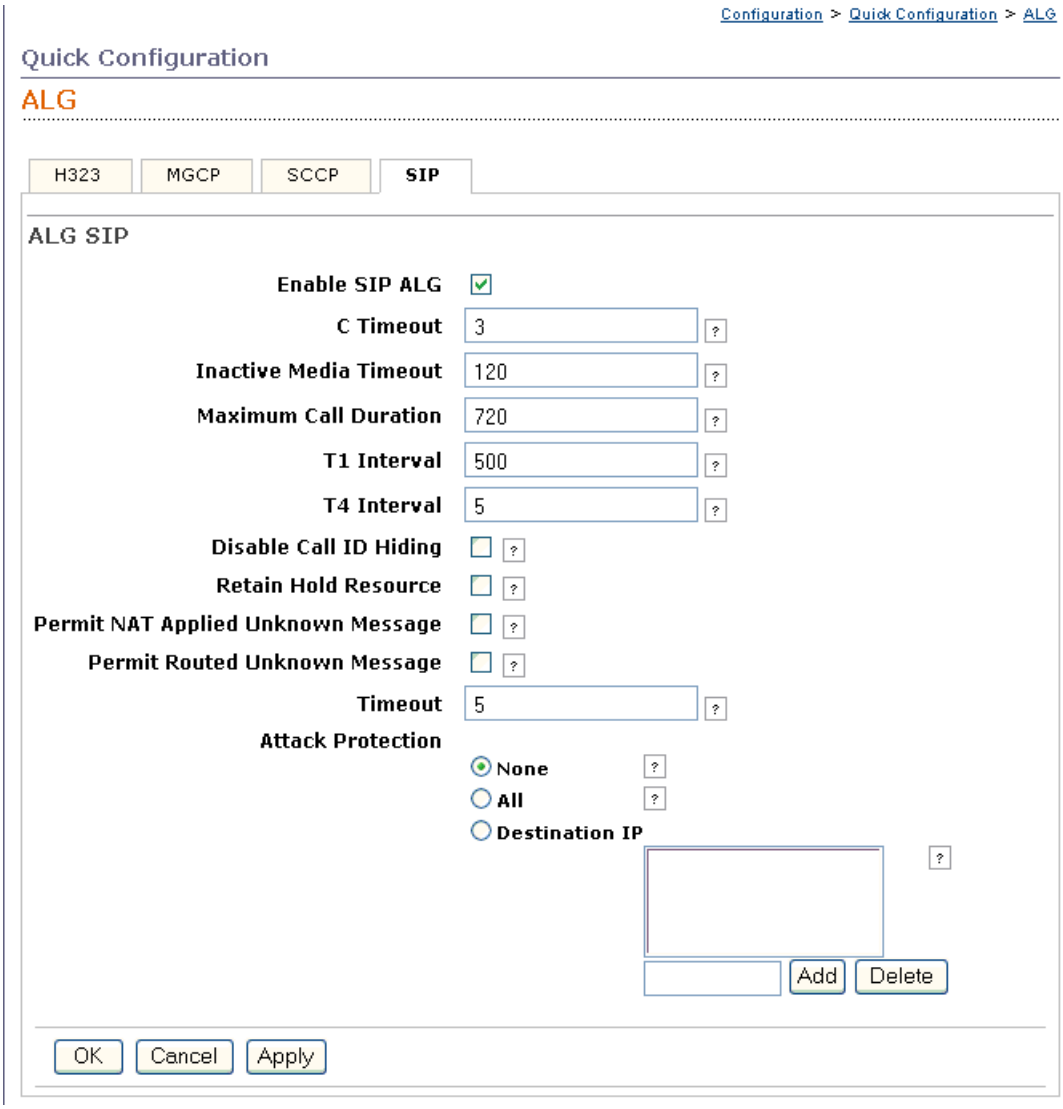

To configure the SIP ALG with Quick Configuration:

1. In the J-Web user interface, select **Configure > ALGs > VoIP ALG**.
2. Select the **Sip** tab if it is not selected.
3. Fill in the parameter settings as described in Table 37 on page 195 and click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel.**

Table 37: SIP ALG Configuraiton Options

Enable SIP ALG	Enables or disables the SIP ALG.	Click the check box.
C Timeout	Specifies the INVITE transaction timeout at the proxy, in minutes; the default is 3. Because the SIP ALG is in the middle, instead of using the INVITE transaction timer value B (which is (64 * T1) = 32 seconds), the SIP ALG gets its timer value from the proxy.	Select a value between 3 and 10 minutes.
Inactive Media Timeout	Specifies the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the temporary openings (pinholes) in the firewall SIP ALG opened for media are closed. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.	Select a value between 10 and 2,550 seconds.
Maximum Call Duration	Sets the absolute maximum length of a call. When a call exceeds this parameter setting, the SIP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, the range is from 3 to 7200 minutes.	Select a value between 3 and 7,200 minutes.
T1 Interval	Specifies the roundtrip time estimate, in seconds, of a transaction between endpoints. The default is 500 milliseconds. Because many SIP timers scale with the T1-Interval (as described in RFC 3261), when you change the value of the T1-Interval timer, those SIP timers also are adjusted.	Select a value between 500 and 5,00 milliseconds.
T4 Interval	Specifies the maximum time a message remains in the network. The default is 5 seconds, the range is 5 to 10 seconds. Because many SIP timers scale with the T4-Interval (as described in RFC 3261), when you change the value of the T4-Interval timer, those SIP timers also are adjusted.	Select a value between 5and 10 seconds..
Disable Call ID Hiding	Enables or disable translation of the host IP address in the call-ID header. Translation is enabled by default.	Click the check box.
Retain Hold Resource	Enable or disables whether the device frees media resources for a Session Initiation Protocol (SIP) Application Layer Gateway (ALG), even when a media stream is placed on hold. By default, media stream resources are released when the media stream is held.	Click the check box.

Table 37: SIP ALG Configuraiton Options *(continued)*

Permit NAT Applied Unknown Message	Specifies how unidentified SIP messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown SIP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped. This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.	Click the check box.
Permit Routed Unknown Message	Specifies that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)	Click the check box.
Timeout	Specifies the amount of time (in seconds) to make an attack table entry for each INVITE, which is listed in the application screen.	Enter a value between 1 and 3,600 seconds.
Attack Protection	Protects servers against INVITE attacks. Configure the SIP application screen to protect the server at some or all destination IP addresses against INVITE attacks. You can include up to 16 destination IP addresses of servers to be protected.	Select None, All or, if you select Destination IP, enter or select an IP address.

Understanding SIP ALG Call Duration and Timeouts

The call duration and timeout features give you control over SIP call activity and help you to manage network resources.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191

Typically a call ends when one of the clients sends a BYE or CANCEL request. The SIP ALG intercepts the BYE or CANCEL request and removes all media sessions for that call. There could be reasons or problems preventing clients in a call from sending BYE or CANCEL requests, for example, a power failure. In this case, the call might go on indefinitely, consuming resources on the device.

A call can have one or more voice channels. Each voice channel has two sessions (or two media streams), one for Real-Time Transport Protocol (RTP) traffic and one for Real-Time Control Protocol (RTCP) signalling. When managing the sessions, the device considers the sessions in each voice channel as one group. Timeouts and call duration settings apply to a group as opposed to each session.

The following parameters govern SIP call activity:

- inactive-media-timeout—This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the temporary openings (pinholes) in the firewall the SIP ALG opened for media are closed. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.
- maximum-call-duration—This parameter sets the absolute maximum length of a call. When a call exceeds this parameter setting, the SIP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, the range is from 3 to 7200 minutes.
- t1-interval—This parameter specifies the roundtrip time estimate, in seconds, of a transaction between endpoints. The default is 500 milliseconds. Because many SIP timers scale with the T1-Interval (as described in RFC 3261), when you change the value of the T1-Interval timer, those SIP timers also are adjusted.
- t4-interval—This parameter specifies the maximum time a message remains in the network. The default is 5 seconds, the range is 5 to 10 seconds. Because many SIP timers scale with the T4-Interval (as described in RFC 3261), when you change the value of the T4-Interval timer, those SIP timers also are adjusted.
- c-timeout—This parameter specifies the INVITE transaction timeout at the proxy, in minutes; the default is 3. Because the SIP ALG is in the middle, instead of using the INVITE transaction timer value B (which is (64 * T1) = 32 seconds), the SIP ALG gets its timer value from the proxy.

Related Topics

- Setting SIP Call Duration and Inactive Media Timeout on page 197

Setting SIP Call Duration and Inactive Media Timeout

The call duration and inactive media timeout features help you to conserve network resources and maximise throughput.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- Understanding SIP ALG Call Duration and Timeouts on page 196

You use the maximum-call-duration parameter to set the maximum allowable length of time a call can be active. When the duration is exceeded, the SIP ALG tears down the call and releases the media sessions. This setting also frees up bandwidth in cases where calls fail to properly terminate.

This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the SIP ALG temporary openings (pinholes) in the firewall opened for media are closed. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

NOTE: You must use the CLI to set SIP-signaling and media-inactivity timeouts.

Use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 198
- CLI Configuration on page 199
- Related Topics on page 199

J-Web Configuration

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, select **Configure** or **Edit**.
3. Next to Alg, select **Configure** or **Edit**.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.
5. Enter a value for any of the following parameters:
   - C timeout
   - Inactive media timeout
   - Maximum call duration
   - T1 interval
   - T4 interval
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In the following example, you set call duration to 3000 minutes and the media inactivity timeout to 90 seconds:

```
user@host# set security alg sip maximum-call-duration 3000
user@host# set security alg sip inactive-media-timeout 90
```

Related Topics

- Understanding SIP ALG Call Duration and Timeouts on page 196
- Configuring SIP Denial of Service (DoS) Attack Protection on page 199
- Allowing Unknown SIP Message Types on page 200

Configuring SIP Denial of Service (DoS) Attack Protection

The ability of the SIP proxy server to process calls can be impacted by repeat SIP INVITE requests—requests that it initially denied. The DoS protection feature enables you to configure the device to monitor INVITE requests and proxy server replies to them. If a reply contains a *3xx*, *4xx*, or *5xx* response code (see "Classes of SIP Responses" on page 212), the ALG stores the source IP address of the request and the IP address of the proxy server in a table. Subsequently, the device checks all INVITE requests against this table and, for a configurable number of seconds (the default is 3), discards any packets that match entries in the table. You can configure the device to monitor and deny repeat INVITE requests to all proxy servers, or you can protect a specific proxy server by specifying the destination IP address. SIP attack protection is configured globally.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191

To configure DoS attack protection, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 200
- CLI Configuration on page 200
- Related Topics on page 200

J-Web Configuration

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, select **Configure** or **Edit**.
3. Next to Alg, select **Configure** or **Edit**.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.
5. Next to Application screen, click Configure.
6. Next to Protect, click Configure.
7. Click the Deny check box.
8. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In this example, you configure the device to protect a single SIP proxy server (1.1.1.3) from repeat INVITE requests to which it has already denied service. Packets are dropped for a period of 5 seconds, after which the device resumes forwarding INVITE requests from those sources.

```
user@host# set security alg sip application-screen protect deny destination-ip 1.1.1.3
user@host# set security alg sip application-screen protect deny timeout 5
```

Related Topics

- Configuring SIP Denial of Service (DoS) Attack Protection on page 199
- Allowing Unknown SIP Message Types on page 200

Allowing Unknown SIP Message Types

To accommodate on-going development of the Session Initiation Protocol (SIP), you might want to allow traffic containing new SIP message types. The unknown SIP message type feature enables you to configure the device to accept SIP traffic containing unknown message types in both NAT and route modes.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191

This feature enables you to specify how unidentified SIP messages are handled by the device. The default is to drop unknown (unsupported) messages. We do not recommend permitting unknown messages because they can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown SIP messages can help you get your network operational. so you can later analyze your VoIP traffic to determine why some messages were being dropped.

Note that this command applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

- permit-nat-applied specifies that unknown messages be allowed to pass if the session is in NAT mode.
- permit-routed specifies that unknown messages be allowed to pass if the session is in Route mode.

To allow unknown messages, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 201
- CLI Configuration on page 202
- Related Topics on page 202

J-Web Configuration

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, select **Configure** or **Edit**.
3. Next to Alg, select **Configure** or **Edit**.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.
5. Next to Application screen, click Configure.
6. Next to Unknown message, click Configure.
7. Click one of the following check boxes:
   - To allow unknown message types in NAT mode, click the Permit nat applied Yes check box.
   - To allow unknown message types in route mode, click the Permit routed Yes check box.
8. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In this example, you configure the device to allow unknown message types in SIP traffic in both route and NAT modes.

```
user@host# set security alg sip application-screen unknown-message permit-nat-applied
    permit-routed
```

Related Topics

- Understanding SIP ALG Call Duration and Timeouts on page 196
- Setting SIP Call Duration and Inactive Media Timeout on page 197
- Configuring SIP Denial of Service (DoS) Attack Protection on page 199

Disabling SIP Call ID Hiding

The Call-ID header in a SIP packet contains the hostname or IP address of the calling entity. In NAT mode, this identifying information is replaced by a public IP address and port number. This is a SIP requirement, and the default behavior of the device. If your implementation of SIP includes proprietary packet headers that you want to reveal, you can use the disable call ID hiding feature to disable this function and enable those call flows. This behavior does not conform to the SIP standard.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191

To disable SIP call hiding, use either the J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 202
- CLI Configuration on page 203
- Related Topics on page 203

J-Web Configuration

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, select **Configure** or **Edit**.
3. Next to Alg, select **Configure** or **Edit**.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.

5. Click the Disable call id hiding Yes check box.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

Use the following command to allow proprietary SIP call flows:

```
user@host# set security alg sip disable-call-id-hiding
```

Related Topics

- Understanding SIP ALG Call Duration and Timeouts on page 196
- Setting SIP Call Duration and Inactive Media Timeout on page 197
- Configuring SIP Denial of Service (DoS) Attack Protection on page 199

Retaining SIP Hold Resources

When a user puts a call on hold, the SIP ALG releases SDP media resources, such as pinholes and translation contexts. When the user resumes the call, an INVITE request message negotiates a new SDP offer and answer and the SIP ALG reallocates resources for the media stream. This can result in new translated IP address and port numbers for the media description even when the media description is the same as the previous description. This is compliant with *RFC 3264 An Offer/Answer Model with the Session Description Protocol (SDP)*.

Some proprietary SIP implementations have designed call flows so that the user agent (UA) module ignores the new SDP INVITE offer and continues to use the SDP offer of the previous negotiation. To accommodate this functionality, you must configure the device to retain SDP media resources when a call is put on hold for reuse when the call is resumed.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191

To retain SIP hold resources, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 204
- CLI Configuration on page 204
- Related Topics on page 204

J-Web Configuration

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, select **Configure** or **Edit**.
3. Next to Alg, select **Configure** or **Edit**.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.
5. Click the Retain hold resources check box.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

Use the following command to accommodate proprietary SIP call flows:

```
user@host# set security alg sip retain-hold-resource
```

Related Topics

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- SDP Session Descriptions on page 189
- Pinhole Creation on page 190

Understanding SIP with Network Address Translation (NAT)

The Network Address Translation (NAT) protocol enables multiple hosts in a private subnet to share a single public IP address to access the Internet. For outgoing traffic, NAT replaces the private IP address of the host in the private subnet with the public IP address. For incoming traffic, the public IP address is converted back into the

private address, and the message is routed to the appropriate host in the private subnet.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191

Using NAT with the SIP service is more complicated because SIP messages contain IP addresses in the SIP headers as well as in the SIP body. When using NAT with the SIP service, the SIP headers contain information about the caller and the receiver, and the device translates this information to hide it from the outside network. The SIP body contains the Session Description Protocol (SDP) information, which includes IP addresses and port numbers for transmission of the media. The device translates SDP information for allocating resources to send and receive the media.

How IP addresses and port numbers in SIP messages are replaced depends on the direction of the message. For an outgoing message, the private IP address and port number of the client are replaced with the public IP address and port number of the Juniper Networks firewall. For an incoming message, the public address of the firewall is replaced with the private address of the client.

When an INVITE message is sent out across the firewall, the SIP ALG collects information from the message header into a call table, which it uses to forward subsequent messages to the correct endpoint. When a new message arrives, for example an ACK or 200 OK, the ALG compares the " From:, To:, and " Call-ID: fields against the call table to identify the call context of the message. If a new INVITE message arrives that matches the existing call, the ALG processes it as a REINVITE.

When a message containing SDP information arrives, the ALG allocates ports and creates a NAT mapping between them and the ports in the SDP. Because the SDP requires sequential ports for the Real-Time Protocol (RTP) and Real-Time Control Protocol (RTCP) channels, the ALG provides consecutive even-odd ports. If it is unable to find a pair of ports, it discards the SIP message.

This topic covers:

- Outgoing Calls on page 206
- Incoming Calls on page 206
- Forwarded Calls on page 207
- Call Termination on page 207
- Call Re-INVITE Messages on page 207
- Call Session Timers on page 207
- Call Cancellation on page 207
- Forking on page 208
- SIP Messages on page 208

- SIP Headers on page 208
- SIP Body on page 210
- SIP NAT Scenario on page 210
- Classes of SIP Responses on page 212
- Related Topics on page 213

Outgoing Calls

When a SIP call is initiated with a SIP request message from the internal to the external network, NAT replaces the IP addresses and port numbers in the SDP and binds the IP addresses and port numbers to the Juniper Networks firewall. Via, Contact, Route, and Record-Route SIP header fields, if present, are also bound to the firewall IP address. The ALG stores these mappings for use in retransmissions and for SIP response messages.

The SIP ALG then opens pinholes in the firewall to allow media through the device on the dynamically assigned ports negotiated based on information in the SDP and the Via, Contact, and Record-Route header fields. The pinholes also allow incoming packets to reach the Contact, Via, and Record-Route IP addresses and ports. When processing return traffic, the ALG inserts the original Contact, Via, Route, and Record-Route SIP fields back into the packets.

Incoming Calls

Incoming calls are initiated from the public network to public static NAT addresses or to interface IP addresses on the device. Static NATs are statically configured IP addresses that point to internal hosts; interface IP addresses are dynamically recorded by the ALG as it monitors REGISTER messages sent by internal hosts to the SIP registrar. When the device receives an incoming SIP packet, it sets up a session and forwards the payload of the packet to the SIP ALG.

The ALG examines the SIP request message (initially an INVITE) and, based on information in the SDP, opens gates for outgoing media. When a 200 OK response message arrives, the SIP ALG performs NAT on the IP addresses and ports and opens pinholes in the outbound direction. (The opened gates have a short time-to-live, and they time out if a 200 OK response message is not received quickly.)

When a 200 OK response arrives, the SIP proxy examines the SDP information and reads the IP addresses and port numbers for each media session. The SIP ALG on the device performs NAT on the addresses and port numbers, opens pinholes for outbound traffic, and refreshes the timeout for gates in the inbound direction.

When the ACK arrives for the 200 OK, it also passes through the SIP ALG. If the message contains SDP information, the SIP ALG ensures that the IP addresses and port numbers are not changed from the previous INVITE—if they are, the ALG deletes old pinholes and creates new pinholes to allow media to pass through. The ALG also monitors the Via, Contact, and Record-Route SIP fields and opens new pinholes if it determines that these fields have changed.

Forwarded Calls

A forwarded call is when, for example, user A outside the network calls user B inside the network, and user B forwards the call to user C outside the network. The SIP ALG processes the INVITE from user A as a normal incoming call. But when the ALG examines the forwarded call from B to C outside the network and notices that B and C are reached using the same interface, it does not open pinholes in the firewall, because media will flow directly between user A and user C.

Call Termination

The BYE message terminates a call. When the device receives a BYE message, it translates the header fields just as it does for any other message. But because a BYE message must be acknowledged by the receiver with a 200 OK, the ALG delays call teardown for five seconds to allow time for transmission of the 200 OK.

Call Re-INVITE Messages

Re-INVITE messages are used to add new media sessions to a call, and to remove existing media sessions. When new media sessions are added to a call, new pinholes are opened in the firewall and new address bindings created. The process is identical to the original call setup. When one or more media sessions are removed from a call, pinholes are closed and bindings released just as with a BYE message.

Call Session Timers

The SIP ALG uses the Session-Expires value to time out a session if a Re-INVITE or UPDATE message is not received. The ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE and uses this value for signaling timeout. If the ALG receives another INVITE before the session times out, it resets all timeout values to this new INVITE or to default values, and the process is repeated.

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the device is protected should one of the following events occur:

- End systems crash during a call and a BYE message is not received.
- Malicious users never send a BYE in an attempt to attack a SIP ALG.
- Poor implementations of SIP proxy fail to process Record-Route and never send a BYE message.
- Network failures prevent a BYE message from being received.

Call Cancellation

Either party can cancel a call by sending a CANCEL message. Upon receiving a CANCEL message, the SIP ALG closes pinholes through the firewall—if any have been opened—and releases address bindings. Before releasing the resources, the ALG delays the control channel age-out for approximately five seconds to allow time for

the final 200 OK to pass through. The call is terminated when the five second timeout expires, regardless of whether a 487 or non-200 response arrives.

Forking

Forking enables a SIP proxy to send a single INVITE message to multiple destinations simultaneously. When the multiple 200 OK response messages arrive for the single call, the SIP ALG parses but updates call information with the first 200 OK message it receives.

SIP Messages

The SIP message format consists of a SIP header section and the SIP body. In request messages, the first line of the header section is the request line, which includes the method type, request-URI, and protocol version. In response messages, the first line is the status line, which contains a status code. SIP headers contain IP addresses and port numbers used for signaling. The SIP body, separated from the header section by a blank line, is reserved for session description information, which is optional. J Series devices currently support the Session Description Protocol (SDP) only. The SIP body contains IP addresses and port numbers used to transport the media.

SIP Headers

In the following sample SIP request message, NAT replaces the IP addresses in the header fields—shown in bold font—to hide them from the outside network.

```
INVITE bob@10.150.20.5 SIP/2.0
Via: SIP/2.0/UDP 10.150.20.3:5434
From: alice@10.150.20.3
To: bob@10.150.20.5
Call-ID: a12abcde@10.150.20.3
Contact: alice@10.150.20.3:5434
Route: <sip:netscreen@10.150.20.3:5060>
Record-Route: <sip:netscreen@10.150.20.3:5060>
```

How IP address translation is performed depends on the type and direction of the message. A message can be any of the following:

- Inbound request
- Outbound response
- Outbound request
- Inbound response

Table 38 on page 209 shows how NAT is performed in each of these cases. Note that for several of the header fields the ALG must know more than just whether the messages comes from inside or outside the network. It must also know what client initiated the call, and whether the message is a request or response.

Table 38: Requesting Messages with NAT Table continued on next page

Inbound Request (from public to private)	To:	Replace domain with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	Replace ALG address with local address
	Contact:	None
	Record-Route:	None
	Route:	None
Outbound Response (from private to public)	To:	Replace ALG address with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	N/A
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	None
Outbound Request (from private to public)	To:	None
	From:	Replace local address with ALG address
	Call-ID:	Replace local address with ALG address
	Via:	Replace local address with ALG address
	Request-URI:	None
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	Replace ALG address with local address

Table 38: Requesting Messages with NAT Table continued on next page *(continued)*

Outbound Response	To:	None
(from public to private)	From:	Replace ALG address with local address
	Call-ID:	Replace ALG address with local address
	Via:	Replace ALG address with local address
	Request-URI:	N/A
	Contact:	None
	Record-Route:	Replace ALG address with local address
	Route:	Replace ALG address with local address

SIP Body

The SDP information in the SIP body includes IP addresses the ALG uses to create channels for the media stream. Translation of the SDP section also allocates resources, that is, port numbers to send and receive the media.

The following excerpt from a sample SDP section shows the fields that are translated for resource allocation.

```
o=user 2344234 55234434 IN IP4 10.150.20.3
c=IN IP4 10.150.20.3
m=audio 43249 RTP/AVP 0
```

SIP messages can contain more than one media stream. The concept is similar to attaching multiple files to an e-mail message. For example, an INVITE message sent from a SIP client to a SIP server might have the following fields:

```
c=IN IP4 10.123.33.4
m=audio 33445 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33447 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33449 RTP/AVP 0
```

J Series devices support up to 6 SDP channels negotiated for each direction, for a total of 12 channels per call. For more information, see “SDP Session Descriptions” on page 189.

SIP NAT Scenario

Figure 23 on page 211 and Figure 24 on page 212 show a SIP call INVITE and 200 OK. In Figure 23 on page 211, ph1 sends a SIP INVITE message to ph2. Note how the IP addresses in the header fields—shown in bold font—are translated by the device.

The SDP section of the INVITE message indicates where the caller is willing to receive media. Note that the Media Pinhole contains two port numbers, 52002 and 52003, for RTCP and RTP. The Via/Contact Pinhole provides port number 5060 for SIP signaling.

Observe how, in the 200 OK response message in Figure 24 on page 212, the translations performed in the INVITE message are reversed. The IP addresses in this message, being public, are not translated, but gates are opened to allow the media stream access to the private network.

Figure 23: SIP NAT Scenario 1

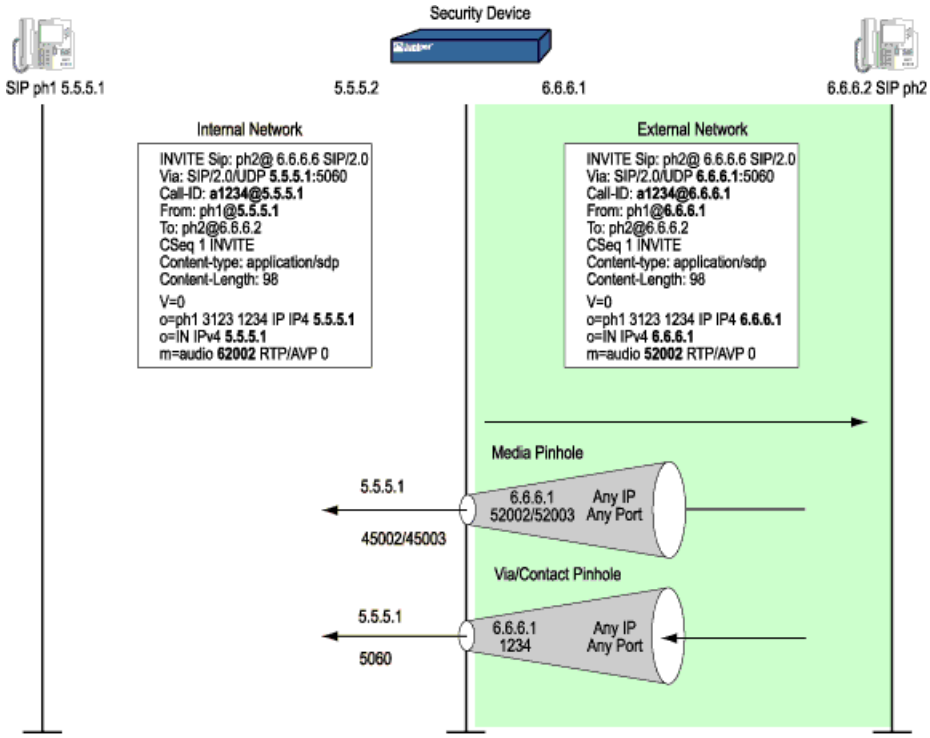

Figure 24: SIP NAT Scenario 2

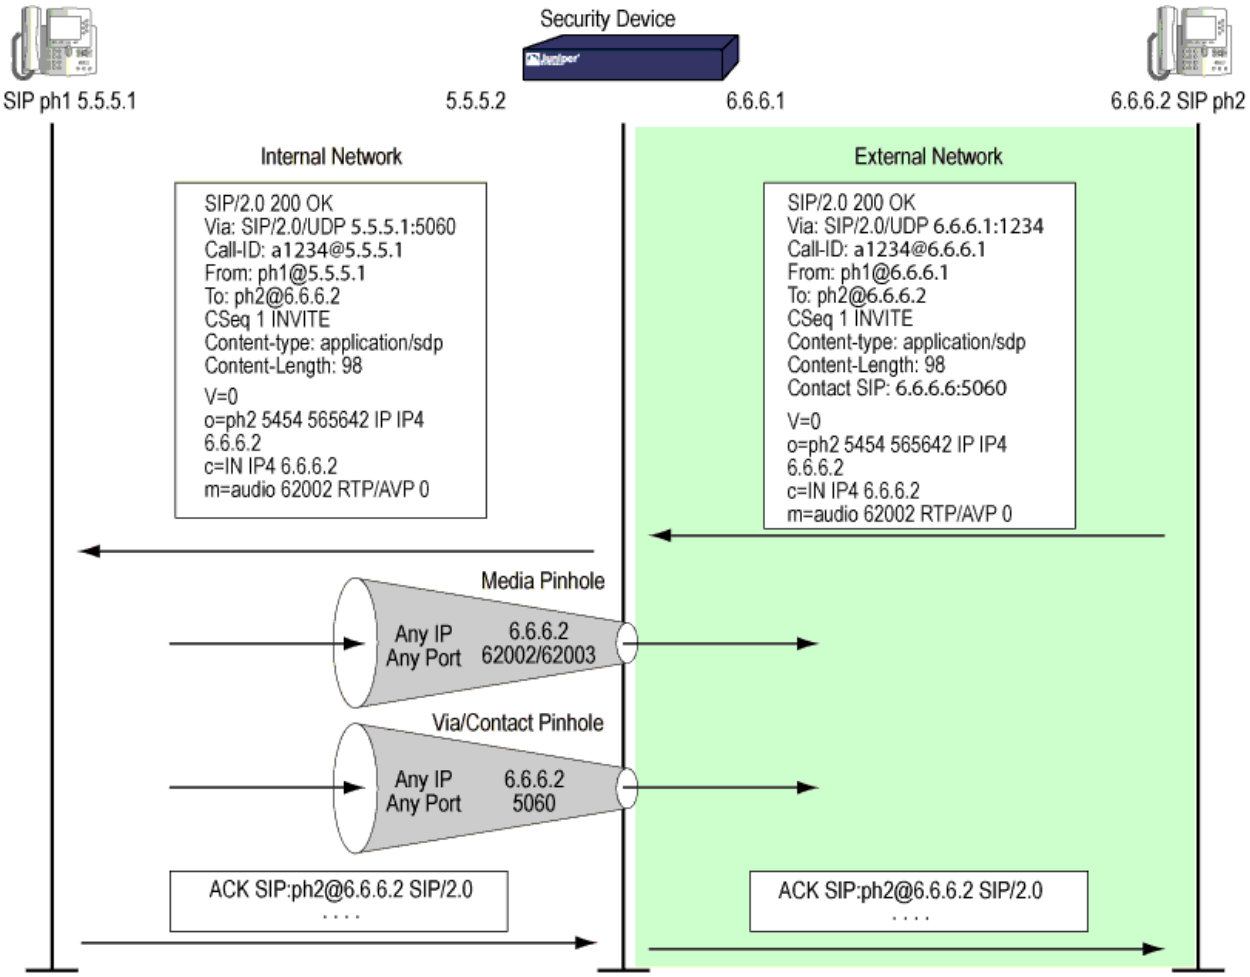

Classes of SIP Responses

SIP responses provide status information about SIP transactions and include a response code and a reason phrase. SIP responses are grouped into the following classes:

- Informational (100 to 199)—Request received, continuing to process the request.
- Success (200 to 299)—Action successfully received, understood, and accepted.
- Redirection (300 to 399)—Further action required to complete the request.
- Client Error (400 to 499)—Request contains bad syntax or cannot be fulfilled at this server.
- Server Error (500 to 599)—Server failed to fulfill an apparently valid request.
- Global Failure (600 to 699)—Request cannot be fulfilled at any server.

Table 39 on page 213 provides a complete list of current SIP responses.

Table 39: SIP Responses

Informational	100 Trying	180 Ringing	181 Call is being forwarded
	182 Queued	183 Session progress	
Success	200 OK	202 Accepted	
Redirection	300 Multiple choices	301 Moved permanently	302 Moved temporarily
	305 Use proxy	380 Alternative service	
Client Error	400 Bad request	401 Unauthorized	402 Payment required
	403 Forbidden	404 Not found	405 Method not allowed
	406 Not acceptable	407 Proxy authentication required	408 Request time-out
	409 Conflict	410 Gone	411 Length required
	413 Request entity too large	414 Request-URL too large	415 Unsupported media type
	420 Bad extension	480 Temporarily not available	481 Call leg/transaction does not exist
	482 Loop detected	483 Too many hops	484 Address incomplete
	485 Ambiguous	486 Busy here	487 Request canceled
	488 Not acceptable here		
Server Error	500 Server internal error	501 Not implemented	502 Bad gateway
	502 Service unavailable	504 Gateway time-out	505 SIP version not supported
Global Failure	600 Busy everywhere	603 Decline	604 Does not exist anywhere
	606 Not acceptable		

Related Topics

- Configuring Host Inbound Traffic on page 53
- Understanding Security Zones on page 47

Understanding Incoming SIP Call Support Using the SIP Registrar

SIP registration provides a discovery capability by which SIP proxies and location servers are able to identify the location or locations where users want to be contacted. A user registers one or more contact locations by sending a REGISTER message to the registrar. The To and Contact fields in the REGISTER message contain the address-of-record Uniform Resource Identifier (URI) and one or more contact URIs,

as shown in Figure 25 on page 215. Registration creates bindings in a location service that associates the address-of-record with the contact address or addresses.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- Understanding SIP with Network Address Translation (NAT) on page 204

The device monitors outgoing REGISTER messages, performs NAT on these addresses, and stores the information in an Incoming Incoming NAT table. Then, when an INVITE message is received from outside the network, the device uses the Incoming NAT table to identify which internal host to route the INVITE message to. You can take advantage of SIP proxy registration service to allow incoming calls by configuring interface source NAT or NAT pools on the egress interface of the device. Interface source NAT is adequate for handling incoming calls in a small office, while we recommend setting up source NAT pools for larger networks or an enterprise environment.

NOTE: Incoming call support using interface source NAT or a source NAT pool is supported for SIP and H.323 services only.For incoming calls, J Series devices currently support UDP and TCP only. Domain name resolution is also currently not supported; therefore, URIs must contain IP addresses, as shown in Figure 25 on page 215.

Figure 25: Using the SIP Registrar

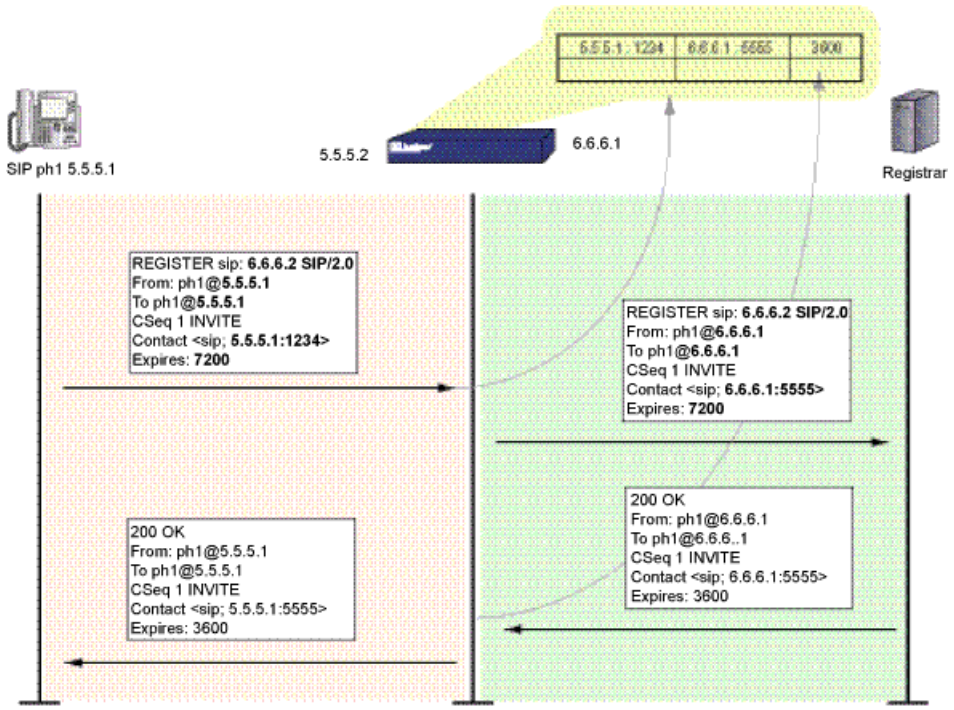

Related Topics

- Configuring Interface Source NAT for Incoming SIP Calls on page 216
- Configuring a Source NAT Pool for Incoming SIP Calls on page 218
- Configuring Static NAT for Incoming SIP Calls on page 223
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring the SIP Proxy in the Public Zone on page 230
- Configuring a Three-Zone SIP Scenario on page 235
- Incoming Calls on page 206
- Outgoing Calls on page 206

Configuring Interface Source NAT for Incoming SIP Calls

In a two-zone scenario with the SIP proxy server in an external, or public zone, you can use NAT for incoming calls by configuring source NAT on the interface to the public zone.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213

In this example, phone1 is on the ge-0/0/0 interface in the private zone, and phone2 and the proxy server are on the ge-0/0/2 interface in the public zone. You configure interface source NAT on ge-0/0/2.0 for incoming calls, then create a policy permitting SIP traffic from the public zone to the private zone and reference the source NAT in the policy. You also create a policy that permits SIP traffic from the private to the public zone, again referencing the source NAT address pool. This enables phone1 in the private zone to register with the proxy in the public zone. See Figure 26 on page 216.

Figure 26: Source NAT for Incoming Calls

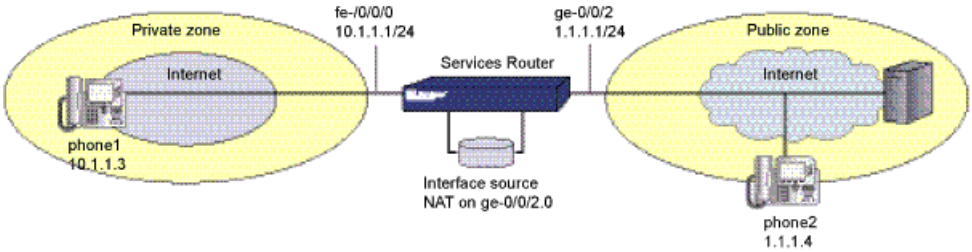

To configure interface source NAT for incoming calls, use either the J-Web or the CLI configuration editor.

This topic covers:

- CLI Configuration on page 216
- Related Topics on page 217

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

```
user@host# set security zones security-zone private interface ge-0/0/0.0
```

2. Configure addresses.

```
user@host# set security zones security-zone private address-book address phone1
  10.1.1.3/32
user@host# set security zones security-zone public address-book address proxy
  10.1.1.3/32
user@host# set security zones security-zone public address-book address phone2
  1.1.1.4/32
```

3. Configure zones.

```
user@host# set security zones security-zone private
user@host# set security zones security-zone public
user@host# set security zones security-zone private interfaces ge-0/0/0.0
user@host# set security zones security-zone public interfaces ge-0/0/2.0
```

4. Configure source NAT.

```
user@host# set security nat interface ge-0/0/2.0 source-nat allow-incoming
user@host# set security nat source-nat address-persistent
```

5. Configure policies.

```
user@host# set security policies from-zone private to-zone public policy outgoing
  match source-address phone1 destination-address any application junos-sip
user@host# set security policies from-zone private to-zone public policy outgoing
  then permit source-nat interface
user@host# set security policies from-zone public to-zone private policy incoming
  match source-address any destination-address incoming-nat-fe0/0/2.0
  application junos-sip
user@host# set security policies from-zone public to-zone private policy incoming
  then permit
```

Related Topics

- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213
- Configuring Static NAT for Incoming SIP Calls on page 223
- Configuring a Source NAT Pool for Incoming SIP Calls on page 218
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring the SIP Proxy in the Public Zone on page 230
- Configuring a Three-Zone SIP Scenario on page 235

Configuring a Source NAT Pool for Incoming SIP Calls

In a two-zone scenario with the SIP proxy server in an external, or public zone, you can use NAT for incoming calls by configuring a NAT pool on the interface to the public zone.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213

In this example, phone1 is in the private zone, and phone2 and the proxy server are in the public zone. You configure a source NAT pool on the ge-0/0/2.0 interface to do NAT on incoming calls, then set a policy permitting SIP traffic from the public zone to the private zone and reference the NAT pool in the policy. You also create a policy that permits SIP traffic from the private to the public zone.. This enables phone1 in the private zone to register with the proxy in the public zone. See Figure 27 on page 218.

Figure 27: Source NAT Pool for Incoming Calls

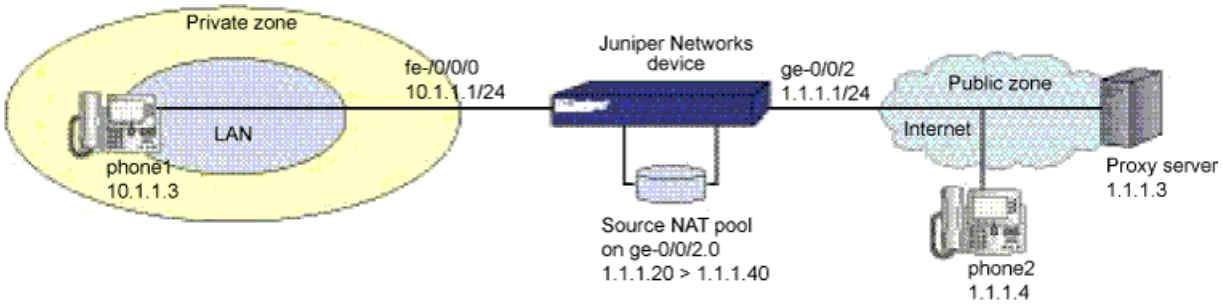

To configure a source NAT pool for incoming calls, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 218
- CLI Configuration on page 222
- Related Topics on page 223

J-Web Configuration

To configure interfaces:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Interfaces, click **Configure** or **Edit**.
3. Next to Interface, click **Add new entry**.
4. In the Interface name box, type **ge-0/0/0**.
5. Next to Unit, click **Add new entry**.
6. In the Interface unit number box, type **0**.
7. Under Family, select inet and click **Configure**.
8. Next to Address, click **Add new entry**.
9. To specify the source address, next to Source box, type **10.1.1.1/24** and click **OK**.
10. To configure other interface, ge-0/0/2, and to add address, repeat Step 2 through Step 9 and click **OK**.
11. To save and commit the configuration, click **Commit**.

To configure a private zone and assign an interface to it:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zones, click **Add new entry**.
5. In the Name box, type **private**.
6. Next to Interfaces, click **Add new entry**.
7. In the Interface unit box, type **ge-0/0/0.0** and click **OK**.
8. To save and commit the configuration, click **Commit**.

To configure addresses:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**
3. Next to Zones, click **Configure**.
4. Next to Security zones, click **Add new entry**.
5. In the Name box, type **private**.
6. Next to Address book click **Configure**.
7. Next to Address, click **Add new entry**.
8. In the Address name box, type **phone1 10.1.1.3/32** and click **OK**.
9. To configure more security zones, public, and address books entries such as proxy 10.1.1.3/32 and phone2 1.1.1.4/32, repeat Step 3 through Step 7and click **OK**.
10. To save and commit the configuration, click **Commit**.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zones, click **Add new entry**.
5. In the Name box, type **private** and click **OK**.
6. To specify the name of the another security zone, click **Add new entry** next to Security zone.
7. Next to the Name box, type **public** and click **OK**.
8. To configure an interface to the private zone, click **private**.
9. Next to Interfaces, click **Add new entry**.
10. In the Interface unit box, type **ge-0/0/0.0** and click **OK**.
11. To configure an interface to the public zone, click **public**.
12. Next to Interfaces, click **Add new entry**.
13. In the Interface unit box, type **ge-0/0/2.0** and click **OK**.
14. To save and commit the configuration, click **Commit**.

To configure source NAT pool:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Nat, click **Configure**.
4. Next to Source, click **Configure**.
5. Next to Address persistent, select the check box and click **OK**.
6. Next to Interface, click **Add new entry**.
7. In the Name box, type **ge-0/0/2.0**.
8. Next to Source nat, click **Configure**.
9. Next to Pool, click **Add new entry**.
10. In the Name box, type **sip-pool**.
11. Next to Address range, click **Add new entry**.
12. Next to High box, type **1.1.1.60** and next to Low box, type **1.1.1.20** and click **OK**.
13. Next to Allow incoming, select the check box and click OK.
14. To save and commit the configuration, click **Commit**.

To configure policies:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.

3. Next to Policies, select the check box and click **Configure**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type **private**.
6. In the To zone name box, type **public**.
7. Next to Policy, click **Add new entry**.
8. In the Policy name box, type **outgoing**.
9. Select the **Match** check box.
10. Select the **Then** check box.
11. Next to Match check box, click **Configure**.
12. From the Source address choice list, select **Source address**.
13. Next to Source address, click **Add new entry**.
14. From the Value keyword list, select **Enter Specific Value**.
15. In the Address box, type **phone1** and click **OK**.
16. From the Destination address choice list, select **Destination address**.
17. Next to Destination address, click **Add new entry**.
18. Next to Value keyword list, select **any** and click **OK**.
19. From the Application choice list, select **Application**.
20. Next to Application, click **Add new entry**.
21. In the Value keyword box, type **junos-sip** and click **OK**.
22. Next to Then, click **Configure**.
23. Next to Action, select **permit**.
24. Click **Configure** next to Permit.
25. Next to Source nat, select the check box and click **Configure**.
26. From the Source nat choice list, select interface and click **OK**.
27. In the From zone name box, type **private**.
28. In the To zone name box, type **public**.
29. Next to Policy, click **Add new entry**.
30. In the Policy name box, type **incoming**.
31. Select the **Match** check box.
32. Select the **Then** check box.
33. Click **Configure** next to Match check box.
34. Next to Source address choice list select **Source address**.
35. Next to Source address, click **Add new entry**.
36. From the Value keyword list, select **any** and click **OK**.
37. From the Destination address choice list, select **Destination address**.

38. Next to Destination address, click **Add new entry**.
39. Next to Value keyword list, select **Enter Specific Value**.
40. To specify the address, type **incoming-nat-fe0/0/2.0** and click **OK**.
41. From the Application choice list, select **Application**.
42. Next to Application, click **Add new entry**.
43. Next to Value keyword box, type **junos-sip** and click **OK**.
44. Next to Then, click **Configure**.
45. Next to Action, select permit.
46. To save and commit the configuration, click **Commit**.

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
user@host# set security zones security-zone private interface ge-0/0/0.0
```

2. Configure addresses.

```
user@host# set security zones security-zone private address-book address phone1
   10.1.1.3/32
user@host# set security zones security-zone public address-book address proxy
   10.1.1.3/32
user@host# set security zones security-zone public address-book address phone2
   1.1.1.4/32
```

3. Configure zones.

```
user@host# set security zones security-zone private
user@host# set security zones security-zone public
user@host# set security zones security-zone private interfaces ge-0/0/0.0
user@host# set security zones security-zone public interfaces ge-0/0/2.0
```

4. Configure the source NAT pool.

```
user@host# set security nat source-nat address-persistent
user@host# set security nat interface ge-0/0/2.0 source-nat pool sip-pool
   address-range low 1.1.1.20 high 1.1.1.60
user@host# set security nat interface ge-0/0/2.0 source-nat pool sip-pool allow
   incoming
```

5. Configure policies.

```
user@host# set security policies from-zone private to-zone public policy outgoing
   match source-address phone1 destination-address any application junos-sip
user@host# set security policies from-zone private to-zone public policy outgoing
   then permit source-nat pool sip-pool
```

```
user@host# set security policies from-zone private to-zone public policy incoming
    match source-address any destination-address incoming-nat-sip-pool application
    junos-sip
user@host# set security policies from-zone public to-zone private policy incoming
    then permit
```

Related Topics

- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213
- Configuring Interface Source NAT for Incoming SIP Calls on page 216
- Configuring Static NAT for Incoming SIP Calls on page 223
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring the SIP Proxy in the Public Zone on page 230
- Configuring a Three-Zone SIP Scenario on page 235

Configuring Static NAT for Incoming SIP Calls

When you locate the SIP proxy server in an external, or public, zone, static NAT configured on the interface to the public will enable callers in the internal, or private, zone to register with the proxy.

For background information, read

- "Understanding the SIP ALG" on page 187
- "SIP ALG Request Methods Overview" on page 191
- "Understanding SIP with Network Address Translation (NAT)" on page 204
- "Understanding Incoming SIP Call Support Using the SIP Registrar" on page 213

In this example, phone1 is on the ge-0/0/0 interface in the private zone, and phone2 and the proxy server are on the ge-0/0/2 interface in the public zone. You configure static NAT on the ge-0/0/2.0 interface to phone1, then create policies that allow SIP traffic from the public zone to the private zone, and reference the static NAT in the policy. This example is similar to the ("Configuring Interface Source NAT for Incoming SIP Calls" on page 216 and "Configuring a Source NAT Pool for Incoming SIP Calls" on page 218, except that with static NAT you need one public address for each private address in the private zone, while with a DIP pool a single interface address can serve multiple private addresses. See Figure 28 on page 224.

Figure 28: Static NAT for Incoming Calls

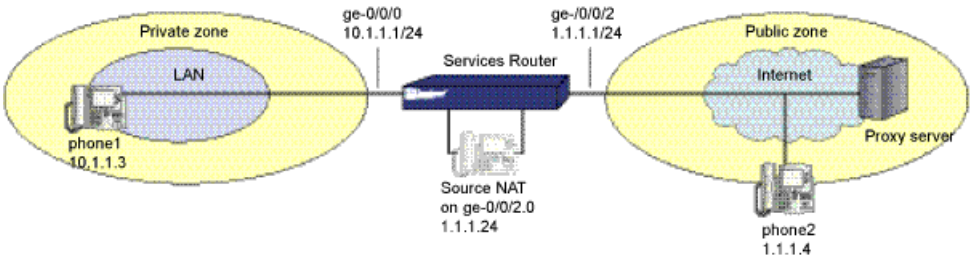

To configure static NAT for incoming calls, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 224
- CLI Configuration on page 227
- Related Topics on page 227

J-Web Configuration

To configure interfaces:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Interfaces, click Configure or Edit.
3. Next to Interface, click Add new entry.
4. In the Interface name box, type **ge-0/0/0**.
5. Next to Unit, click Add new entry.
6. In the Interface unit number box, type **0**.
7. Under Family, select inet and click Configure.
8. Next to Address, click Add new entry.
9. In the Source box, type **10.1.1.1/24** and click OK.
10. To configure other interface, ge-0/0/2, and to add address, repeat Step 2 through Step 9, and click OK.
11. To save and commit the configuration, click Commit.

To configure a zone and assign an interface:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zones, click Add new entry.
5. In the Name box, type **private**.
6. Next to Interfaces, click Add new entry.

7. In the Interface unit box, type **ge-0/0/0.0** and click OK.
8. To save and commit the configuration, click Commit.

To configure addresses:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zones, click Add new entry.
5. In the Name box, type **private**.
6. Next to Address book, click Configure.
7. Next to Address, click Add new entry.
8. In the Address name box, type **phone1 10.1.1.3/32**.
9. To configure more security zones, public, and address books entries such as proxy 10.1.1.3/32 and phone2 1.1.1.4/32, repeat Step 3 through Step 8 and click OK.
10. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zones, click Add new entry.
5. In the Name box, type **private** and click OK.
6. To specify the name of the another security zone, next to Security zone, click Add new entry.
7. In the Name box, type **public** and click OK.
8. To configure an interface to the private zone, click private.
9. Next to Interfaces, click Add new entry.
10. In the Interface unit box, type **ge-0/0/0.0** and click OK.
11. To configure an interface to the public zone, click public.
12. Next to Interfaces, click Add new entry.
13. In the Interface unit box, type **ge-0/0/0.0** and click OK.
14. To configure an interface to the private zone, click private.
15. Next to Interfaces, click Add new entry.
16. In the Interface unit box, type **ge-0/0/2.0** and click OK.
17. To save and commit the configuration, click Commit.

To configure static NAT:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click Configure or Edit.
3. Next to Nat, click Configure or Edit.
4. Next to Interfaces, click Add new entry.
5. In the Name box, type **ge-0/0/2.0** and click OK.
6. Under the Name column, click ge-0/0/2.0.
7. Next to Static nat, click Add new entry.
8. In the Address box, type **1.1.1.3/32**.
9. In the Host box, type **10.1.1.3/32** and click OK.
10. To save and commit the configuration, click Commit.

To configure policies:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **public**.
6. In the To zone name box, type **private** and click OK.
7. Under the From zone name column, click public.
8. Next to Policy, click Add new entry.
9. In the Policy name box, type **incoming**.
10. Select the Match check box.
11. Select the Then check box.
12. Next to Match, click Configure.
13. Next to Source address, select Source address.
14. Next to Source address, click Add new entry.
15. From the Value keyword list, select any and click OK.
16. From the Destination address choice list, select Destination address.
17. Next to Destination address, click Add new entry.
18. From the Value keyword list, select Enter Specific Value.
19. In the Address box, type **static_nat_1.1.1.3-32** and click OK.
20. From the Application choice list, select Application.
21. Next to Application, click Add new entry.
22. In the Value keyword box, type **junos-jsrp** and click OK.

23. Next to Then, click Configure.
24. Next to Action, select permit and click OK.
25. To save and commit the configuration, click Commit.

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
user@host# set security zones security-zone private interface ge-0/0/0.0
```

2. Configure addresses.

```
user@host# set security zones security-zone private address-book address phone1
   10.1.1.3/32
user@host# set security zones security-zone public address-book address proxy
   10.1.1.3/32
user@host# set security zones security-zone public address-book address phone2
   1.1.1.4/32
```

3. Configure zones.

```
user@host# set security zones security-zone private
user@host# set security zones security-zone public
user@host# set security zones security-zone private interfaces ge-0/0/0.0
user@host# set security zones security-zone public interfaces ge-0/0/2.0
```

4. Configure static NAT.

```
user@host# set security nat interface ge-0/0/2.0 static-nat 1.1.1.3/32 host
   10.1.1.3/32
```

5. Configure Policies.

```
user@host# set security policies from-zone public to-zone private policy incoming
   match source-address any destination-address static_nat_1.1.1.3-32 application
   junos-jsrp
user@host# set security policies from-zone public to-zone private policy incoming
   then permit
```

Related Topics

- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213
- Configuring Interface Source NAT for Incoming SIP Calls on page 216
- Configuring a Source NAT Pool for Incoming SIP Calls on page 218
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring a Three-Zone SIP Scenario on page 235

Configuring the SIP Proxy in the Private Zone

With the SIP proxy server in the internal, or private, zone, static NAT on the interface to the external, or public, zone is sufficient to allow callers in the public zone to register with the proxy server.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213

In this example, phone1 and the SIP proxy server are on the ge-0/0/0 interface in the private zone, and phone2 is on the ge-0/0/2 interface in the public zone. You configure static NAT on the ge-0/0/2 interface to the proxy server to allow phone2 to register with the proxy, then create a policy allowing SIP traffic from the public to the private zone to enable callers in the public zone to register with the proxy, and a policy from the private to the public zone to allow phone1 to call out. See Figure 29 on page 228.

Figure 29: Proxy in the Private Zone

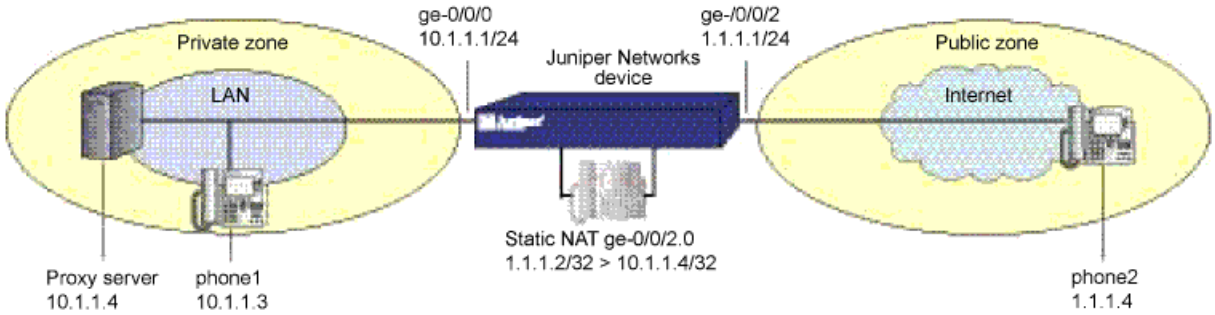

Use either the J-Web or CLI configuration editor.

This topic covers:

- CLI Configuration on page 228
- Related Topics on page 229

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

2. Configure zones.

```
user@host# set security zones security-zone private
user@host# set security zones security-zone public
user@host# set security zones security-zone private interfaces ge-0/0/0.0
user@host# set security zones security-zone public interfaces ge-0/0/2.0
```

3. Configure addresses.

```
user@host# set security zones security-zone private address-book address phone1
  10.1.1.3/32
user@host# set security zones security-zone private address-book address proxy
  10.1.1.4/32
user@host# set security zones security-zone public address-book address phone2
  1.1.1.4/32
```

4. Configure static-NAT.

```
user@host# set security nat interface ge-0/0/2.0 static-nat 1.1.1.2/32 host
  10.1.1.4/32
```

5. Configure policies.

```
user@host# set security policies from-zone private to-zone public policy outgoing
  match source-address any
user@host# set security policies from-zone private to-zone public policy outgoing
  match destination-address phone2
user@host# set security policies from-zone private to-zone public policy outgoing
  match application junos-sip
user@host# set security policies from-zone private to-zone public policy outgoing
  then permit source-nat interface
user@host# set security policies from-zone public to-zone private policy incoming
  match source-address phone2
user@host# set security policies from-zone public to-zone private policy incoming
  match destination-address static_nat_1.1.1.2_32
user@host# set security policies from-zone public to-zone private policy incoming
  match application junos-sip
user@host# set security policies from-zone public to-zone private policy incoming
  then permit
```

Related Topics

- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213
- Configuring a Source NAT Pool for Incoming SIP Calls on page 218
- Configuring Static NAT for Incoming SIP Calls on page 223
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring a Three-Zone SIP Scenario on page 235
- Configuring Static NAT for Incoming SIP Calls on page 223

Configuring the SIP Proxy in the Public Zone

When you locate the SIP proxy server in an external, or public, zone, you will typically want to configure NAT on the interface to that zone.

Before You Begin

For background information, read

- Understanding the SIP ALG on page 187
- SIP ALG Request Methods Overview on page 191
- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213

In this example, phone1 is on the ge-0/0/0.0 interface in the private zone, and the proxy server and phone2 are on the ge-0/0/2.0 interface in the public zone. You configure source NAT on interface 0/0/2.0 in the public zone, then create a policy permitting SIP traffic from the public zone to the private zone and reference the NAT interface. You also create a policy from private to public to allow phone1 to register with the proxy server in the public zone. See Figure 30 on page 230.

Figure 30: Proxy in the Public Zone

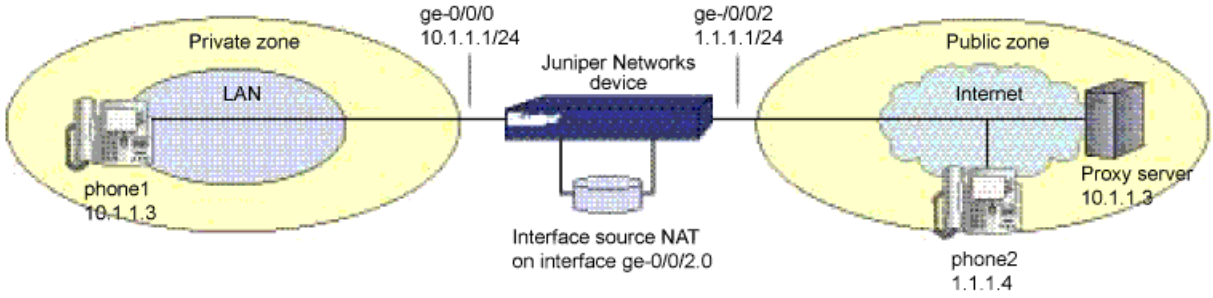

To configure the SIP proxy in the public zone, use either the J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 230
- CLI Configuration on page 234
- Related Topic on page 235

J-Web Configuration

To configure interfaces:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Interfaces, click Configure or Edit.
3. Next to Interface, click Add new entry.
4. In the Interface name box, type **ge-0/0/0**.
5. Next to Unit, click Add new entry.
6. Next to Interface unit number, type **0**.
7. Next to Inet, select the check box and click Configure.
8. Next to Address, click Add new entry.
9. Next to Source, type **10.1.1.1/24** and click OK.
10. To configure another interface, ge-0/0/2, and address, 1.1.1.1/24, repeat Step 2 through Step 9 and click OK.
11. To save and commit the configuration, click Commit.

To configure zones:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zones, click Add new entry.
5. In the Name box, type **private** and click OK.
6. Next to Security zone, click Add new entry.
7. In the Name box, type **public** and click OK.
8. To configure an interface to the private zone, click private.
9. Next to Interfaces, click Add new entry.
10. Next to Interface unit box, type **ge-0/0/0.0** and click OK.
11. To configure an interface to the public zone, click public.
12. Next to Interfaces, click Add new entry.
13. Next to Interface unit box, type **ge-0/0/2.0** and click OK.
14. To save and commit the configuration, click Commit.

To configure addresses:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **private**.
6. Next to Address book, click Configure or Edit.

7. Next to Address, click Add new entry.
8. In the Address name box, type **phone1 10.1.1.3/32** and click OK.
9. To configure another security zone public and address books entries such as phone2 1.1.1.4/32 and proxy 1.1.1.3/32, repeat Step 4 through Step 8 and click OK.
10. To save and commit the configuration, click Commit.

To configure interface Source-Nat:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Nat, click Configure.
4. Next to Source, click Configure.
5. Next to Address persistent, select the check box and click OK.
6. Next to Interface, click Add new entry.
7. In the Name box, type **ge-0/0/2.0**.
8. Next to Allow incoming, select the check box and click OK.
9. To save and commit the configuration, click Commit.

To configure policies:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **private**.
6. In the To zone name box, type **public** and click OK.
7. Under the From zone name column, click private.
8. Next to Policy, click Add new entry.
9. In the Policy name box, type **outgoing**.
10. Select the Match check box.
11. Select the Then check box.
12. Next to Match, click Configure.
13. Next to Source address, select Source address.
14. Next to Source address, click Add new entry.
15. From the Value keyword list, select Enter Specific Value.
16. In the Address box, type **phone1** and click OK.
17. From the Destination address choice list, select Destination address.

18. Next to Destination address, click Add new entry.
19. From the Value keyword list, select any and click OK.
20. From the Application choice list, select Application.
21. Next to Application, click Add new entry.
22. In the Value keyword box, type **junos-sip** and click OK.
23. Next to Then, click Configure.
24. Next to Action, select permit and click OK.
25. Next to Permit, click Configure.
26. Select the Source Nat check box, and click Configure.
27. From the Source nat list, select Interface and click OK.
28. To save and commit the configuration, click Commit.

To configure another policy From-zone, public, and To zone, private, follow the sequence of steps listed below:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **public**.
6. In the To zone name box, type **private** and click OK.
7. Under the From zone name column, click public.
8. Next to Policy, click Add new entry.
9. In the Policy name box, type **incoming**.
10. Select the Match check box.
11. Select the Then check box.
12. Next to Match, click Configure.
13. Next to Source address, select Source address.
14. Next to Source address, click Add new entry.
15. From the Value keyword list, select any and click OK.
16. In the Address box, type **phone2** and click OK.
17. From the Destination address choice list, select Destination address.
18. Next to Destination address, click Add new entry.
19. From the Value keyword list, select Enter Specific Value.
20. In the Address box, type **incoming_nat_ge-0/0/2.0** and click OK.
21. From the Application choice list, select Application.

22. Next to Application, click Add new entry.
23. In the Value keyword box, type **junos-sip** and click OK.
24. Next to Then, click Configure.
25. Next to Action, select permit and click OK.
26. To save and commit the configuration, click Commit.

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

2. Configure zones.

```
user@host# set security zones security-zone private
user@host# set security zones security-zone public
user@host# set security zones security-zone private interfaces ge-0/0/0.0
user@host# set security zones security-zone public interfaces ge-0/0/2.0
```

3. Configure addresses.

```
user@host# set security zones security-zone private address-book address phone1
  10.1.1.3/32
user@host# set security zones security-zone public address-book address phone2
  1.1.1.4/32
user@host# set security zones security-zone public address-book address proxy
  1.1.1.3/32
```

4. Configure interface source-Nat.

```
user@host# set security nat source-nat address-persistent
user@host# set security nat interface ge-0/0/2.0 allow-incoming
```

5. Configure policies.

```
user@host# set security policies from-zone private to-zone public policy outgoing
  match source-address phone1
user@host# set security policies from-zone private to-zone public policy outgoing
  match destination-address any
user@host# set security policies from-zone private to-zone public policy outgoing
  match application junos-sip
user@host# set security policies from-zone private to-zone public policy outgoing
  then permit source-nat interface
user@host# set security policies from-zone public to-zone private policy incoming
  match source-address any
user@host# set security policies from-zone public to-zone private policy incoming
  match destination-address incoming_nat_ge-0/0/2.0
user@host# set security policies from-zone public to-zone private policy incoming
  match application junos-sip
user@host# set security policies from-zone public to-zone private policy incoming
  then permit
```

Related Topic

- "Understanding SIP with Network Address Translation (NAT)" on page 204
- "Understanding Incoming SIP Call Support Using the SIP Registrar" on page 213
- "Configuring Interface Source NAT for Incoming SIP Calls" on page 216
- "Configuring a Source NAT Pool for Incoming SIP Calls" on page 218
- "Configuring Static NAT for Incoming SIP Calls" on page 223
- "Configuring the SIP Proxy in the Private Zone" on page 228
- "Configuring a Three-Zone SIP Scenario" on page 235

Configuring a Three-Zone SIP Scenario

In a three-zone SIP configuration, the SIP proxy server is typically in a different zone than the calling and called parties. Such a scenario requires additional address and zone configuration, and policies to ensure that all parties have access to each other and to the proxy server.

Before You Begin

For background information, read

- "Understanding the SIP ALG" on page 187.
- "SIP ALG Request Methods Overview" on page 191.
- "Understanding SIP with Network Address Translation (NAT)" on page 204.
- "Understanding Incoming SIP Call Support Using the SIP Registrar" on page 213.

In this example, phone1 is on the ge-0/0/0 interface in the private zone, phone2 is on the ge-0/0/0/2 interface in the public zone, and the proxy server is on the ge-0/0/1.0 interface in the DMZ. You configure static NAT on the ge-0/0/1 interface to phone1 in the private zone. You then create policies from the private zone to the DMZ and from the DMZ ot the private zone, from the public zone to the DMZ and from the DMZ to the public zone, and from the pirvate zone to the public zone. The arrows in Figure 31 on page 236 show the flow of SIP signaling traffic when phone2 in the public zone places a call to phone1 in the private zone. After the session is initiated, the media flows directly between phone1 and phone2.

Figure 31: Three-Zone, Proxy in the DMZ

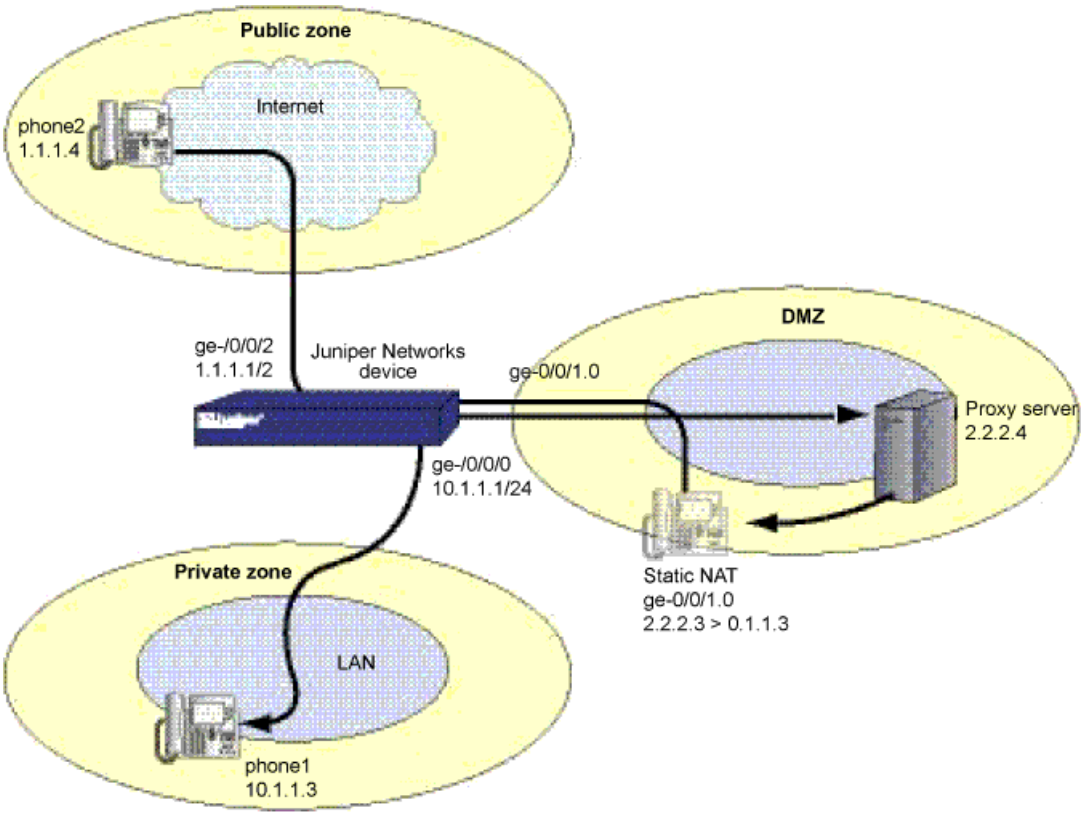

To configure a three-zone SIP scenario, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 236
- CLI Configuration on page 242
- Related Topics on page 243

J-Web Configuration

To configure interfaces:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Interfaces, click Configure or Edit.
3. Next to Interface, click Add new entry.
4. In the Interface name box, type **ge-0/0/0**.
5. Next to Unit, click Add new entry.
6. Next to Interface unit number, type **0**.
7. Next to Inet, select the check box and click Configure.

8. Next to Address, click Add new entry.
9. Next to Source, type **10.1.1.1/24** and click OK.
10. To configure another interface, ge-0/0/1 and ge-0/0/2, and addresses, 2.2.2.2/24 and 1.1.1.1/24, repeat Step 2 through Step 9 and click OK.
11. To save and commit the configuration, click Commit.

To configure zones:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zones, click Add new entry.
5. In the Name box, type **private** and click OK.
6. Next to Security zone, click Add new entry.
7. In the Name box, type **public** and click OK.
8. Next to Security zone, click Add new entry.
9. In the Name box, type **dmz** and click OK.
10. To configure an interface to the private zone, click private.
11. Next to Interfaces, click Add new entry.
12. Next to Interface unit box, type **ge-0/0/0.0** and click OK.
13. To configure an interface to the public zone, click public.
14. Next to Interfaces, click Add new entry.
15. Next to Interface unit box, type **ge-0/0/2.0** and click OK.
16. To configure an interface to the dmz, click dmz.
17. Next to Interfaces, click Add new entry.
18. Next to Interface unit box, type **ge-0/0/1.0** and click OK.
19. To save and commit the configuration, click Commit.

To configure addresses:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **private**.
6. Next to Address book, click Configure.
7. Next to Address, click Add new entry.
8. In the Address name box, type **phone1 10.1.1.3/32** and click OK.

9. To configure more security zones, public and dmz, and address books entries such as phone2 1.1.1.4/32 and proxy 2.2.2.4/32, repeat Step 4 through Step 8, and click OK.
10. To save and commit the configuration, click Commit.

To configure static NAT:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Nat, click Configure.
4. Next to Interfaces, click Add new entry.
5. In the Name box, type **ge-0/0/1.0**.
6. Next to Static nat, click Add new entry.
7. In the Address box, type **2.2.2.3/32**.
8. In the Host box, type **10.1.1.3/32** and click OK.
9. To save and commit the configuration, click Commit.

To configure policies:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **private**.
6. In the To zone name box, type **dmz**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **private**-to-proxy.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match check box, click Configure.
12. From the Source address choice list, select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, select Enter Specific Value.
15. In the Address box, type **phone1** and click OK.
16. From the Destination address choice list, select Destination address.
17. Next to Destination address, click Add new entry.
18. Next to Value keyword list, select proxy and click OK.
19. From the Application choice list, select Application.

20. Next to Application, click Add new entry.
21. In the Value keyword box, type **junos-sip** and click OK.
22. Next to Then, click Configure.
23. Next to Action, select permit.
24. Click Configure next to Permit.
25. Next to Source nat, select the check box and click Configure.
26. From the Source nat choice list, select interface and click OK.

To configure from zone, public, and to zone, dmz, and the respective source address, destination address, and application:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry
5. In the From zone name box, type **public**.
6. In the To zone name box, type **private**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **public**-to-proxy.
9. Select the Match check box.
10. Select the Then check box.
11. Click Configure next to Match check box.
12. Next to Source address choice list select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, type **phone2** and click OK.
15. From the Destination address choice list, select Destination address.
16. Next to Destination address, click Add new entry.
17. Next to Value keyword list, select Enter Specific Value.
18. To specify the address, type **proxy** and click OK.
19. From the Application choice list, select Application.
20. Next to Application, click Add new entry.
21. Next to Value keyword box, type **junos-sip** and click OK.
22. Next to Then, click Configure.
23. Next to Action, select permit and click OK.

To configure from zone, private, and to zone, public, and the respective source address, destination address, and application:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **private**.
6. In the To zone name box, type **public**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **private**-to-public.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match check box, click Configure.
12. From the Source address choice list, select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, select Enter Specific Value.
15. In the Address box, type **phone1** and click OK.
16. From the Destination address choice list, select Destination address.
17. Next to Destination address, click Add new entry.
18. Next to Value keyword list, type **phone2** and click OK.
19. From the Application choice list, select Application.
20. Next to Application, click Add new entry.
21. In the Value keyword box, type **junos-sip** and click OK.
22. Next to Then, click Configure.
23. Next to Action, select permit.
24. Click Configure next to Permit.
25. Next to Source nat, select the check box and click Configure.
26. From the Source nat choice list, select interface and click OK.

To configure from zone, dmz, and to zone, private, and the respective source address, destination address, and application:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry
5. In the From zone name box, type **dmz**.
6. In the To zone name box, type **private**.

7. Next to Policy, click Add new entry.
8. In the Policy name box, type **proxy-to-private**.
9. Select the Match check box.
10. Select the Then check box.
11. Click Configure next to Match check box.
12. Next to Source address choice list select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, type **proxy** and click OK.
15. From the Destination address choice list, select Destination address.
16. Next to Destination address, click Add new entry.
17. Next to Value keyword list, select Enter Specific Value.
18. To specify the address, type **static_nat_2.2.2.3_32** and click OK.
19. From the Application choice list, select Application.
20. Next to Application, click Add new entry.
21. Next to Value keyword box, type **junos-sip** and click OK.
22. Next to Then, click Configure.
23. Next to Action, select permit and click OK.

To configure from zone, dmz, and to zone, public, and the respective source address, destination address, and application:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry
5. In the From zone name box, type **dmz**.
6. In the To zone name box, type **public**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **proxy-to-public**.
9. Select the Match check box.
10. Select the Then check box.
11. Click Configure next to Match check box.
12. Next to Source address choice list select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, type **proxy** and click OK.
15. From the Destination address choice list, select Destination address.

16. Next to Destination address, click Add new entry.
17. Next to Value keyword list, select Enter Specific Value.
18. To specify the address, type **phone2** and click OK.
19. From the Application choice list, select Application.
20. Next to Application, click Add new entry.
21. Next to Value keyword box, type **junos-sip** and click OK.
22. Next to Then, click Configure.
23. Next to Action, select **permit** and click OK.
24. To save and commit the configuration, click Commit.
25. To check the configuration, see "Verifying the SIP Configuration" on page 522.

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

2. Configure zones.

```
user@host# set security zones security-zone private interfaces ge-0/0/0.0
user@host# set security zones security-zone public interfaces ge-0/0/2.0
user@host# set security zones security-zone dmz interfaces ge-0/0/1.0
```

3. Configure addresses.

```
user@host# set security zones security-zone private address-book address phone1
   10.1.1.3/32
user@host# set security zones security-zone public address-book address phone2
   1.1.1.4/32
user@host# set security zones security-zone dmz address-book address proxy
   2.2.2.4/32
```

4. Configure static-NAT.

```
user@host# set security nat interface ge-0/0/1.0 static-nat 2.2.2.3/32 host
   10.1.1.3/32
```

5. Configure policies.

```
user@host# set security policies from-zone private to-zone dmz policy
   private-to-proxy match source-address phone1
user@host# set security policies from-zone private to-zone dmz policy
   private-to-proxy match destination-address proxy
user@host# set security policies from-zone private to-zone dmz policy
   private-to-proxy match application junos-sip
user@host# set security policies from-zone private to-zone dmz policy
   private-to-proxy then permit source-nat interface
```

```
user@host# set security policies from-zone public to-zone dmz policy public-to-proxy
    match source-address phone2
user@host# set security policies from-zone public to-zone dmz policy public-to-proxy
    match destination-address proxy
user@host# set security policies from-zone public to-zone dmz policy public-to-proxy
    match application junos-sip
user@host# set security policies from-zone public to-zone dmz policy public-to-proxy
    then permit
user@host# set security policies from-zone private to-zone public policy
    private-to-public match source-address phone1
user@host# set security policies from-zone private to-zone public policy
    private-to-public match destination-address phone2
user@host# set security policies from-zone private to-zone public policy
    private-to-public match application junos-sip
user@host# set security policies from-zone private to-zone public policy
    private-to-public then permit source-nat interface
user@host# set security policies from-zone dmz to-zone private policy
    proxy-to-private match source-address proxy
user@host# set security policies from-zone dmz to-zone private policy
    proxy-to-private match destination-address static_nat_2.2.2.3_32
user@host# set security policies from-zone dmz to-zone private policy
    proxy-to-private match application junos-sip
user@host# set security policies from-zone dmz to-zone private policy
    proxy-to-private then permit
user@host# set security policies from-zone dmz to-zone public policy proxy-to-public
    match source-address proxy
user@host# set security policies from-zone dmz to-zone public policy proxy-to-public
    match destination-address phone2
user@host# set security policies from-zone dmz to-zone public policy proxy-to-public
    match application junos-sip
user@host# set security policies from-zone dmz to-zone public policy proxy-to-public
    then permit
```

Related Topics

- Understanding SIP with Network Address Translation (NAT) on page 204
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 213
- Configuring Interface Source NAT for Incoming SIP Calls on page 216
- Configuring a Source NAT Pool for Incoming SIP Calls on page 218
- Configuring Static NAT for Incoming SIP Calls on page 223
- Configuring the SIP Proxy in the Private Zone on page 228
- Configuring the SIP Proxy in the Public Zone on page 230

Verifying the SIP Configuration

To verify the SIP configuration, perform these tasks:

- "Verifying SIP Calls" on page 244
- "Verifying SIP Call Detail" on page 245

- "Verifying SIP Transactions" on page 246
- "Verifying SIP Counters" on page 247
- "Verifying the Rate of SIP Messages" on page 248

Verifying the SIP ALG

Purpose Display SIP verification options.

Action From the CLI, enter the show security alg sip ? command.

```
user@host> show security alg sip ?
Possible completions:
  calls                Show SIP calls
  counters             Show SIP counters
  rate                 Show SIP rate
  transaction          Show SIP transactions
```

What it Means The output shows a list of all SIP verification parameters. Verify the following information:

- Calls—Lists all SIP calls.
- Counters—Provides counters of response codes for each SIP request method and error type.
- Rate—Provides speed and periodicity of SIP signaling messages.
- Transaction—provides passage time for last message, for all messages, and the rate at which messages transit the network.

Related Topics

- Understanding the SIP ALG on page 187
- Verifying SIP Calls on page 244
- Verifying SIP Call Detail on page 245
- Verifying SIP Transactions on page 246
- Verifying SIP Counters on page 247
- Verifying the Rate of SIP Messages on page 248

Verifying SIP Calls

Purpose Display information about active calls.

Action From the J-Web interface, select **Monitor > ALGs > SIP > Calls**. Alternatively, from the CLI, enter the show security alg sip calls command.

```
user@host> show security alg sip calls
Total number of calls: 1
    Call leg 1 zone: 1
     UAS call ID: 000ed748-5533005e-11d97de9-77759865@10.10.10.254 (pending
```

```
transactions 0)
     Local tag  : 000f90542e7e005c7807c3a0-0647d41e
     Remote tag : 000ed748553300724648b313-2ccd0d9a
     State      : STATE_ESTABLISHED
    Call leg 2 zone: 1
     UAC call ID: 000ed748-5533005e-11d97de9-77759865@10.10.10.254 (pending
transactions 0)
     Local tag  : 000ed748553300724648b313-2ccd0d9a
     Remote tag : 000f90542e7e005c7807c3a0-0647d41e
     State      : STATE_ESTABLISHED
```

What it Means The output shows a list of all active SIP calls. Verify the following information:

- Each call leg and associated zone.
- The User Agent Server (UAS) call ID and local and remote tags, and the state of the call.

Related Topics

- Understanding the SIP ALG on page 187
- Verifying the SIP ALG on page 244
- Verifying SIP Call Detail on page 245
- Verifying SIP Transactions on page 246
- Verifying SIP Counters on page 247
- Verifying the Rate of SIP Messages on page 248

Verifying SIP Call Detail

Purpose Display address and Session Description Protocol (SDP) about active calls.

Action From the J-Web interface, select **Monitor > ALGs > SIP > Calls**. Alternatively, from the CLI, enter the show security alg sip calls detail command.

```
user@host> show security alg sip calls detail
Total number of calls: 1
       Call ID    : 000ed748-5533005e-11d97de9-77759865@10.10.10.254
       Local tag  : 000f90542e7e005c7807c3a0-0647d41e
       Remote tag : 000ed748553300724648b313-2ccd0d9a
       State      : STATE_ESTABLISHED               RM Group:2048
                    Local Info              Remote Info            RM Info
                    ----------              -----------            -------
       Call leg 1:  IP             Port     IP             Port     RSC id

         Host       10.10.10.10    5060     10.10.10.10    5060
         Contact    10.10.10.100   5060     10.10.10.254   1025     8191
         Contact maddr-                      -
         SDP:c      10.10.10.100             10.10.10.254
         SDP:m      10.10.10.100   18902    10.10.10.254   64518 8192 ,8185
       Call leg 2:
         Host       10.10.10.10    5060     10.10.10.10    5060
         Contact    10.10.10.254   1025     10.10.10.100   5060     8188
         Contact maddr-                      -
```

```
        SDP:c       10.10.10.254             10.10.10.100
        SDP:m       10.10.10.254   64518     10.10.10.100   18902 8186 ,8187
```

What it Means The output provides detail about all active SIP calls. Verify the following information:

- The total number of calls, their ID and tag information, and state
- Remote group ID
- The IP addresses and port numbers for each call leg, and Session Description Protocol (SDP) connection and media details

Related Topics

- Understanding the SIP ALG on page 187
- Verifying the SIP ALG on page 244
- Verifying SIP Calls on page 244
- Verifying SIP Transactions on page 246
- Verifying SIP Counters on page 247
- Verifying the Rate of SIP Messages on page 248

Verifying SIP Transactions

Purpose Display information about SIP transactions.

Action From the J-Web interface, select **Monitor > ALGs > SIP > Transactions**. Alternatively, from the CLI, enter the show security alg sip transactions command.

```
user@host> show security alg sip transaction
Total number of transactions: 1
       Transaction Name   Method    CSeq      State        Timeout  VIA RSC ID
       UAS:tsx0x4b06ddf4   BYE      101     Proceeding       -1        -
       UAC:tsx0x4b06f610   BYE      101     Calling          27        8185
```

What it Means The output provides a history of SIP call transactions. Verify the following information:

- The total number of transactions
- Information about the User Agent Client (UAC) and User Agent Server (UAS) for each transaction.

Related Topics

- Understanding the SIP ALG on page 187
- Verifying the SIP ALG on page 244
- Verifying SIP Calls on page 244
- Verifying SIP Call Detail on page 245

- Verifying SIP Counters on page 247
- Verifying the Rate of SIP Messages on page 248

Verifying SIP Counters

Purpose Display information about SIP counters.

Action From the J-Web interface, select **Monitor > ALGs > SIP > Counters**. Alternatively, from the CLI, enter the show security alg sip counters command.

```
user@host> show security alg sip counters
Method          T       1xx       2xx       3xx       4xx       5xx       6xx
               RT        RT        RT        RT        RT        RT        RT
INVITE          2         4         2         0         0         0         0
                1         1         0         0         0         0         0
CANCEL          0         0         0         0         0         0         0
                0         0         0         0         0         0         0
ACK             2         0         0         0         0         0         0
                0         0         0         0         0         0         0
BYE             2         0         1         0         0         0         0
                0         0         0         0         0         0         0
REGISTER        0         0         0         0         0         0         0
                0         0         0         0         0         0         0
OPTIONS         0         0         0         0         0         0         0
                0         0         0         0         0         0         0
INFO            0         0         0         0         0         0         0
                0         0         0         0         0         0         0
MESSAGE         0         0         0         0         0         0         0
                0         0         0         0         0         0         0
NOTIFY          0         0         0         0         0         0         0
                0         0         0         0         0         0         0
PRACK           0         0         0         0         0         0         0
                0         0         0         0         0         0         0
PUBLISH         0         0         0         0         0         0         0
                0         0         0         0         0         0         0
REFER           0         0         0         0         0         0         0
                0         0         0         0         0         0         0
SUBSCRIBE       0         0         0         0         0         0         0
                0         0         0         0         0         0         0
UPDATE          0         0         0         0         0         0         0
                0         0         0         0         0         0         0
BENOTIFY        0         0         0         0         0         0         0
                0         0         0         0         0         0         0
SERVICE         0         0         0         0         0         0         0
                0         0         0         0         0         0         0
OTHER           0         0         0         0         0         0         0
                0         0         0         0         0         0         0
    SIP Error Counters
    ------------------------------
       Total Pkt-in                     :15
       Total Pkt dropped on error       :0
       Transaction error                :0
       Call error                       :0
       IP resolve error                 :0
       NAT error                        :0
       Resource manager error           :0
       RR header exceeded max           :0
       Contact header exceeded max      :0
```

```
    Invite Dropped due to call limit    :0
    SIP msg not processed by stack      :0
    SIP msg not processed by alg        :0
    SIP unknown method dropped          :0
    Decoding error                      :0
    Request for disconnected call       :0
    Request out of state                :0
```

What it Means The output provides a count of all SIP response codes transmitted and received, and of SIP errors. Verify the following information:

- A count of transmissions of response codes for each SIP request method
- A count of all possible error types

Related Topics

- Understanding the SIP ALG on page 187
- Verifying the SIP ALG on page 244
- Verifying SIP Calls on page 244
- Verifying SIP Call Detail on page 245
- Related Topics on page 246
- Verifying the Rate of SIP Messages on page 248

Verifying the Rate of SIP Messages

Purpose Display information about SIP message rate.

Action From the J-Web interface, select **Monitor > ALGs > SIP > Rate**. Alternatively, from the CLI, enter the show security alg sip rate command.

```
user@host> show security alg sip rate
CPU ticks per us is 166
    Time taken for the last message is 1103 us
    Total time taken for 3124 messages is 6221482 us(in less than 10 minutes)
    Rate: 502 messages/second
```

What it Means The output provides information about CPU usage for messages, and speed and periodicity of SIP signaling messages. Verify the following information:

- CPU ticks per US
- Passage time for last message, for all messages, and the rate at which messages transit the network

Related Topics

- Understanding the SIP ALG on page 187
- Verifying the SIP ALG on page 244
- Verifying SIP Calls on page 244

- Verifying SIP Call Detail on page 245
- Verifying SIP Transactions on page 246
- Verifying SIP Counters on page 247

Understanding the SCCP ALG

The Skinny Client Control Protocol (SCCP) is a Cisco proprietary protocol for call signaling. Skinny is based on a call-agent-based call-control architecture. The control protocol uses binary-coded frames encoded on TCP frames sent to well-known TCP port number destinations to set up and tear down RTP media sessions.

The SCCP protocol just as other call control protocols, negotiates media endpoint parameters—specifically the RTP port number and the IP address of media termination—by embedding information in the control packets. The SCCP ALG implemented on a J Series device (or firewall) parses these control packets and facilitates media and control packets to flow through the J Series device.

The SCCP ALG also implements rate limiting of calls and helps protect critical resources from overloading and denial of service attacks.

The following functions are implemented by the SCCP ALG in JUNOS Software:

- Validation of SCCP protocol data units
- Translation of embedded IP address and port numbers
- Allocation of firewall resources (pinholes and gates) to pass media
- Aging out idle calls
- Configuration API for SCCP ALG parameters
- Operational mode API for displaying counters, status and statistics

Before You Begin

For background information, read

- Application Layer Gateways (ALGs) on page 157
- Pinhole Creation on page 190

In the SCCP architecture, a proxy, known as the CallManager, does most of the processing. IP phones, also called End Stations, run the SCCP client and connect to a primary (and, if available, a secondary) CallManager over TCP on port 2000 and register with the primary CallManager. This connection is then used to establish calls coming to or from the client.

The SCCP ALG supports the following:

- Call flow from a SCCP client, through the CallManager, to another SCCP client.
- Seamless failover—Switches over all calls in process to the standby firewall during failure of the primary.

- VoIP signaling payload inspection—Fully inspects the payload of incoming VoIP signaling packets. Any malformed packet attack is blocked by the ALG.
- SCCP signaling payload inspection—Fully inspects the payload of incoming SCCP signaling packets. Any malformed-packet attack is blocked by the ALG.
- Stateful processing—Invokes the corresponding VoIP-based state machines to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- Network Address Translation (NAT)—Translates any embedded IP address and port information in the payload, based on the existing routing information and network topology, with the translated IP address and port number, if necessary.
- Pinhole creation and management for VoIP traffic—Identifies IP address and port information used for media or signaling and dynamically opens (and closes) pinholes to securely stream the media.

This topic covers:

- SCCP Security on page 250
- SCCP Components on page 250
- SCCP Transactions on page 251
- SCCP Control Messages and RTP Flow on page 253
- SCCP Messages on page 254
- Related Topics on page 255

SCCP Security

The SCCP ALG includes the following security features:

- Stateful inspection of SCCP control messages over TCP and validation of the message format, and message validity for the current call state. Invalid messages are dropped.
- Security policy enforcement between Cisco IP phones and Cisco CallManager
- Protect against call flooding by rate limiting the number of calls processed by the ALG
- Seamless failover of calls, including the ones in progress in case of J Series device failure in a clustered deployment.

SCCP Components

The principal components of the SCCP VoIP architecture include the following:

- “SCCP Client” on page 251
- “CallManager” on page 251
- “Cluster” on page 251

SCCP Client

The SCCP client runs on an IP phone, also called an *End Station*, which uses SCCP for signaling and for making calls. For an SCCP client to make a call, it must first register with a Primary CallManager (and a secondary, if available). The connection between the client and the CallManager is over TCP on port 2000. This connection is then used to establish calls to or from the client. Transmission of media is over RTP, UDP, and IP.

CallManager

The CallManager implements SCCP call control server software and has overall control of all devices and communication in the SCCP VoIP network. Its functions include defining, monitoring and controlling SCCP groups, regions of numbers, and route plans; providing initialization, admission, and registration of devices on the network; providing a redundant database that contains addresses, phone numbers, and number formats; and initiating contact with called devices or their agents to establish logical sessions in which voice communication can flow.

Cluster

A Cluster is a collection of SCCP clients and a CallManager. The CallManager in the cluster knows about all SCCP clients in the cluster. There can be more than one CallManager for backup in a cluster. CallManager behavior varies in each of the following cluster scenarios:

- Intra-Cluster, in which the CallManager knows about each SCCP client, and the call is between SCCP clients of the same cluster.
- Inter-Cluster, in which the CallManager needs to communicate with another CallManager using H.323 for call setup.
- Inter-Cluster calls using the gatekeeper for admission control and address resolution.

  CallManager behavior also varies with calls between an SCCP client and a phone in a public switched telephone network (PSTN), and with calls between an SCCP client and a phone in another administrative domain that is using H.323.

SCCP Transactions

SCCP transactions are the processes that need to take place in order for an SCCP call to proceed. SCCP transactions include the following processes:

- "Client Initialization" on page 252
- "Client Registration" on page 252
- "Call Setup" on page 253
- "Media Setup" on page 253

Client Initialization

To initialize, the SCCP client needs to know the IP address of the CallManager, its own IP address, and other information about the IP gateway and DNS servers. Initialization takes place on the local LAN. The client sends a Dynamic Host Control Protocol (DHCP) request to get an IP address, the DNS server address, and the TFTP server name and address. The client needs the TFTP server name to download the configuration file called sep*macaddr*.cnf. If the TFTP name is not given, the client uses the default filename in the IP phone. The client then downloads the .cnf (xml) configuration file from TFTP server. CNF files contain the IP address or addresses of the primary and secondary Cisco CallManager. With this information, the client contacts the CallManager to register.

Client Registration

The SCCP client, after initialization, registers with the CallManager over a TCP connection on well-known default port 2000. The client registers by providing the CallManager with its IP address, the MAC address of the phone, and other information, such as protocol and version. The client cannot initiate or receive calls until it is registered. Keepalive messages keep this TCP connection open between the client and CallManager so that the client can initiate or receive calls at any time, provided that a policy on the device allows this.

Table 40 on page 252 lists SCCP messages and indicates messages that are of interest to the device.

Table 40: SCCP Registration Messages

RegisterMessage		b
IPortMessage		b
	RegisterAckMessage	b
	CapabilititsRequest	
CapabilitiesResMessage		
ButtonTemplateReqMessage		
	ButtonTemplateResMessage	
SoftKeyTemplateReqMessage		
	SoftKeyTemplateResMessage	
LineStatReqMessage		b
	LineStatMessage	b

Call Setup

IP phone-to-IP phone call setup using SCCP is always handled by the CallManager. Messages for call setup are sent to the CallManager, which returns messages appropriate to the status of the call. If call setup is successful, and a policy on the device allows the call, the CallManager sends the media setup messages to the client.

Media Setup

The CallManager sends the IP address and port number of the called party to the calling party. The CallManager also sends the media IP address and port number of the calling party to the called party. After media setup, media is transmitted directly between clients. When the call ends, the CallManager is informed and terminates the media streams. At no time during this process does the CallManager hand over call-setup function to the client. Media is streamed directly between clients through RTP/UDP/IP.

SCCP Control Messages and RTP Flow

Figure 32 on page 254 shows the SCCP control messages used to set up and tear down a simple call between *Phone1* and *Phone2*. Except for the OffHook message initiating the call from *Phone1* and the OnHook message signaling the end of the call, all aspects of the call are controlled by the CallManager.

Figure 32: Call Setup and Teardown

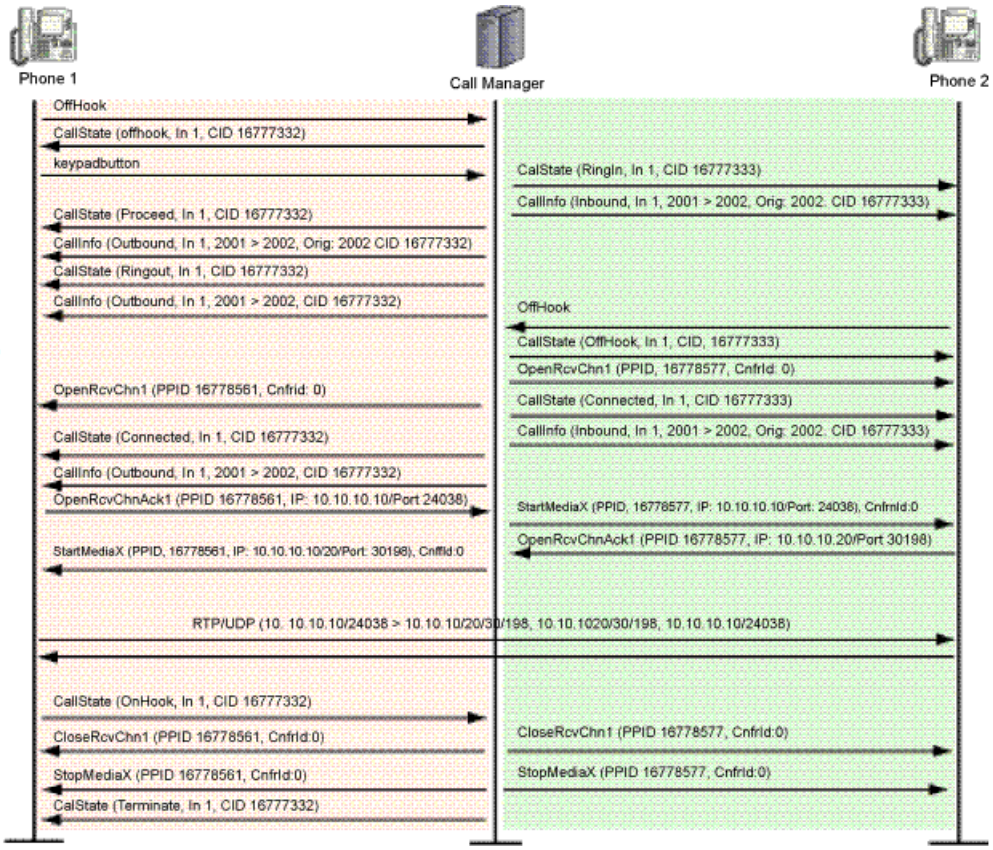

SCCP Messages

Table 41 on page 254, Table 42 on page 254, Table 43 on page 255, and Table 44 on page 255 list the SCCP call message IDs in the four intervals allowed by the J Series device.

Table 41: Station to CallManager Messages

#define STATION_REGISTER_MESSAGE	0x00000001
#define STATION_IP_PORT_MESSAGE	0x00000002
#define STATION_ALARM_MESSAGE	0x00000020
#define STATION_OPEN_RECEIVE_CHANNEL_ACK	0x00000022

Table 42: CallManager to Station Messages

#define STATION_START_MEDIA_TRANSMISSION	0x00000001
#define STATION_STOP_MEDIA_TRANSMISSION	0x00000002

Table 42: CallManager to Station Messages *(continued)*

#define STATION_CALL_INFO_MESSAGE	0x00000020
#define STATION_OPEN_RECEIVE_CHANNEL_ACK	0x00000022
#define STATION_CLOSE_RECEIVE_CHANNEL	0x00000106

Table 43: CallManager 4.0 Messages and Post Sccp 6.2

#define STATION_REGISTER_TOKEN_REQ_MESSAGE	0x00000029
#define STATION_MEDIA_TRANSMISSION_FAILURE	0x0000002A
#define STATION_OPEN_MULTIMEDIA_RECEIVE_CHANNEL_ACK	0x00000031

Table 44: CallManager to Station

#define STATION_OPEN_MULTIMEDIA_RECEIVE_CHANNEL	0x00000131
#define STATION_START_MULTIMEDIA_TRANSMISSION	0x00000132
#define STATION_STOP_MULTIMEDIA_TRANSMISSION	0x00000133
#define STATION_CLOSE_MULTIMEDIA_RECEIVE_CHANNEL	0x00000136

Related Topics

- Configuring the H.323 ALG—Quick Configuration on page 164
- Understanding the SIP ALG on page 187
- Understanding the MGCP ALG on page 266

Configuring the SCCP ALG—Quick Configuration

You can use J-Web Quick Configuration to quickly configure SCCP ALG Parameters

Before You Begin
For background information, read “Understanding the SCCP ALG” on page 249

Figure 33 on page 256 shows the SCCP ALG configuration page.

Figure 33: SCCP ALG Configuration

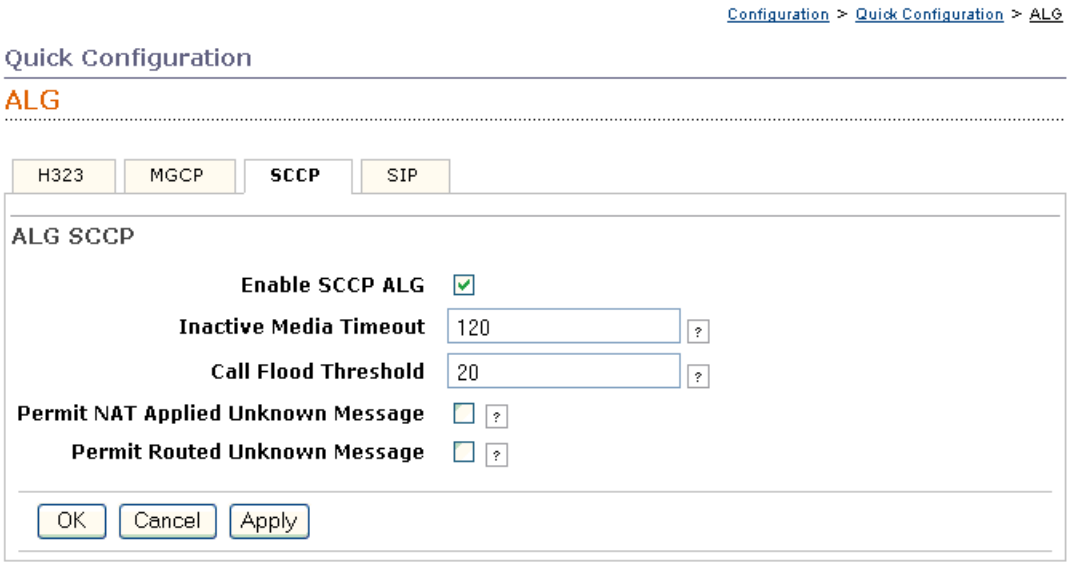

To configure the SCCP ALG with Quick Configuration:

1. In the J-Web user interface, select **Configure > Security > ALGs > VoIP**.
2. Select the **SCCP** tab if it is not selected.
3. Enter your parameter settings as described in Table 45 on page 256 and click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel.**

Table 45: SCCP ALG Configuration Options

Enable SCCP ALG	Enables or disables the SCCP ALG.	Click the check box.
Inactive Media Timeout	Indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the Skinny Client Control Protocol (SCCP) ALG the gates opened for media are closed.	Select a value from 10 to 600 seconds.
Call Flood Threshold	Protect Skinny Client Control Protocol (SCCP) ALG clients from flood attacks by limiting the number of calls they attempt to process	Select a value from 2 to 1,000.

Table 45: SCCP ALG Configuration Options *(continued)*

Permit NAT Applied Unknown Message	Specifies how unidentified SCCP messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown SCCP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped. This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.	Click the check box.
Permit Routed Unknown Message	Specifies that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)	Click the check box.

Setting SCCP Inactive Media Timeout

The inactive media timeout feature helps you to conserve network resources and maximize throughput.

Before You Begin
For background information, read "Understanding the SCCP ALG" on page 249.

This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the gates the Skinny Client Control Protocol (SCCP) opened for media are closed. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

NOTE: You must use the CLI to set SIP-signaling and media-inactivity timeouts.

To configure inactive media timeout, use either the J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 258
- CLI Configuration on page 258
- Related Topics on page 258

J-Web Configuration

1. Select Configure>Security>ALGs>VoIP.
2. Select the SCCP tab if it is not selected.
3. Enter a value in the Inactive media timeout field,
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In the following example, you set the media inactivity timeout to 90 seconds:

```
user@host# set security alg sccp inactive-media-timeout 90
```

Related Topics

- Understanding the SCCP ALG on page 249
- Allowing Unknown SCCP Message Types on page 258
- Configuring SCCP Denial of Service (DoS) Attack Protection on page 260
- Verifying the SCCP Configuration on page 263

Allowing Unknown SCCP Message Types

To accommodate on-going development of the Skinny Client Control Protocol (SCCP), you might want to allow traffic containing new SCCP message types. The unknown SCCP message type feature enables you to configure the device to accept SCCP traffic containing unknown message types in both NAT and route modes.

Before You Begin
For background information, read “Understanding the SCCP ALG” on page 249.

This feature enables you to specify how unidentified SCCP messages are handled by the device. The default is to drop unknown (unsupported) messages. We do not recommend permitting unknown messages because they can compromise security and is not recommended. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown SCCP messages can help you get your network operational so that you can later analyze your VoIP traffic to determine why some messages were being dropped.

Note that this command applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

- permit-nat-applied specifies that unknown messages be allowed to pass if the session is in NAT mode.
- permit-routed specifies that unknown messages be allowed to pass if the session is in route mode.

To allow unknown message types, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 259
- CLI Configuration on page 260
- Related Topics on page 260

J-Web Configuration

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, select Configure or Edit.
3. Next to Alg, select Configure or Edit.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.
5. Next to Application screen, click Configure.
6. Next to Unknown message, click Configure.
7. Click one of the following check boxes:
   - To allow unknown message types in NAT mode, click Permit nat applied Yes check box.
   - To allow unknown message types in route mode, click Permit routed Yes check box.
8. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In this example, you configure the device to allow unknown SCCP message types in both route and NAT modes.

```
user@host# set security alg sccp application-screen unknown-message
   permit-nat-applied permit-routed
```

Related Topics

- Understanding the SCCP ALG on page 249
- Setting SCCP Inactive Media Timeout on page 257
- Configuring SCCP Denial of Service (DoS) Attack Protection on page 260
- Verifying the SCCP Configuration on page 263

Configuring SCCP Denial of Service (DoS) Attack Protection

You can protect Skinny Client Control Protocol (SCCP) ALG clients from flood attacks by limiting the number of calls they attempt to process.

Before You Begin
For background information, read "Understanding the SCCP ALG" on page 249.

When you configure SCCP call flood protection, the SCCP ALG drops any calls exceeding the threshold you set. The range is 2 to 1,000 calls per second per client, the default is 20.

To configure DoS attack protection, use the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 260
- CLI Configuration on page 261
- Related Topics on page 261

J-Web Configuration

To configure the Services Router to drop any calls exceeding 500 per second per client:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, select Configure or Edit.
3. Next to Alg, select Configure or Edit.
4. Select the Sip check box if it is not selected and click **Configure** or **Edit**.

5. Next to Application screen, click Configure.
6. In the Threshold box, type **500** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

In this example, you configure the device to drop any calls exceeding 500 per second per client.

```
user@host# set security alg sccp application-screen call-flood threshold 500
```

Related Topics

- Understanding the SCCP ALG on page 249
- Setting SCCP Inactive Media Timeout on page 257
- Allowing Unknown SCCP Message Types on page 258
- Verifying the SCCP Configuration on page 263

Configuring Call Manager/TFTP Server in the Private Zone

When the same device serves as both the Call Manager and the TFTP server and are located in the private network, you might want to configure static NAT on the outgoing interface of the Juniper Networks device.

Before You Begin

For background information, read

- Understanding the SCCP ALG on page 249
- Configuring the SCCP ALG—Quick Configuration on page 255

In this example, phone1 and the Call Manager/TFTP Server are on the ge-0/0/0.0 interface in the private zone, and phone2 is on the ge-0/0/1.0 interface in the public zone. You configure static NAT for the Call Manager/TFTP Server on the ge-0/0/1.0 interface, so that when phone2 boots up it can contact the TFTP Server and obtain the IP address of the Call Manager. (We recommend that you change the IP address of the Call Manager in the TFTP Server config file (sep < mac_addr > .cnf) to the NAT IP address of the Call Manager.) You then create a policy allowing SCCP traffic from the public to the private zone and reference that NAT in the policy. You also create a policy from the Trust to the Untrust zone to allow phone1 to call out. See Figure 34 on page 262.

Figure 34: Call Manager/TFTP Server in the Private Zone

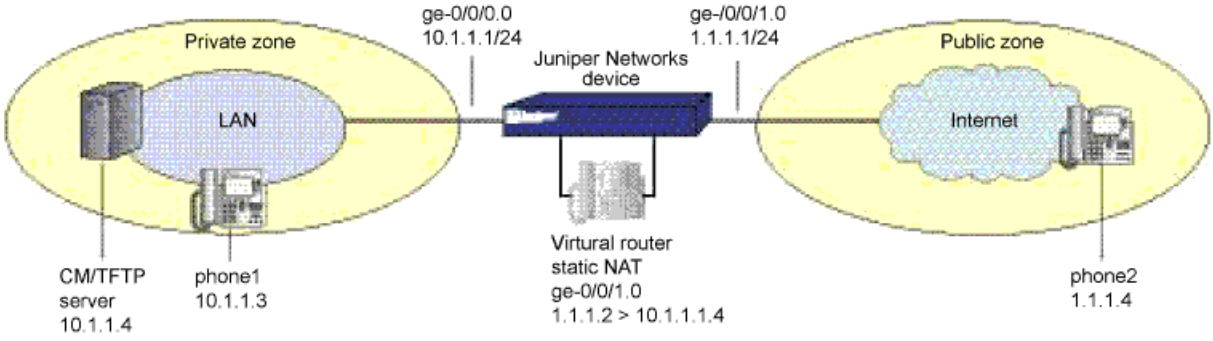

Use either the J-Web or CLI configuration editor.

This topic covers:

■ CLI Configuration on page 262

CLI Configuration

1. Configure interfaces.

   user@host# **set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24**
   user@host# **set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24**

2. Configure zone.

   user@host# **set security zones security-zone private interface ge-0/0/0.0**
   user@host# **set security zones security-zone private address-book address phone1 10.1.1.3/32**
   user@host# **set security zones security-zone private address-book address cm-tftp_server 10.1.1.4/32**
   user@host# **set security zones security-zone public interface ge-0/0/1.0**
   user@host# **set security zones security-zone public address-book address phone2 1.1.1.4/32**

3. Configure static NAT.

   user@host# **set security nat interface ge-0/0/1.0 static 1.1.1.2 host 10.1.1.4**

4. Configure policies.

   user@host# **set security policies from-zone private to-zone public policy out-pol match source-address any**
   user@host# **set security policies from-zone private to-zone public policy out-pol match destination-address phone2**
   user@host# **set security policies from-zone private to-zone public policy out-pol match application junos-sccp**
   user@host# **set security policies from-zone private to-zone public policy out-pol then permit source-nat interface**
   user@host# **set security policies from-zone public to-zone junos-global policy in-pol match source-address phone2**

```
user@host# set security policies from-zone public to-zone junos-global policy
  in-pol match destination-address static_nat_1.1.1.2_32
user@host# set security policies from-zone public to-zone junos-global policy
  in-pol match application junos-sccp
user@host# set security policies from-zone public to-zone junos-global policy
  in-pol then permit
```

Verifying the SCCP Configuration

To verify the SCCP configuration, perform these tasks:

- "Verifying the SCCP ALG" on page 263
- "Verifying SCCP Call Details" on page 264
- "Verifying SCCP Counters" on page 265

Verifying the SCCP ALG

Purpose Display SCCP verification options.

Action From the CLI, enter the show security alg sccp command.

```
user@host> show security alg sccp ?
Possible completions:
  calls                Show SCCP calls
  counters             Show SCCP counters
```

What it Means The output shows a list of all SCCP verification parameters. Verify the following information:

- All SCCP calls
- Counters for all SCCP calls

Related Topics

- Understanding the SCCP ALG on page 249
- Verifying SCCP Calls on page 263
- Verifying SCCP Call Details on page 264
- Verifying SCCP Counters on page 265

Verifying SCCP Calls

Purpose Display a list of all SCCP calls

Action From the CLI, enter the show security alg sccp calls command.

```
user@host> show security alg sccp calls
Possible completions:
  calls                Show SCCP calls
```

```
  counters              Show SCCP counters
  endpoints             Show SCCP endpoints
```

What it Means The output shows a list of all SCCP verification parameters. Verify the following information:

- All SCCP calls
- Counters for all SCCP c alls
- Information about all SCCP endpoints

Related Topics

- Understanding the SCCP ALG on page 249
- Verifying the SCCP ALG on page 263
- Verifying SCCP Call Details on page 264
- Verifying SCCP Counters on page 265

Verifying SCCP Call Details

Purpose Display details about all SCCP calls.

Action From the CLI, enter the show security alg sccp calls detail command.

```
user@host> show security alg sccp calls detail
Client IP address: 11.0.102.91
  Client zone: 7
  CallManager IP: 13.0.99.226
  Conference ID: 16789504
  Resource manager group: 2048
  SCCP channel information:
    Media transmit channel address (IP address/Port): 0.0.0.0:0
    Media transmit channel translated address (IP address/Port): 0.0.0.0:0
    Media transmit channel pass-through party ID (PPID): 0
    Media transmit channel resource ID: 0
    Media receive channel address (IP address/Port): 11.0.102.91:20060
    Media receive channel translated address (IP address/Port): 25.0.0.1:1032
    Media receive channel pass-through party ID (PPID): 16934451
    Media receive channel resource ID: 8185
    Multimedia transmit channel address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel pass-through party ID (PPID): 0
    Multimedia transmit channel resource ID: 0
    Multimedia receive channel address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel pass-through party ID (PPID): 0
    Multimedia receive channel resource ID: 0
Total number of calls = 1
```

What it Means The output shows a list of all SCCP verification parameters. Verify the following information:

- Client zone
- CallManager IP address: 13.0.99.226

- Conference ID
- Resource manager group
- SCCP channel information
- Total number of calls

Related Topics

- Understanding the SCCP ALG on page 249
- Verifying the SCCP ALG on page 263
- Verifying SCCP Calls on page 263
- Verifying SCCP Counters on page 265

Verifying SCCP Counters

Purpose Display a list of all SCCP counters

Action From the J-Web interface, select Monitor>ALGs>SCCP>Counters. Alternatively, from the CLI, enter the show security alg sccp counters command.

```
user@host> show security alg sccp ounters
SCCP call statistics:
  Active client sessions         : 0
  Active calls                   : 0
  Total calls                    : 0
  Packets received               : 0
  PDUs processed                 : 0
  Current call rate              : 0
Error counters:
  Packets dropped                     : 0
  Decode errors                       : 0
  Protocol errors                     : 0
  Address translation errors          : 0
  Policy lookup errors                : 0
  Unknown PDUs                        : 0
  Maximum calls exceeded              : 0
  Maximum call rate exceeded          : 0
  Initialization errors               : 0
  Internal errors                     : 0
  Nonspecific error                   : 0
  No active calls to delete           : 0
  No active client sessions to delete : 0
  Session cookie create errors        : 0
  Invalid NAT cookie detected         : 0
```

What it Means The output shows a list of all SCCP verification parameters. Verify the following information:

- SCCP call statistics
- Error counters

Related Topics

- Understanding the SCCP ALG on page 249
- Verifying the SCCP ALG on page 263
- Verifying SCCP Calls on page 263
- Verifying SCCP Call Details on page 264

Understanding the MGCP ALG

The Media Gateway Control Protocol (MGCP) is a text-based Application Layer protocol used for call setup and call control between the media gateway and the media gateway controller (MGC).

Before You Begin

For background information, read

- Application Layer Gateways (ALGs) on page 157
- Pinhole Creation on page 190

The protocol is based on a master/slave call control architecture: the media gateway controller (call agent) maintains call control intelligence, and media gateways carry out the instructions from the call agent. Both signaling packets and media packets are transmitted over UDP. JUNOS Software supports MGCP in route mode and Network Address Translation (NAT) mode.

The MGCP ALG performs the following procedures:

- Conducts VoIP signaling payload inspection. The payload of the incoming VoIP signaling packet is fully inspected based on related RFCs and proprietary standards. Any malformed packet attack is blocked by the ALG.
- Conducts MGCP signaling payload inspection. The payload of the incoming MGCP signaling packet is fully inspected in accordance with RFC 3435. Any malformed-packet attack is blocked by the ALG.
- Provides stateful processing. The corresponding VoIP-based state machines are invoked to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- Performs NAT. Any embedded IP address and port information in the payload is properly translated based on the existing routing information and network topology, and is then replaced with the translated IP address and port number, if necessary.
- Manages pinholes for VoIP traffic. To keep the VoIP network secure, the IP address and port information used for media or signaling is identified by the ALG, and any needed pinhole is dynamically created and closed during call setup.

This topic covers:

- MGCP Security on page 267
- Entities in MGCP on page 267
- Commands on page 268
- Response Codes on page 271
- Related Topics on page 272

MGCP Security

The MGCP ALG includes the following security features:

- Denial of Service (DoS) attack protection.—the ALG performs stateful inspection at the UDP packet level, the transaction level, and at the call level. MGCP packets matching the RFC3435 message format, transaction state, and call state, are processed. All other messages are dropped.
- Security policy enforcement between gateway and gateway controller (signaling policy).
- Security policy enforcement between gateways (media policy).
- Per-gateway MGCP message flooding control. Any malfunctioning or hacked gateway will not disrupt the whole VoIP network. Combined with per-gateway flooding control, damage is contained within the impacted gateway.
- Per-gateway MGCP connection flooding control.
- Seamless switchover/failover if calls, including calls in progress, are switched to the standby firewall in case of system failure.

Entities in MGCP

There are four basic entities in MGCP:

- “Endpoint” on page 267
- “Connection” on page 268
- “Call” on page 268
- “Call Agent” on page 268

Endpoint

A media gateway is a collection of endpoints. An endpoint can be an analog line, trunk, or any other access point. An endpoint contains the following elements:

local-endpoint-name@domain-name

The following examples are some valid endpoint IDs:

group1/Trk8@mynetwork.net
group2/Trk1/*@[192.168.10.8] (wild-carding)
$@voiptel.net (any endpoint within the media gateway)

*@voiptel.net (all endpoints within the media gateway)

Connection

Connections are created on each endpoint by an MG during call setup. A typical VoIP call involves two connections. A complex call, for example a three-party call or conference call, might require more connections. The media gateway controller (MGC) can instruct media gateways to create, modify, delete and audit a connection.

A connection is identified by its connection ID which is created by the MG when it is requested to create a connection. Connection ID is presented as a hexadecimal string, and its maximum length is 32 characters

Call

A call is identified by its call ID, which is created by the MGC when establishing a new call. Call ID is a hexadecimal string with a maximum length of 32 characters. Call ID is unique within the MGC. Two or more connections can have the same call ID if they belong to the same call.

Call Agent

One or more call agents (also called media gateway controllers) are supported in MGCP to enhance reliability in the VoIP network. The following two examples are of call agent names:

CallAgent@voipCA.mynetwork.com
voipCA.mynetwork.com

Several network addresses can be associated under one domain name in the Domain Name System (DNS). By keeping track of the time to live (TTL) of DNS query/response data and implementing retransmission using other alternative network addresses, switchover and failover is achieved in MGCP.

The concept of a *notified entity* is essential in MGCP. The notified entity for an endpoint is the call agent currently controlling that endpoint. An endpoint should send any MGCP command to its notified entity. However, different call agents might send MGCP commands to this endpoint.

The notified entity is set to a provisioned value upon startup, but can be changed by a call agent through the use of the NotifiedEntity parameter contained in an MGCP message. If the notified entity for an endpoint is empty or has not been set explicitly, its value defaults to the source address of the last successful non-audit MGCP command received for that endpoint.

Commands

The MGCP protocol defines nine commands for controlling endpoints and connections. All commands are composed of a command header, optionally followed by Session Description Protocol (SDP) information. A command header has the following elements:

- A command line: command verb + transaction ID + endpointId + MGCP version.
- Zero or more parameter lines, composed of a parameter name followed by a parameter value.

Table 46 on page 269 lists supported MGCP commands and includes a description of each, the command syntax, and examples. Refer to RFC 2234 for a complete explanation of command syntax.

Table 46: MGCP Commands

EPCF	EndpointConfiguration—used by a call agent to inform a gateway of coding characteristics (a-law or mu-law) expected by the line side of the endpoint.	ReturnCode[PackageList] EndpointConfiguration (EndpointId,[BearerInformation])	EPCF 2012 wxx/T2@mynet.com MGCP 1.0B: e:mu
CRCX	CreateConnection—used by a call agent to instruct the gateway to create a connection with, and endpoint inside, the gateway.	ReturnCode, [ConnectionId,] [SpecificEndPointId,] [LocalConnectionDescriptor,] [SecondEndPointId,] [SecondConnectionId,] [PackageList] CreateConnection (CallId, EndpointId, [NotifiedEntity,] [LocalConnectionOption,] Mode, [{RemoteConnectionDescriptor | SecondEndpoindId},] [encapsulated RQNT,] [encapsulated EPCF])	CRCX 1205 aaln/1@gw-25.att.net MGCP 1.0C: A3C47F21456789F0L: p:10, a:PCMUM: sendrecvX: 0123456789ADR: L/hdS: L/rgv = 0o = - 25678 753849 IN IP4 128.96.41.1s = -c = IN IP4 128.96.41.1t = 0 0m = audio 3456 RTP/AVP 0
MDCX	ModifyConnection—used by a call agent to instruct a gateway to change the parameters for an existing connection.	ReturnCode[LocalConnectionDescriptor][PackageList] ModifyConnection (CallId, EndpointId, ConnectionId, [NotifiedEntity,] [LocalConnectionOption,] [Mode,] [RemoteConnectionDescriptor,] [encapsulated RQNT,] [encapsulated EPCF])	MDCX 1210 aaln/1@rgw-25.att.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8M: recvonlyX: 0123456789AER: L/huS: G/rtv = 0o = - 4723891 7428910 IN IP4 128.96.63.25s = -c = IN IP4 128.96.63.25t = 0 0m = audio 3456 RTP/AVP 0

Table 46: MGCP Commands *(continued)*

DLCX	DeleteConnection—used by a call agent to instruct a gateway to delete an existing connection. DeleteConnection can also be used by a gateway to release a connection that can no longer be sustained.	ReturnCode,ConnectionParameters,[PackageList] DeleteConnection (CallId, EndpointId, ConnectionId, [NotifiedEntity,] [encapsulated RQNT,] [encapsulated EPCF])	Example 1: MGC -> MG DLCX 9210 aaln/1@rgw-25.att.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8 Example 2: MG -> MGC DLCX 9310 aaln/1@rgw-25.att.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8E: 900 - Hardware errorP: PS = 1245, OS = 62345, PR = 780, OR = 45123, PL = 10, JI = 27, LA = 48
RQNT	NotificationRequest command—used by a call agent to instruct an MG to monitor for certain event(s) or signal(s) for a specific endpoint.	ReturnCode,[PackageList] NotificationRequest[(EndpointId, [NotifiedEntity,] [RequestedEvents,] RequestIdentifier, [DigitMap,] [SignalRequests,] [QuarantineHandling,] [DetectEvents,] [encapsulated EPCF])	RQNT 1205 aaln/1@rgw-25.att.net MGCP 1.0N: ca-new@callagent-ca.att.netX: 0123456789AAR: L/hd(A, E(S(L/dl),R(L/oc,L/hu,D/[0-9#*T](D)))D: (0T|00T|xx|91xxxxxxxxxx|9011x.T)S:T: G/ft
NTFY	Notify—used by a gateway to inform the call agent when requested event(s) or signal(s) occur.	ReturnCode,[PackageList] Notify (EndpointID, [NotifiedEntity,] RequestIdentifier, ObservedEvents)	NTFY 2002 aaln/1@rgw-25.att.net MGCP 1.0N: ca@ca1.att.net:5678X: 0123456789ACO: L/hd,D/9,D/1,D/2,D/0,D/1,D/8,D/2,D/9,D/4, D/2,D/6,D/6
AUEP	AuditEndpoint—used by a call agent to audit the status of the endpoint.	ReturnCode, EndPointIdList, | { [RequestedEvents,] [QuarantineHandling,] [DigitMap,] [SignalRequests,] [RequestedIdentifier,] [NotifiedEntity,] [ConnectionIdentifier,] [DetectEvents,] [ObservedEvents,] [EventStats,] [BearerInformation,] [BearerMethod,] [RestartDelay,] [ReasonCode,] [MaxMGCPDatagram,] [Capabilities]} [PackageList] AuditEndpoint (EndpointId, [RequestedInfo])	Example 1: AUEP 1201 aaln/1@rgw-25.att.net MGCP 1.0F: A, R,D,S,X,N,I,T,OExample 2:AUEP 1200 *@rgw-25.att.net MGCP 1.0

Table 46: MGCP Commands *(continued)*

AUCX	AuditConnection—used by a call agent to collect the parameters applied to a connection.	ReturnCode, [CallId,] [NotifiedEntity,] [LocalConnectionOptions,] [Mode,] [RemoteConnectionDescriptor,] [LocalConnectionDescriptor,] [ConnectionParameters,] [PackageList] AuditConnection (EndpointId, ConnectionId, RequestedInfo)	AUCX 3003 aaln/1@rgw-25.att.net MGCP 1.0I: 32F345E2F: C,N,L,M,LC,P
RSIP	RestareInProgress—used by a gateway to notify a call agent that one or more endpoints are being taken out of service or placed back in service.	ReturnCode,[NotifiedEntity,][PackageList] RestartInProgress (EndpointId, RestartMethod, [RestartDelay,] [ReasonCode])	RSIP 5200 aaln/1@rg2-25.att.net MGCP 1.0RM: gracefulRD: 300

Response Codes

Every command sent by the calling agent or gateway, whether successful or not, requires a response code. The response code is in the header of the response message, and optionally is followed by session description information.

The response header is composed of a response line, followed by zero or more parameter lines, each containing a parameter name letter followed by its value. The response header is composed of a three-digit response code, transaction ID, and optionally followed by commentary. The response header in the following response message shows the response code 200 (successful completion), followed by ID 1204 and the comment:OK.

```
200 1204 OK
I: FDE234C8
v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=audio 3456 RTP/AVP 96
a=rtpmap:96 G726-32/8000
```

The ranges of response codes are defined as follows:

- 000 — 099 indicate a response acknowledgement.
- 100 — 199—indicate a provisional response.
- 200 — 299 indicate a successful completion (final response).
- 400 — 499 indicate a transient error (final response).
- 500 — 599 indicate a permanent error (final response).

Refer to RFC 3661 for detailed information about response codes.

A response to a command is sent to the source address of the command, not to the current notified entity. A media gateway can receive MGCP commands from various network addresses simultaneously, and send back responses to corresponding network addresses. However, it sends all MGCP commands to its current notified entity.

Related Topics

- Understanding the H.323 ALG on page 162
- Understanding the SIP ALG on page 187
- Configuring a Media Gateway in Subscribers' Homes on page 282

Configuring the MGCP ALG—Quick Configuration

You can use J-Web Quick Configuration to quickly configure MGCP ALG Parameters

Before You Begin
For background information, read “Understanding the MGCP ALG” on page 266

Figure 35 on page 272 shows the MGCP ALG configuration page.

Figure 35: MGCP ALG Configuration

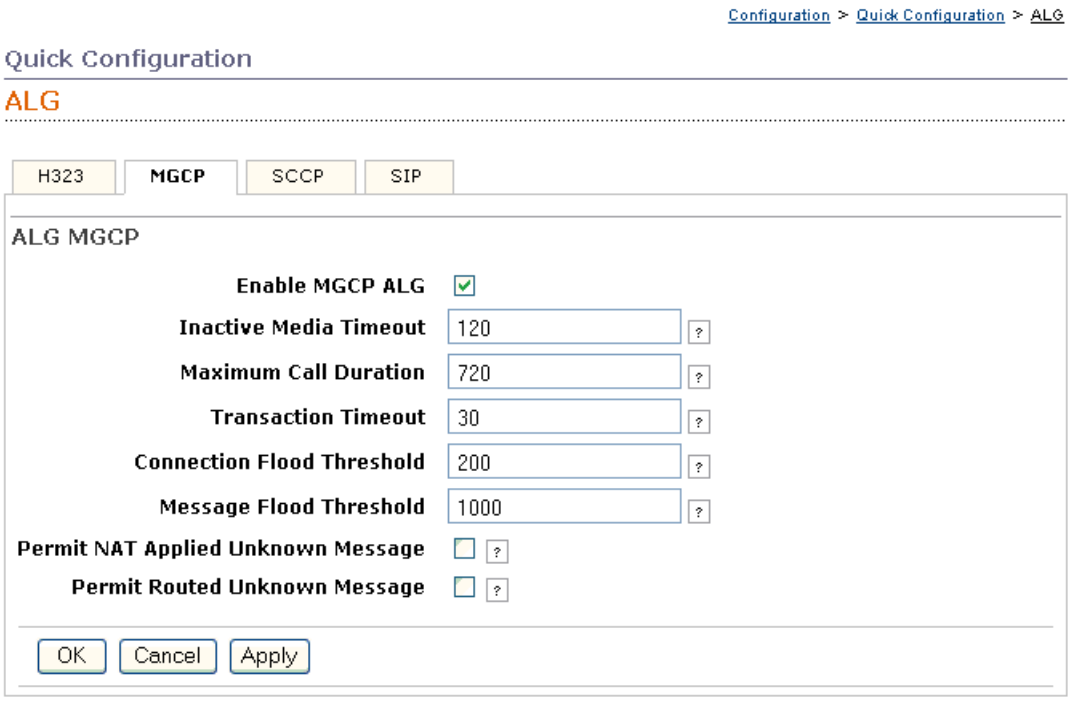

To configure the MGCP ALG with Quick Configuration:

1. Select Configure>Security>ALGs>VoIP ALG.
2. Select the MGCP tab if it is not selected.
3. Fill in the options as described in Table 47 on page 273 and click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click Apply.
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

Table 47: MGCP Configuration Options

Enable MGCP ALG	Enables or disables the MGCP ALG.	Click the check box.
Inactive Media Timeout	Specifies the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the temporary openings (pinholes) in the firewall MGCP ALG opened for media are closed. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.	Select a value between 10 and 2,550 seconds.
Maximum Call Duration	Sets the absolute maximum length of a call. When a call exceeds this parameter setting, the MGCP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, the range is from 3 to 7200 minutes.	Select a value between 3 and 7,200 minutes.
Transaction Timeout	Specifies a timeout value for MGCP transactions. A transaction is a signalling message, for example, a NTFY from the gateway to the call agent or a 200 OK from the call agent to the gateway. The Juniper Networks device tracks these transactions, and clears them when they time out.	Enter a value from 3 to 50 seconds.
Connection Flood Threshold	Limits the number of new connection requests allowed per Media Gateway (MG) per second. Messages exceeding the ALG.	Enter a value from 2 to 10,000.
Message Flood Threshold	Limits the rate per second at which message requests to the Media Gateway are processed. Messages exceeding the threshold are dropped by the Media Gateway Control Protocol (MGCP) Application Layer Gateway (ALG). This feature is disabled by default.	Enter a value from 2 to 50,000 seconds per media gateway.

Table 47: MGCP Configuration Options *(continued)*

Permit NAT Applied Unknown Message	Specifies how unidentified MGCP messages are handled by the Juniper Networks device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown MGCP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped. This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.	Click the check box.
Permit Routed Unknown Message	Specifies that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)	Click the check box.

Understanding MGCP ALG Call Duration and Timeouts

The call duration feature gives you control over MGCP call activity and helps you to manage network resources.

Before You Begin

For background information, read "Understanding the MGCP ALG" on page 266.

Typically a DLCX message will be sent out to delete a connection. The MCGP ALG intercepts it and removes all media sessions for that connection.

A call can have one or more voice channels. Each voice channel has two sessions (or two media streams), one for Real-Time Transport Protocol (RTP) traffic and one for Real-Time Control Protocol (RTCP) signaling. When managing the sessions, the device considers the sessions in each voice channel as one group. Timeouts and call duration settings apply to a group as opposed to each session.

The following parameters govern MGCP call activity:

- inactive-media-timeout—This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the MGCP ALG gates opened for media are closed. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

- transaction-timeout—The device tracks transactions between the gateway and the call agent and clears transactions when they time out. Timeout range for MGCP transactions is from 3 to 50 seconds, the default is 30 seconds.
- maximum-call-duration—This parameter sets the absolute maximum length of a call. When a call exceeds this parameter setting, the MGCP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, the range is from 3 to 7200 minutes.

Related Topics

- Understanding the MGCP ALG on page 266
- Setting MGCP Call Duration on page 275
- Setting MGCP Inactive Media Timeout on page 276
- Setting the MGCP Transaction Timeout on page 278
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 279
- Allowing Unknown MGCP Message Types on page 280

Setting MGCP Call Duration

The call duration feature helps you to conserve network resources and maximize throughput.

Before You Begin

For background information, read

- Understanding the MGCP ALG on page 266
- Understanding MGCP ALG Call Duration and Timeouts on page 274

You use the call duration parameter to set the maximum allowable length of time a call can be active. When the duration is exceeded, the MGCP ALG tears down the call and releases the media sessions. This setting also frees up bandwidth in cases where calls fail to properly terminate.

Use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 276
- CLI Configuration on page 276
- Related Topics on page 276

J-Web Configuration

1. Select Configure>Security>ALGs>VoIP ALG.
2. Select the MGCP tab if it is not selected.
3. Enter a value in the Maximum call duration box.
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In the following example, you set call duration to 3000 minutes.

```
user@host# set security alg mgcp maximum-call-duration 3000
```

Related Topics

- Understanding the MGCP ALG on page 266
- Understanding MGCP ALG Call Duration and Timeouts on page 274
- Setting MGCP Inactive Media Timeout on page 276
- Setting the MGCP Transaction Timeout on page 278
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 279
- Allowing Unknown MGCP Message Types on page 280

Setting MGCP Inactive Media Timeout

The inactive media timeout feature helps you to conserve network resources and maximize throughput.

Before You Begin
For background information, read ■ Understanding the MGCP ALG on page 266 ■ Understanding MGCP ALG Call Duration and Timeouts on page 274 ■ Setting the MGCP Transaction Timeout on page 278

You use the inactive media parameter to set the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group.

Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the MGCP ALG closes the gates it opened for media. The default setting is 120 seconds, the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

To set an inactive media timeout value, use either the J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 277
- CLI Configuration on page 277
- Related Topics on page 277

J-Web Configuration

1. Select Configure>Security>ALGs>VoIP ALG.
2. Select the MGCP tab if it is not selected.
3. Enter a value in the Inactive Media Timeout box.
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In the following example, you set the inactive media timeout to 90 seconds:

```
user@host# set security alg mgcp inactive-media-timeout 90
```

Related Topics

- Understanding the MGCP ALG on page 266
- Understanding MGCP ALG Call Duration and Timeouts on page 274
- Setting MGCP Call Duration on page 275
- Setting the MGCP Transaction Timeout on page 278
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 279
- Allowing Unknown MGCP Message Types on page 280

Setting the MGCP Transaction Timeout

The transaction timeout feature helps you to conserve network resources and maximize throughput.

Before You Begin

For background information, read

- "Understanding the MGCP ALG" on page 266.
- "Understanding MGCP ALG Call Duration and Timeouts" on page 274.
- "Setting MGCP Inactive Media Timeout" on page 276.

A transaction is a signaling message, for example, a NTFY from the gateway to the call agent or a 200 OK from the call agent to the gateway. The Juniper Networks device tracks these transactions, and clears them when they time out. The timeout range for MGCP transactions is from 3 to 50 seconds, the default is 30 seconds.

Use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 278
- CLI Configuration on page 278
- Related Topics on page 279

J-Web Configuration

1. Select Configure>Security>ALGs>VoIP ALG.
2. Select the MGCP tab if it is not selected.
3. Enter a value in the Transaction Timeout box.
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In the following example, you set the transaction timeout to 20 seconds:

```
user@host# set security alg mgcp transaction-timeout 20
```

Related Topics

- Understanding the MGCP ALG on page 266
- Understanding MGCP ALG Call Duration and Timeouts on page 274
- Setting MGCP Call Duration on page 275
- Setting MGCP Inactive Media Timeout on page 276
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 279
- Allowing Unknown MGCP Message Types on page 280

Configuring MGCP Denial of Service (DoS) Attack Protection

You can protect the MGCP media gateway from flood attacks by limiting the number of remote access service (RAS) messages and connections per second it will attempt to process.

Before You Begin

For background information, read

- "Understanding the MGCP ALG" on page 266.
- "Understanding MGCP ALG Call Duration and Timeouts" on page 274.

When you configure MGCP message flood protection, the MGCP ALG drops any messages exceeding the threshold you set. The range is 2 to 50,000 messages per second per media gateway, the default is 1,000 messages per second per media gateway.

When you configure MGCP connection flood protection, the MGCP ALG drops any connection request exceeding the threshold you set. This limits the rate of processing of CreatConnection (CRCX) commands, thereby indirectly limiting pinhole creation. The range is 2 to 10,000 connection requests per second per media gateway, the default is 200.

Use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 279
- CLI Configuration on page 280
- Related Topics on page 280

J-Web Configuration

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, select Configure or Edit.

3. Next to Alg, select Configure or Edit.
4. Select the Mgcp check box if it is not selected and click Configure or Edit.
5. Click Application Screen.
6. Under Connection flood, enter a value in the Threshold box.
7. Under Message flood, enter a value in the Threshold box.
8. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In this example, you configure the device to drop any message requests exceeding 10,000 requests per second and to drop any connection requests exceeding 4,000 per second.

```
user@host# set security alg mgcp application-screen message-flood threshold 10000
user@host# set security alg mgcp application-screen connection-flood threshold 4000
```

Related Topics

- Understanding the MGCP ALG on page 266
- Understanding MGCP ALG Call Duration and Timeouts on page 274
- Setting MGCP Call Duration on page 275
- Setting MGCP Inactive Media Timeout on page 276
- Setting the MGCP Transaction Timeout on page 278
- Allowing Unknown MGCP Message Types on page 280

Allowing Unknown MGCP Message Types

To accommodate on-going development of the Media Gateway Control Protocol (MGCP), you might want to allow traffic containing new MGCP message types. The unknown MGCP message type feature enables you to configure the Juniper Networks device to accept MGCP traffic containing unknown message types in both NAT and route modes.

Before You Begin
For background information, read “Understanding the MGCP ALG” on page 266.

This feature enables you to specify how unidentified MGCP messages are handled by the Juniper Networks device. The default is to drop unknown (unsupported) messages. Unknown messages can compromise security. However, in a secure test or production environmnet, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown MGCP messages can help you get your network operational so that you can later analyze your VoIP traffic to determine why some messages were being dropped.

Note that this command applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

- permit-nat-applied—Specifies that unknown messages be allowed to pass if the session is in NAT mode.
- permit-routed—Specifies that unknown messages be allowed to pass if the session is in Route mode.

To allow unknown message types, use either the J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 281
- CLI Configuration on page 282
- Related Topics on page 282

J-Web Configuration

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to ALG, click Configure or Edit.
4. Select the Mgcp check box if it is not selected and click Configure.
5. Click one of the following check boxes:
   - To allow unknown message types in NAT mode, click the Permit nat applied Yes check box.
   - To allow unknown message types in route mode, click the Permit routed Yes check box.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To apply the configuration, click Commit.
   - To cancel your entries and return to the main page, click Cancel.

CLI Configuration

In this example, you configure the device to allow unknown MGCP message types in both route and NAT modes.

```
user@host# set security alg mgcp application-screen unknown-message
   permit-nat-applied permit-routed
```

Related Topics

- Understanding the MGCP ALG on page 266
- Understanding SIP ALG Call Duration and Timeouts on page 196
- Setting MGCP Call Duration on page 275
- Setting MGCP Inactive Media Timeout on page 276
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 279

Configuring a Media Gateway in Subscribers' Homes

When a cable service provider offers MGCP services to residential subscribers, typically they locate the Juniper Networks device and the call agent on their premises and intstall an Integrated Access Device (IAD), or set-top box, in each subscriber's home. The IADs act as a gateways for the residences

Before You Begin
For background information, read “Understanding the MGCP ALG” on page 266.

After creating zones—external_subscriber for the customer and internal_ca for the service provider—you configure addresses, then interfaces, and finally policies to allow signaling between endpoints. Note that although gateways frequently reside in different zones, requiring policies for media traffic, in this example both gateways are in the same subnet. Note also that because Real-time Transport Protocol (RTP) traffic between the gateways never passes through the device, no policy is needed for the media. See Figure 36 on page 283.

Use either the J-Web or CLI configuration editor.

Figure 36: Media Gateway in Subscribers' Home

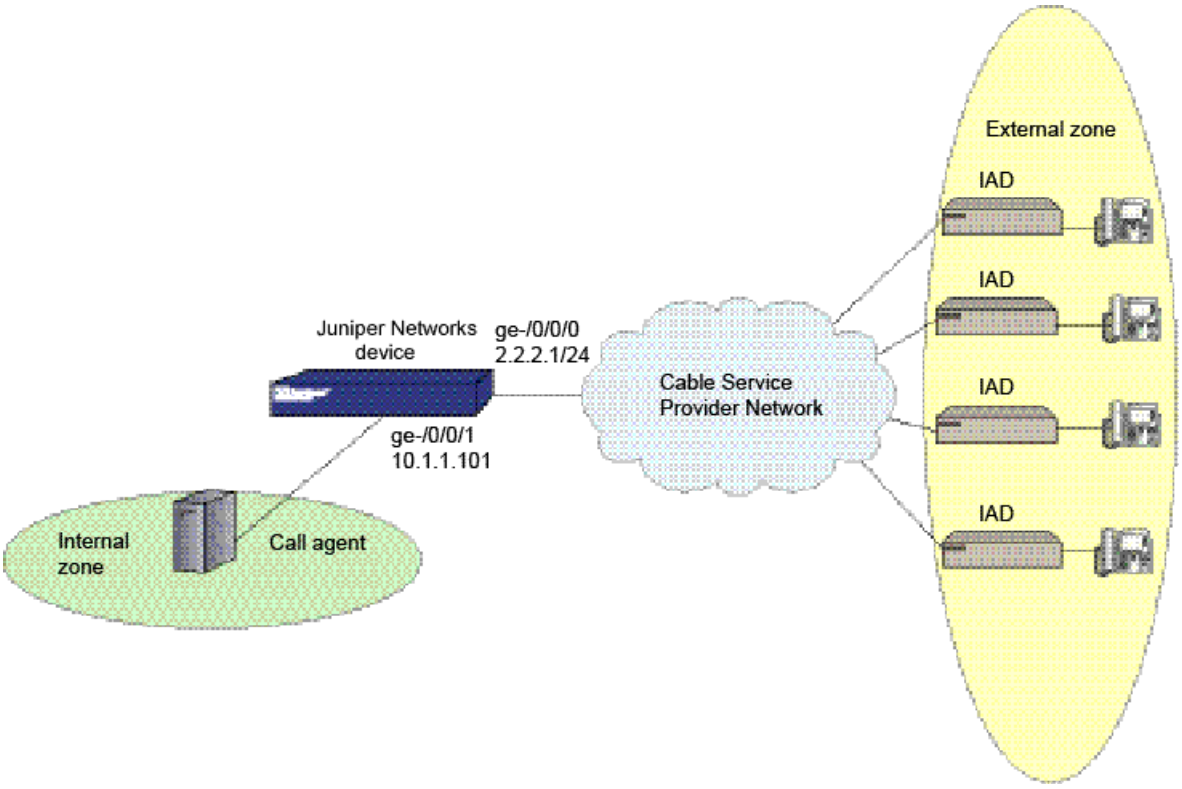

This topic covers:

- J-Web Configuration on page 283
- CLI Configuration on page 288
- Related Topics on page 289

J-Web Configuration

To configure zones:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zones, click Add new entry.
5. In the Name box, type **external_subscriber**.
6. Next to Host inbound traffic, click Configure.
7. Next to System services, click Add new entry.
8. Next to Service name, select all and click OK.
9. Next to Protocols, click Add new entry.
10. Next to Protocol name box, select all and click OK.

11. Next to Security zones, click Add new entry.
12. In the Name box, type **internal_ca**.
13. Next to Host inbound traffic, click Configure.
14. Next to System services, click Add new entry.
15. Next to Service name, select all and click OK.
16. Next to Protocols, click Add new entry.
17. Next to Protocol name box, select all and click OK.
18. To save and commit the configuration, click Commit.

To configure addresses:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **internal_ca**.
6. Next to Address book, click Configure or Edit.
7. Next to Address, click Add new entry.
8. In the Address name box, type **ca_agent1 10.1.1.101/32** and click OK.
9. To configure another security zone external_subscriber, repeat Step 2 through Step 9 and click OK.
10. Next to Address book, click Configure.
11. Next to Address, click Add new entry.
12. In the Address name box, type **SubscriberSubNet 2.2.2.1/24** and click OK.
13. To save and commit the configuration, click Commit.

To configure interfaces:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Zones, click Configure.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **internal_ca**
6. Next to Interfaces, click Add new entry.
7. In the Interface unit box, type **ge-0/0/01** and click OK.
8. Next to Interfaces, click Configure or Edit.
9. Next to Interface, click Add new entry.
10. In the Interface name box, type **ge-0/0/1**.

11. Next to unit, click Add new entry.
12. In the Interface unit number box, type **0**.
13. Under Family, next to Inet, select the check box and click Configure.
14. Next to Address, click Add new entry.
15. Next to Source, type **2.2.2.1/24** and click OK.
16. To configure another security zone, external _subscriber, and interface, ge-0/0/0, Step 1 through Step 7 and click OK.
17. To configure another interface, ge-0/0/0, repeat Step h to Step l.
18. Under Family, next to Inet, select the check box and click OK.
19. To save and commit the configuration, click Commit.

To configure internal-to-external zone policies:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **internal_ca**.
6. In the To zone name box, type **external_subscriber**.
7. Next to Policy, click Add new entry.
8. To specify the Policy name, next to Policy name box, type **Pol-CA-To-Subscribers**.
9. Select the Match check box.
10. Select the Then check box.
11. Click Configure next to Match check box.
12. From the Source address choice list, select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, select Enter Specific Value.
15. In the Address box, type **ca-agent1** and click OK.
16. From the Destination address choice list, select Destination address.
17. Next to Destination address, click Add new entry.
18. Next to Value keyword list, select Enter Specific Value.
19. In the Address box, type **SubscriberSubNet** and click OK.
20. From the Application choice list, select Application.
21. Next to Application, click Add new entry.
22. Next to Value keyword box, type junos-mgcp and click OK.
23. Next to Then, click Configure.

24. Next to Action, select permit and click OK.
25. To save and commit the configuration, click Commit.

To configure from zone, external_subscriber, and to zone, internal_ca:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **external_subscriber**.
6. In the To zone name box, type **internal_ca**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **Pol-Subscribers-To-CA**.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match check box, click Configure.
12. From the Source address choice list, select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, select Enter Specific Value.
15. In the Address name box, type **SubscriberSubnet** and click OK.
16. From the Destination address choice list, select Destination address.
17. Next to Destination address, click Add new entry.
18. Next to Value keyword list, select Enter Specific Value.
19. In the Address name box, type **call_agent1** and click OK.
20. From the Application choice list, select Application.
21. Next to Application, click Add new entry.
22. Next to Value keyword box, type **junos-mgcp** and click OK.
23. Next to Then, click Configure.
24. Next to Action, select permit and click OK.
25. To save and commit the configuration, click Commit.

To configure from zone and to zone as internal_ca:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.

5. In the From zone name box, type **internal_ca**.
6. In the To zone name box, type **internal_ca**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **Pol-Intra-CA**.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match check box, click Configure.
12. From the Source address choice list, select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, select any and click OK.
15. From the Destination address choice list, select Destination address.
16. Next to Destination address, click Add new entry.
17. Next to Value keyword list, select any and click OK.
18. From the Application choice list, select Application.
19. Next to Application, click Add new entry.
20. Next to Value keyword box, select any and click OK.
21. Next to Then, click Configure.
22. Next to Action, select permit.
23. To save and commit the configuration, click Commit.

To configure from zone and to zone as external_subscriber:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click Configure or Edit.
3. Next to Policies, select the check box and click Configure or Edit.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **external_subscriber**.
6. In the To zone name box, type **external_subscriber**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **Pol-Intra-subscriber**.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match check box, click Configure.
12. From the Source address choice list, select Source address.
13. Next to Source address, click Add new entry.
14. From the Value keyword list, select any and click OK.

15. From the Destination address choice list, select Destination address.
16. Next to Destination address, click Add new entry.
17. Next to Value keyword list, select any and click OK.
18. From the Application choice list, select Application.
19. Next to Application, click Add new entry.
20. Next to Value keyword box, select any and click OK.
21. Next to Then, click Configure.
22. Next to Action, select permit.
23. To save and commit the configuration, click Commit.

CLI Configuration

1. Configure zones.

```
user@host# set security zones security-zone external_subscriber
user@host# set security zones security-zone external_subscriber
  host-inbound-traffic system-services all
user@host# set security zones security-zone external_subscriber
  host-inbound-traffic protocols all
user@host# set security zones security-zone internal_ca
user@host# set security zones security-zone internal_ca host-inbound-traffic
  system-services all
user@host# set security zones security-zone internal_ca host-inbound-traffic
  protocols all
```

2. Configure addresses.

```
user@host# set security zones security-zone internal_ca address-book address
  ca_agent1 10.1.1.101/32
user@host# set security zones security-zone external_subscriber address-book
  address SubscriberSubNet 2.2.2.1/24
```

3. Configure interfaces.

```
user@host# set security zones security-zone internal_ca interfaces ge-0/0/1
user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.1/24
user@host# set security zones security-zone external_subscriber interfaces
  ge-0/0/0
user@host# set interfaces ge-0/0/0 unit 0 family inet
```

4. Configure internal-to-external zone policies.

```
user@host# set security policies from-zone internal_ca to-zone external_subscriber
  policy Pol-CA-To-Subscribers match source-address ca_agent1
user@host# set security policies from-zone internal_ca to-zone external_subscriber
  policy Pol-CA-To-Subscribers match destination-address SubscriberSubNet
user@host# set security policies from-zone internal_ca to-zone external_subscriber
  policy Pol-CA-To-Subscribers match application junos-mgcp
user@host# set security policies from-zone internal_ca to-zone external_subscriber
  policy Pol-CA-To-Subscribers then permit
```

```
user@host# set security policies from-zone external_subscriber to-zone internal_ca
    policy Pol-Subscribers-To-CA match source-address SubscriberSubNet
user@host# set security policies from-zone external_subscriber to-zone internal_ca
    policy Pol-Subscribers-To-CA match destination-address call_agent1
user@host# set security policies from-zone external_subscriber to-zone internal_ca
    policy Pol-Subscribers-To-CA match application junos-mgcp
user@host# set security policies from-zone external_subscriber to-zone internal_ca
    policy Pol-Subscribers-To-CA then permit
user@host# set security policies from-zone internal_ca to-zone internal_ca policy
    Pol-Intra-CA match source-address any
user@host# set security policies from-zone internal_ca to-zone internal_ca policy
    Pol-Intra-CA match destination-address any
user@host# set security policies from-zone internal_ca to-zone internal_ca policy
    Pol-Intra-CA match application any
user@host# set security policies from-zone internal_ca to-zone internal_ca policy
    Pol-Intra-CA then permit
user@host# set security policies from-zone external_subscriber to-zone
    external_subscriber policy Pol-Intra-subscriber match source-address any
user@host# set security policies from-zone external_subscriber to-zone
    external_subscriber policy Pol-Intra-subscriber match destination-address any
user@host# set security policies from-zone external_subscriber to-zone
    external_subscriber policy Pol-Intra-subscriber match application any
user@host# set security policies from-zone external_subscriber to-zone
    external_subscriber policy Pol-Intra-subscriber then permit
```

Related Topics

- Understanding the MGCP ALG on page 266
- Configuring Three-Zone ISP-Hosted Service Using Source and Static NAT on page 289

Configuring Three-Zone ISP-Hosted Service Using Source and Static NAT

When an Internet service provider (ISP) in one geographical location provides service to two networks in different geographical locations, a three-zone configuration might be necessary.

Before You Begin
For background information, read “Understanding the MGCP ALG” on page 266.

In this example, (see Figure 37 on page 290) an ISP located on the American west coast provides MGCP service to customers in separate networks in Asia and San Francisco. Asia customers are in the asia_3 zone and supported by the asia_gw gateway; San Francisco customers are in the sf_2 zone and supported by the sf_gw gateway; and the west_ca,call agent is in the DMZ. The gateways and the call agent are listed in Table 48 on page 290, showing the corrsponding IP address, interface, and zone.

Table 48: Three-Zone ISP-Host Service

sf_gw	2.2.2.201	ge-0/0/0	sf_2
asia-gw	3.3.3.110	ge-0/0/1	asia_3
west_ca	10.1.1.101	ge-0/0/2	DMZ

After creating zones and setting addresses for the gateways and the call agent, you associate the zones and addresses to interfaces, and then configure NAT (ge-0/0/1.0) and policies.

Figure 37: Three-Zone ISP-Hosted Service

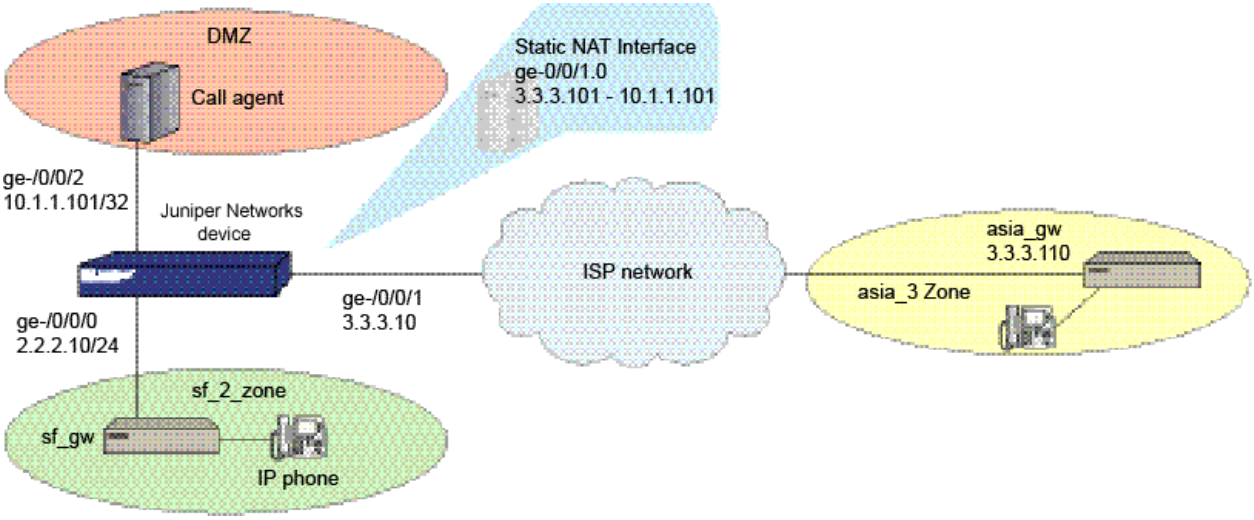

To configure a three-zone ISP-hosted service using source and static NAT, use either the J-Web or CLI configuration editor.

This topic covers:

- CLI Configuration on page 290
- Related Topics on page 292

CLI Configuration

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 2.2.2.10/24
user@host# set interfaces ge-0/0/1 unit 0 family inet address 3.3.3.10/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 10.1.1.2/24
```

2. Configure addresses.

```
user@host# set security zones security-zone sf_2 address-book address sf_gw
    2.2.2.201/32
```

```
user@host# set security zones security-zone asia_3 address-book address asia_gw
  3.3.3.110/32
user@host# set security zones security-zone dmz address-book address west_ca
  10.1.1.101/32
```

3. Associate the zones and addresses to interfaces.

```
user@host# set security zones security-zone sf_2 interfaces ge-0/0/0
user@host# set security zones security-zone asia_3 interfaces ge-0/0/1
user@host# set security zones security-zone dmz interfaces ge-0/0/2
```

4. Configure zones sf_2, asia_3, and DMZ to allow incoming VoIP traffic.

```
user@host# set security zones security-zone sf_2
user@host# set security zones security-zone sf_2 host-inbound-traffic
  system-services all
user@host# set security zones security-zone sf_2 host-inbound-traffic protocols
  all
user@host# set security zones security-zone asia_3
user@host# set security zones security-zone asia_3 host-inbound-traffic
  system-services all
user@host# set security zones security-zone asia_3 host-inbound-traffic protocols
  all
user@host# set security zones security-zone dmz
user@host# set security zones security-zone dmz host-inbound-traffic
  system-services all
user@host# set security zones security-zone dmz host-inbound-traffic protocols
  all
```

5. Configure static NAT on interface ge-0/0/ and source NAT on interface ge-0/0/2.

```
user@host# set security nat interface ge-0/0/1.0 static-nat 3.3.3.101/32 host
  10.1.1.101/32
user@host# set security nat interface ge-0/0/1.0 source-nat pool src-nat-pool
  address 2.2.2.10
```

6. Configure policies.

```
user@host# set security policies from-zone dmz to-zone asia_3 policy
  pol-dmz-to-asia_3 match source-address west_ca
user@host# set security policies from-zone dmz to-zone asia_3 policy
  pol-dmz-to-asia_3 match destination-address asia_gw
user@host# set security policies from-zone dmz to-zone asia_3 policy
  pol-dmz-to-asia_3 match application junos-mgcp
user@host# set security policies from-zone dmz to-zone asia_3 policy
  pol-dmz-to-asia_3 then permit
user@host# set security policies from-zone asia_3 to-zone dmz policy
  pol-asia_3-to-dmz match source-address asia_gw
user@host# set security policies from-zone asia_3 to-zone dmz policy
  pol-asia_3-to-dmz match destination-address 3.3.3.101
user@host# set security policies from-zone asia_3 to-zone dmz policy
  pol-asia_3-to-dmz match application junos-mgcp
user@host# set security policies from-zone asia_3 to-zone dmz policy
  pol-asia_3-to-dmz then permit
user@host# set security policies from-zone sf_2 to-zone dmz policy pol-sf_2-to-dmz
  match source-address sf_gw
```

```
user@host# set security policies from-zone sf_2 to-zone dmz policy pol-sf_2-to-dmz
    match destination-address west-ca
user@host# set security policies from-zone sf_2 to-zone dmz policy pol-sf_2-to-dmz
    match application junos-mgcp
user@host# set security policies from-zone sf_2 to-zone dmz policy pol-sf_2-to-dmz
    then permit
user@host# set security policies from-zone dmz to-zone sf_2 policy pol-dmz-to-sf_2
    match source-address west_ca
user@host# set security policies from-zone dmz to-zone sf_2 policy pol-dmz-to-sf_2
    match destination-address sf_gw
user@host# set security policies from-zone dmz to-zone sf_2 policy pol-dmz-to-sf_2
    match application junos-mgcp
user@host# set security policies from-zone dmz to-zone sf_2 policy pol-dmz-to-sf_2
    then permit
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 match source-address sf_gw
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 match destination-address asia_gw
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 match application junos-mgcp
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 then permit source-nat pool src-nat-pool
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 match source-address sf_gw
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 match destination-address asia_gw
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 match application junos-mgcp
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-sf_2-to-asia_3 then permit source-nat pool src-nat-pool
user@host# set security policies from-zone sf_2 to-zone asia_3 policy pol-intra-sf_2
    match source-address any
user@host# set security policies from-zone sf_2 to-zone asia_3 policy pol-intra-sf_2
    match destination-address any
user@host# set security policies from-zone sf_2 to-zone asia_3 policy pol-intra-sf_2
    match application any
user@host# set security policies from-zone sf_2 to-zone asia_3 policy pol-intra-sf_2
    then permit
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-intra-asia_3 match source-address any
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-intra-asia_3 match destination-address any
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-intra-asia_3 match application any
user@host# set security policies from-zone sf_2 to-zone asia_3 policy
    pol-intra-asia_3 then permit
```

Related Topics

- Understanding the MGCP ALG on page 266
- Configuring a Media Gateway in Subscribers' Homes on page 282

Verifying the MGCP Configuration

To verify the MGCP configuration, perform these tasks:

- Verifying MGCP Calls on page 293
- Verifying MGCP Endpoints on page 294
- Verifying MGCP Counters on page 295

Verifying the MGCP ALG

Purpose Display MGCP verification options.

Action From the CLI, enter the show security alg mgcp ? command.

```
user@host> show security alg mgcp ?
Possible completions:
  calls                Show MGCP calls
  counters             Show MGCP counters
  endpoints            Show MGCP endpoints
```

What it Means The output shows a list of all MGCP verification parameters. Verify the following information:

- All MGCP calls
- Counters for all MGCP c alls
- Information about all MGCP endpoints

Related Topics

- Understanding the MGCP ALG on page 266
- Verifying MGCP Calls on page 293
- Verifying MGCP Endpoints on page 294
- Verifying MGCP Counters on page 295

Verifying MGCP Calls

Purpose Display information about active calls.

Action From the J-Web interface, select **Monitor > ALGs > MGCP > Calls**. Alternatively, from the CLI, enter the show security alg mgcp calls command.

```
user@host> show security alg mgcp calls
Endpoint@GW                    Zone         Call ID                   RM Group
d001@101.50.10.1               Trust        10d55b81140e0f76          512
   Connection Id> 0
   Local SDP>  o: 101.50.10.1                x_o: 101.50.10.1
               c: 101.50.10.1/32206          x_c: 101.50.10.1/32206
   Remote SDP> c: 3.3.3.5/16928              x_c: 3.3.3.5/16928
```

```
Endpoint@GW                    Zone         Call ID                    RM Group
d001@3.3.3.5                   Untrust      3a104e9b41a7c4c9           511
   Connection Id> 0
   Local SDP>  o: 3.3.3.5                    x_o: 3.3.3.5
               c: 3.3.3.5/16928              x_c: 3.3.3.5/16928
   Remote SDP> c: 101.50.10.1/32206          x_c: 101.50.10.1/32206
```

What it Means The output displays information about all MGCP calls. Verify the following information:

- Endpoint
- Zone
- Call identifier
- Resource Manager group

Related Topics

- Understanding the MGCP ALG on page 266
- Verifying the MGCP ALG on page 293
- Verifying MGCP Endpoints on page 294
- Verifying MGCP Counters on page 295

Verifying MGCP Endpoints

Purpose Display information about MGCP endpoints.

Action From the J-Web interface, select Monitor>ALGs>MGCP>Endpoints. Alternatively, from the CLI, enter the show security alg mgcp endpoints command.

```
user@host> show security alg mgcp endpoints
Gateway: 101.50.10.1 Zone: Trust IP: 101.50.10.1 -> 101.50.10.1
   Endpoint            Trans #  Call #   Notified Entity
   d001                1        1        0.0.0.0/0->0.0.0.0/0
Gateway: 3.3.3.5 Zone: Untrust IP: 3.3.3.5 -> 3.3.3.5
   Endpoint            Trans #  Call #   Notified Entity
   d001                1        1        0.0.0.0/0->0.0.0.0/0
```

What it Means The output displays information about all MGCP endpoints. Verify the following information:

- Gateway IP address and zone of both endpoints
- Endpoint identifier, Trans #, Call #, and notified entity for each gateway

Related Topics

- Understanding the MGCP ALG on page 266
- Verifying the MGCP ALG on page 293
- Verifying MGCP Calls on page 293
- Verifying MGCP Counters on page 295

Verifying MGCP Counters

Purpose Display information about MGCP counter.

Action From the J-Web interface, select Monitor ALGs>MGCP>Counters. Alternatively, from the CLI, enter the show security alg mgcp counters command.

```
user@host> show security alg mgcp counters
MGCP counters summary:
Packets received             :284
Packets dropped              :0
Message received             :284
Number of connections        :4
Number of active connections :3
Number of calls              :4
Number of active calls       :3
Number of transactions       :121
Number of active transactions:52
Number of re-transmission    :68
MGCP Error Counters:
Unknown-method               :0
Decoding error               :0
Transaction error            :0
Call error                   :0
Connection error             :0
Connection flood drop        :0
Message flood drop           :0
IP resolve error             :0
NAT error                    :0
Resource manager error       :0
MGCP Packet Counters:
CRCX     :4       MDCX     :9       DLCX     :2
AUEP     :1       AUCX     :0       NTFY     :43
RSIP     :79      EPCF     :0       RQNT     :51
000-199  :0       200-299  :95      300-999  :0
```

What it Means The output displays information about all MGCP counters. Verify the following information:

- Summary of MGCP counters
- MGCP error counters
- MGCP packet counters

Related Topics

- Understanding the MGCP ALG on page 266
- Verifying the MGCP ALG on page 293
- Verifying MGCP Calls on page 293
- Verifying MGCP Endpoints on page 294

Understanding the RPC ALG

JUNOS Software supports basic Remote Procedure Call Application Layer Gateway (RPC ALG) services. RPC is a protocol that allows an application running in one address space to access the resources of applications running in another address space as if the resources were local to the first address space. The RPC ALG is responsible for RPC packet processing.

Before You Begin

For background information, read

- Application Layer Gateways (ALGs) on page 157

The RPC ALG in JUNOS Software supports the following services and features:

- Sun Microsystems RPC Open Network Computing (ONC)
- Microsoft RPC Distributed Computing Environment (DCE)
- Dynamic port negotiation
- Ability to allow and deny specific RPC services
- Static NAT and source NAT (with no port translation)
- RPC applications in security policies

Use the RPC ALG if you need RPC-based applications, such as NFS or Microsoft Outlook from the device. The RPC ALG functionality is enabled by default.

The RPC ALG can be applied by using:

- ALG Default Application

  Use the default application set:

  - application-set junos-sun-rpc (for control sessions)
  - application-set junos-sun-rpc-portmap (for data sessions)

  In the policy, use the predefined application set for Control and Data Connections.

- Default Control Application

  Use the predefined control via junos-sun-rpc:

  - Create an application for Data USER_DEFINED_DATA. You can make a set of your own data (for example, my_rpc_application_set)and use it in the policy.
  - In the policy, use the predefined application set for Control and Customized Data Application.
    - junos-sun-rpc

    - USER_DEFINED_DATA

- Custom Control And Custom Data Application

  Use the customized application:

  - Create an application for Control USER_DEFINED_CONTROL and Data USER_DEFINED_DATA
  - In the policy, use the user-defined application set for Control and Customized Data Application:
    - USER_DEFINED_CONTROL
    - USER_DEFINED_DATA

This topic covers:

- Sun RPC ALG on page 297
- Microsoft RPC ALG on page 299
- Related Topics on page 300

Sun RPC ALG

Sun Microsystems Remote Procedure Call—also known as Open Network Computing Remote Procedure Call (ONC RPC)—provides a way for a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.

J Series devices running JUNOS Software support the Sun RPC as a predefined service and allow and deny traffic based on a security policy you configure. The Application Layer Gateway (ALG) provides the functionality for J Series devices to handle the dynamic transport address negotiation mechanism of the Sun RPC and to ensure program number-based security policy enforcement. You can define a security policy to permit or deny all RPC requests, or to permit or deny by specific program number. The ALG also supports route and NAT mode for incoming and outgoing requests. The following SUN RPC topics are addressed in this section:

- “Typical RPC Call Scenario” on page 297
- “Sun RPC Services” on page 298
- “Customizing Sun RPC Services” on page 298

Typical RPC Call Scenario

When an application or a PC client calls a remote service, it needs to find the transport address of the service. In the case of TCP/UDP, the address is a port number. A typical procedure for this case is as follows:

1. The client sends the GETPORT message to the RPCBIND service on the remote machine. The GETPORT message contains the program number, and version and procedure number of the remote service it is attempting to call.
2. The RPCBIND service replies with a port number.
3. The client calls the remote service using the port number returned.
4. The remote service replies to the client.

A client also can use the CALLIT message to call the remote service directly, without determining the port number of the service. In this case, the procedure is as follows:

1. The client sends a CALLIT message to the RPCBIND service on the remote machine. The CALLIT message contains the program number, and the version and procedure number of the remote service it attempting to call.
2. RPCBIND calls the service for the client.
3. RCPBIND replies to the client if the call has been successful. The reply contains the call result and the services's port number.

Sun RPC Services

Table 49 on page 298 lists predefined Sun RPC services.

Table 49: Predefined Sun RPC Services

junos-sun-rpc-portmap-tcp	junos-sun-rpc-portmap
junos-sun-rpc-portmap-udp	

Customizing Sun RPC Services

All Sun RPC applications can be customized by using a predefined application set. For example, an application can be customized to open the control session only and not allow any data sessions:

```
application-set junos-sun-rpc {

application junos-sun-rpc-tcp;

application junos-sun-rpc-udp;

}
```

In the following example, the predefined application set allows data sessions only. It will not work without the control session.

```
application-set junos-sun-rpc-portmap {

application junos-sun-rpc-portmap-tcp;
```

application junos-sun-rpc-portmap-udp;

}

To customize all Sun RPC applications with predefined application sets, use both application sets in the policy:

application-set [junos-sun-rpc junos-sun-rpc-portmap]

NOTE: Microsoft Remote Procedure Call (MS RPC) applications are customized in the same way as SUN RPC applications.

Microsoft RPC ALG

Microsoft Remote Procedure Call (MS RPC) is the Microsoft implementation of the Distributed Computing Environment (DCE) RPC. Like the Sun RPC (see "Sun RPC ALG" on page 297), MS RPC provides a way for a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service program's Universal Unique IDentifier (UUID). The specific UUID is mapped to a transport address.

J Series devices running JUNOS Software support MS RPC as a predefined service and allow and deny traffic based on a policy you configure. The ALG provides the functionality for J Series devices to handle the dynamic transport address negotiation mechanism of the MS RPC, and to ensure UUID-based security policy enforcement. You can define a security policy to permit or deny all RPC requests, or to permit or deny by specific UUID number. The ALG also supports route and NAT mode for incoming and outgoing requests.

MS RPC Services in Security Policies

- 0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde
- 1453c42c-0fa6-11d2-a910-00c04f990f3b
- 10f24e8e-0fa6-11d2-a910-00c04f990f3b
- 1544f5e0-613c-11d1-93df-00c04fd7bd09

The corresponding TCP/UDP ports are dynamic. To permit them, you use a set applications *application-name* term *term-name* uuid *hex-number* statement for each number. The ALG maps the program numbers into dynamically negotiated TCP/UDP ports based on these four UUIDs and permits or denies the service based on a policy you configure.

Predefined Microsoft RPC Services

Table 50 on page 300 lists predefined Microsoft services, parameters associated with each service, and a brief description of each service. Parameters include Universal Unique Identifiers (UUIDs) and TCP/UDP source and destination ports. A UUID is a

128-bit unique number generated from a hardware address, a timestamp, and seed values.

Table 50: Predefined Microsoft RPC Services

junos-ms-rpc-portmap-tcp	junos-ms-rpc-portmap
junos-ms-rpc-portmap-udp	

Related Topics

- Understanding the SIP ALG on page 187
- Understanding the SCCP ALG on page 249
- Understanding the MGCP ALG on page 266
- Disabling and Enabling RPC ALG on page 300
- Verifying the RPC ALG Tables on page 301

Disabling and Enabling RPC ALG

The RPC ALG functionality is enabled by default and requires no configuration. You can disable the RPC ALG

Before You Begin
For background information, read "Understanding the RPC ALG" on page 296.

This topic covers:

- J-Web Configuration on page 300
- CLI Configuration on page 301
- Related Topics on page 301

J-Web Configuration

The Sun RPC ALG and Microsoft RPC ALG are enabled by default. To disable or re-enable the RPC ALGs:

- To disable or re-enable the Sun RPC ALG:
  1. Select Configure>CLI Tools>Point and Click CLI.
  2. Next to Alg, click Configure or Edit.
  3. Next to Sunrpc, click Configure or Edit.
  4. Select Yes for the Disable or Enable check box, as appropriate.

- To disable or re-enable the Microsoft RPC ALG:
  1. Select Configure>CLI Tools>Point and Click CLI.
  2. Next to Alg, click Configure or Edit.
  3. Next to Msrpc, click Configure or Edit.
  4. Select Yes for the Disable or Enable check box, as appropriate.

CLI Configuration

The Sun RPC ALG and Microsoft RPC ALG are enabled by default.

To disable or re-enable the Sun RPC ALG:

- To disable the ALG, enter the following command:

  user@host# **set security alg sunrpc disable**

- To re-enable the ALG, enter the following command:

  user@host# **delete security alg sunrpc**

To disable or re-enable the Microsoft RPC ALG:

- To disable the ALG, enter the following command:

  user@host# **set security alg msrpc disable**

- To re-enable the ALG, enter the following command:

  user@host# **delete security alg msrpc**

Related Topics

- Understanding the RPC ALG on page 296
- Verifying the RPC ALG Tables on page 301

Verifying the RPC ALG Tables

To verify the RPC ALG, perform these tasks:

- Display the Sun RPC port mapping table
- Display the Microsoft UUID mapping table

Display the Sun RPC Port Mapping Table

Purpose Display the Sun RPC port mapping table. The Sun RPC ALG monitors packets on TCP or UDP port 111.

Action From the CLI, enter the show security alg sunrpc portmap command.

```
user@host> show security alg sunrpc portmap
IP                Port    Protocol   Program
10.209.17.127     32835   TCP        100005
10.209.17.127     2049    UDP        100003
10.209.17.127     111     UDP        100000
10.209.17.127     111     TCP        100000
```

What it Means The output displays the server IP address, TCP or UDP port, and corresponding program number for each program being called with the Sun RPC ALG process:

- Server's IP address and port that maps to the program
- Protocol used to support the process
- Program ID number

Display the Microsoft RPC UUID Mapping Table

Purpose Display the Microsoft RPC UUID-to-Object ID (OID) mapping table. The Microsoft RPC ALG monitors packets on TCP port 135. (You can display the Microsoft RPC port mapping table with the show security alg msrpc portmap command.)

Action From the CLI, enter the show security alg msrpc object-id-map command.

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
1be617c0-31a5-11cf-a7d8-00805f48a135    0x80000020
e3514235-4b06-11d1-ab04-00c04fc2dcd2    0x80000002
67df7c70-0f04-11ce-b13f-00aa003bac6c    0x80000014
```

What it Means The output displays mapping of the UUID to the OIDs. Verify the following information:

- Microsoft RPC UUID
- Microsoft RPC object ID

Related Topics

- Understanding the RPC ALG on page 296
- Disabling and Enabling RPC ALG on page 300

Part 5
User Authentication

- Firewall User Authentication on page 305
- Infranet Authentication on page 335

Chapter 9

Firewall User Authentication

Firewall user authentication enables administrators to restrict and permit users (firewall users) accessing protected resources (different zones) behind a firewall based on their source IP address and other credentials.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Firewall User Authentication Overview on page 305
- Understanding Authentication Schemes on page 307
- Configuring for Pass-Through Authentication on page 310
- Configuring for Web Authentication on page 315
- Understanding Client Groups for Firewall Authentication on page 319
- Configuring for External Authentication Servers on page 322
- Understanding SecurID User Authentication on page 326
- Configuring the SecurID Server on page 327
- Displaying the Authentication Table on page 329
- Understanding Banner Customization on page 330
- Customizing a Banner on page 330
- Configuring Firewall Authentication—Quick Configuration on page 332
- Verifying Firewall User Authentication on page 333

Firewall User Authentication Overview

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. JUNOS Software supports the following types of users:

- Administrators. For more information, see the *JUNOS Software Administration Guide*
- Point-to-Point Protocol (PPP) users. For more information, see the *JUNOS Software Administration Guide*

- Firewall users. Firewall user authentication enables administrators to restrict and permit users (firewall users) accessing protected resources (different zones) behind a firewall based on their source IP address and other credentials.

This topic covers:

- Authentication, Authorization, and Accounting (AAA) Servers on page 306
- Types of Firewall User Authentication on page 306
- Related Topics on page 307

Authentication, Authorization, and Accounting (AAA) Servers

AAA provides an extra level of protection and control for user access in the following ways:

- Authentication determines the firewall user.
- Authorization determines what the firewall user can do.
- Accounting determines what the firewall user did on the network.

You can use authentication alone or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

Once the user's credentials are collected, they are processed in one of the following ways:

- Administrative authentication supports the following types of servers:
  - local
  - RADIUS
  - TACACS+

    For more information on administrative authentication, see the *JUNOS Software Administration Guide*).
- Firewall user authentication supports the following types of servers:
  - Local authentication and authorization
  - RADIUS authentication and authorization (compatible with Funk RADIUS server)
  - LDAP authentication only (supports LDAP version 3 and compatible with Windows AD)
  - SecurID authentication only (using an RSA SecurID external authentication server)

Types of Firewall User Authentication

JUNOS Software supports the following two types of firewall user authentication:

- Pass-Through Authentication—A host or a user from one zone tries to access resources on another zone. You must use an FTP, Telnet, or HTTP client to access the IP address of the protected resource and to get authenticated by the firewall. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.
- Web Authentication—Users try to connect, using HTTP, to an IP address on the device that is enabled for Web authentication; in this scenario, you do not use HTTP to get to the IP address of the protected resource. You are prompted for the username and password that are verified by the device. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

Related Topics

- Understanding Authentication Schemes on page 307
- Understanding Client Groups for Firewall Authentication on page 319
- Understanding Banner Customization on page 330

Understanding Authentication Schemes

After you define firewall users, you can create a policy that requires the users to authenticate themselves through one of two authentication schemes.

Before You Begin
For background information, read "Firewall User Authentication Overview" on page 305.

The first scheme authenticates users when FTP, HTTP, or Telnet traffic matching a policy requiring authentication reaches the Juniper Networks device. In the second scheme, users authenticate themselves before sending traffic (of any kind—not just FTP, HTTP, or Telnet) that has a policy requiring user authentication.

This topic covers:

- Pass-Through Authentication on page 307
- Web Authentication on page 308
- Related Topics on page 310

Pass-Through Authentication

When a user attempts to initiate an HTTP, an FTP, or a Telnet connection request that has a policy requiring authentication, the Juniper Networks device intercepts the request and prompts the user to enter a name and password. Before granting permission, the device validates the username and password by checking them

against those stored in the local database or on an external authentication server. See Figure 38 on page 308.

Figure 38: Policy Lookup for a User

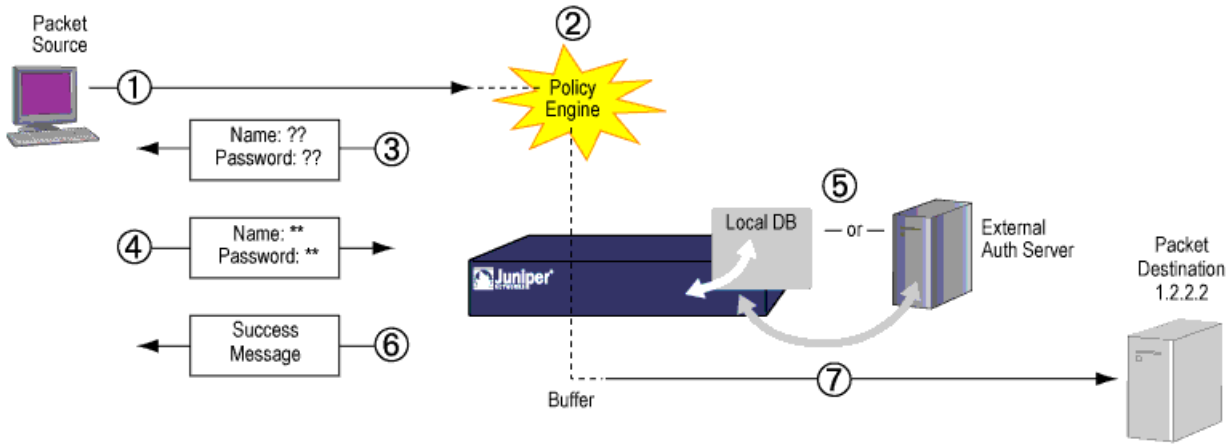

1. A client user sends an FTP, an HTTP, or a Telnet packet to 1.2.2.2.
2. The Juniper Networks device intercepts the packet, notes that its policy requires authentication from either the local database or an external authentication server, and buffers the packet.
3. The Juniper Networks device prompts the user for login information through FTP, HTTP, or Telnet.
4. The user replies with a username and password.
5. The Juniper Networks device either checks for an authentication user account on its local database or it sends the login information to the external authentication server as specified in the policy.
6. Finding a valid match (or receiving notice of such a match from the external authentication server), the Juniper Networks device informs the user that the login has been successful.
7. The Juniper Networks device forwards the packet from its buffer to its destination IP address 1.2.2.2.

After a Juniper Networks device authenticates a user at a particular source IP address, it subsequently permits traffic—as specified in the policy requiring authentication through pass through—from any other user at that same address. This might be the case if the user originates traffic from behind a NAT device that changes all original source addresses to a single translated address.

Web Authentication

Web authentication is an alternate form of firewall user authentication. Instead of pointing to the resource you want to connect to from your client browser, you point the browser to an IP address on the device that is enabled for Web authentication. This initiates an HTTP session to the IP address hosting the Web Authentication

feature on the device. The device then prompts you for your username and password and caches the result in the device. Later when traffic encounters a web-authentication policy, you are allowed or denied access based on the prior Web authentication results as shown in Figure 39 on page 309.

Figure 39: Web Authentication Example

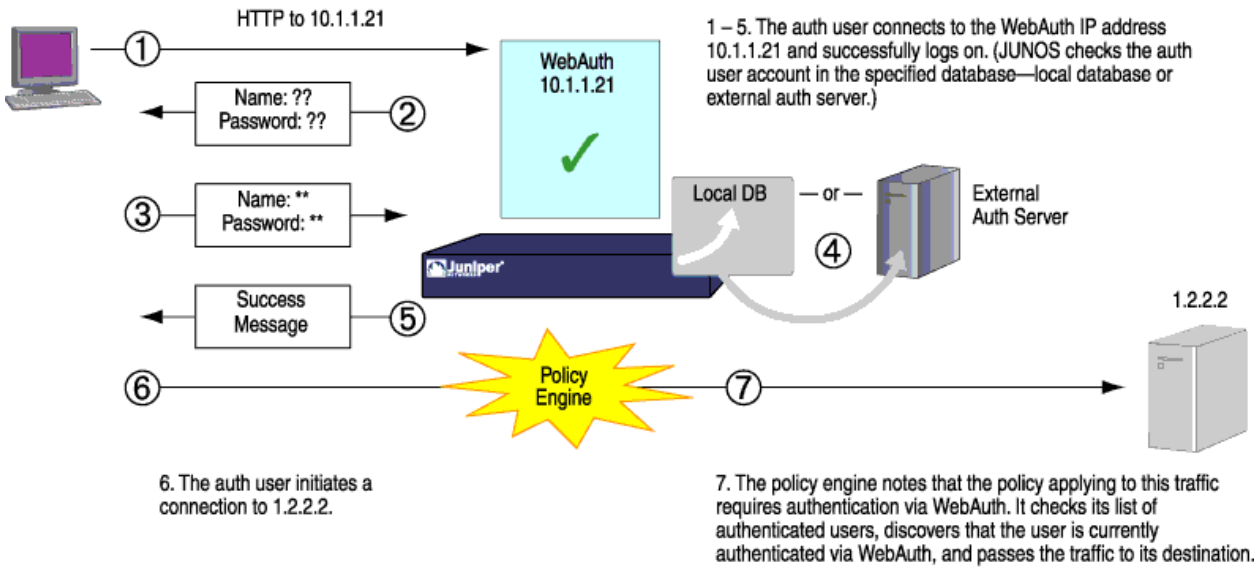

Follow these Web authentication guidelines:

- You can leave the default Web authentication server as the local database or you can choose an external auth server for the role. The default Web Authentication profile determines if the user authenticates using the local database or the external authentication server. An access profile stores usernames and passwords of users or points to external authentication servers where such information is stored.
- The Web authentication address must be in the same subnet as the interface that you want to use to host it. For example, if you want authentication users to connect using Web Authentication through ethernet3, which has IP address 1.1.1.1/24, then you can assign Web authentication an IP address in the 1.1.1.0/24 subnet.
- You can put a Web authentication address in the same subnet as the IP address of any physical interface or virtual security interface (VSI). (For information about different types of interfaces, see “Security Zones and Interfaces” on page 47.)
- You can put Web authentication addresses on multiple interfaces.
- After a device authenticates a user at a particular source IP address, it subsequently permits traffic—as specified in the policy requiring authentication through Web authentication—from any other user at that same address. This might be the case if the user originates traffic from behind a NAT device that changes all original source addresses to a single translated address.
- With Web authentication enabled, any HTTP traffic to the IP address will get the Web Authentication login page instead of the admin login page. Disabling this

option will show the admin login page (assuming that [system services web-management HTTP] is enabled.

- We recommend that you have a separate primary or preferred IP address, if an address is used for Web authentication.

Related Topics

- Configuring for Pass-Through Authentication on page 310
- Configuring for Web Authentication on page 315

Configuring for Pass-Through Authentication

Pass-through firewall user authentication occurs when the client is trying to access a destination on another zone using FTP, Telnet, or HTTP. After authenticating successfully, the firewall acts as a proxy for an FTP, Telnet, or HTTP server so that it can first authenticate the user before allowing access to the actual FTP, Telnet, or HTTP server behind the firewall. See Figure 40 on page 310.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Read "Understanding Authentication Schemes" on page 307.
3. Read "Firewall User Authentication Overview" on page 305.

Figure 40: Configuring Pass-Through Firewall Authentication

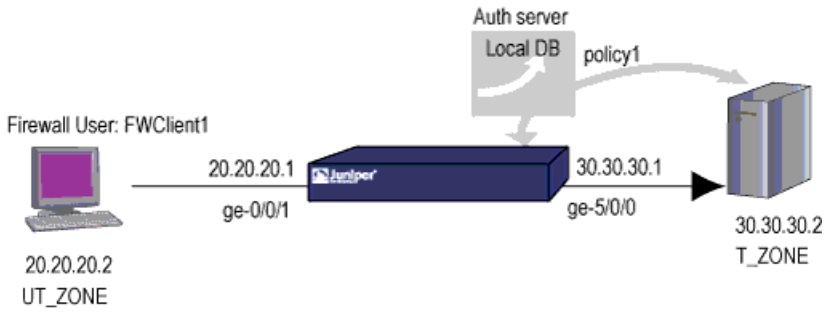

To configure pass-through firewall authentication, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 311
- CLI Configuration on page 313
- Related Topics on page 314

J-Web Configuration

To create IP addresses for the interfaces on the device using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Interfaces, click **Configure** or **Edit**.
3. Next to Interface, click **Add new entry**.
4. In the Interface name box, type ge-0/0/1.
5. Next to Unit, click **Add new entry**.
6. In the interface unit number, type 0.
7. Under Family, select **Inet** and click **OK**.
8. Next to Address book, click **Configure** or **Edit**.
9. Next to Address, click Add new entry.
10. In the Address name box, type 20.20.20.1/24 and click OK.
11. Next to Address, click Add new entry.
12. In the Address name box, type 20.20.20.2/24 and click OK.
13. To configure another interface ge-5/0/0 and more addresses like 30.30.30.1/24 and 30.30.30.1/24, repeat Step 2 through Step 9 and click OK.

To create an access profile:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Profile, click **Add new entry**.
4. In the Profile name box, type FWAuth.
5. Next to Client, click **Add new entry**.
6. In the Name box, type FWClient1.
7. Next to Firewall User, click **Configure** or **Edit**.
8. In the Password box, type pwd and click OK.

To add the FWAuth profile for pass-through firewall authentication:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Firewall Authentication, click **Configure** or **Edit**.
3. Next to Pass through, click Configure or Edit.
4. In the Default profile box, type FWAuth.

To define a success banner for Telnet sessions:

1. Next to Telnet box, click Configure or Edit.
2. Next to Banner, click Configure or Edit.
3. In the Success box, type " WELCOME TO JUNIPER TELNET SESSION" and click OK.

To create security zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type UT-ZONE.
6. Next to Host inbound traffic, click **Configure** or **Edit**.
7. Next to System services, click **Add new entry**.
8. From the Service name list, select **all** and click **OK**.
9. To configure an interface for the created security zone, corresponding to the security zone, click **Edit**.
10. Next to Interfaces, click **Add new entry**.
11. In the Interface unit box, type ge-0/0/1.0 and click **OK**.
12. Next to Protocols, click **Add new entry**.
13. Next to the Protocol name box, type all and click **OK**.
14. To specify another interface fe-5/0/0.0 for the zone, repeat Step 9 and Step 10 and click **OK**.
15. To add another security zone T-ZONE, repeat Step 3 through Step 7 and click **OK**.

To assign a security policy to the zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policy, select the check box and click **Configure**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type UT-ZONE.
6. In the To zone name box, type T-ZONE.
7. Next to Policy, click **Add new entry.**
8. In the Policy name box, type Policy-W.
9. Select the **Match** check box.
10. Select the **Then** check box.
11. Next to Match, click **Configure** or **Edit**.
12. From the Source address choice list, select **Source address**.

13. Next to Source address, click **Add new entry**.
14. From the Value keyword list, select **any** and click **OK**.
15. From the Destination address choice list, select **Destination address**.
16. Next to Destination address, click **Add new entry**.
17. From the Value keyword list, select **any** and click **OK**.
18. From the Application Choice list, select **Application**.
19. Next to Application, click **Add new entry**.
20. In the Value keyword list, type junos-telnet and click **OK**.
21. Next to Then, click **Configure** or **Edit**.
22. From the Action list, select **Permit** and click **OK**.
23. Next to Permit, click **Configure** or **Edit**.
24. Next to Firewall Authentication, click **Configure** or **Edit**.
25. From the Auth type list, select **Pass through**.
26. Next to Pass through, click **Configure** or **Edit**.
27. In the Client match box, type FW Client1 and click **OK**.
28. To authenticate the firewall user authentication, telnet FWClient1 to host2.
29. If you are finished configuring the device, commit the configuration.
30. To check the configuration, see "Verifying Firewall User Authentication" on page 333.

CLI Configuration

To configure the device for pass-through firewall authentication as shown in Figure 40 on page 310, follow these steps:

1. Create IP addresses for the interfaces on the device.

```
user@host# set interfaces ge-0/0/1
user@host# set unit 0 family inet address 20.20.20.1/24
user@host# set unit 0 family inet address 20.20.20.2/24
user@host# set interfaces ge-5/0/0
user@host# set unit 0 family inet address 30.30.30.1/24
user@host# set unit 0 family inet address 30.30.30.2/24
```

2. Create an access profile, FWAUTH, for FWClient1 and specify a password, pwd.

```
user@host# set access profile FWAUTH client FWClient1 firewall-user password
  pwd
```

3. Add the FWAUTH profile for pass-through firewall authentication and define a success banner for Telnet sessions.

```
user@host# set access firewall-authentication pass-through default-profile
  FWAUTH
```

```
user@host# set access firewall-authentication pass-through telnet banner
  success "WELCOME TO JUNIPER TELNET SESSION"
```

4. Create security zones.

```
user@host# set security zones security-zone UT-ZONE host-inbound-traffic
  system-services all
user@host# set security zones security-zone UT-ZONE interfaces ge-0/0/1.0
  host-inbound-traffic protocols all
user@host# set security zones security-zone T-ZONE host-inbound-traffic
  system-services all
user@host# set security zones security-zone T-ZONE interfaces fe-5/0/0.0
  host-inbound-traffic protocols all
```

5. Assign a security policy, policy1, to the zones.

```
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
  policy1 match source-address any
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
  policy1 match destination-address any
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
  policy1 match application junos-telnet
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
  policy1 then permit firewall-authentication pass-through client-match
  FWclient1
```

6. Use Telnet to autheticate firewall user, FWClient1, to host2.

```
regress@FWClient1# run telnet 30.30.30.2
Trying 30.30.30.2...
Connected to 30.30.30.2.
Escape character is '^]'.
Firewall User Authentication
Username: FWClient1
Password:***
WELCOME TO JUNIPER TELNET SESION
Host1 (ttyp0)
login: regress
Password:
— JUNOS 8.5R1.1 built 2007-10-12 13:30:18 UTC
%
```

7. If you are finished configuring the device, commit the configuration.
8. To check the configuration, see “Verifying Firewall User Authentication” on page 333.

Related Topics

- Firewall User Authentication Overview on page 305
- Understanding Client Groups for Firewall Authentication on page 319

Configuring for Web Authentication

To enable Web authentication, you must specify the IP address of the device hosting the HTTP session.

Before You Begin

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Read "Understanding Authentication Schemes" on page 307.
3. Read "Firewall User Authentication Overview" on page 305.

These settings are used if the firewall user accessing a protected resource wants to be authenticated by directly accessing the Web server for Web authentication.

When referencing the following example, note that to enable or disable Web authentication, you must add or remove the Web authentication HTTP flag under the interface's address hierarchy.

The following example shows how to set up a policy allowing access to FWClient1, when traffic encounters a policy that has Web authentication enabled (Policy-W). FWClient1 should already have been authenticated through the Web Authentication login page as explained in "Configuring for External Authentication Servers" on page 322.

To configure a server for Web authentication in the scenario shown in Figure 39 on page 309, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 315
- CLI Configuration on page 318
- Related Topics on page 319

J-Web Configuration

To configure the device for Web authentication using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Interfaces, click **Configure** or **Edit**.
3. Next to Interface, click **Add new entry**.
4. In the Interface name box, type ge-0/0/1.
5. Next to Unit, click **Add new entry**.
6. In the interface unit number, type 0.
7. Under Family, select **Inet** and click **OK**.
8. Next to Address, click **Add new entry**.

9. In the Address name box, type 20.20.20.1/24 and click **OK**.
10. Next to the Web authentication box, click **Configure** or **Edit**.
11. Select the **Http** check box and click **OK**.
12. Next to Address, click **Add new entry**.
13. In the Address name box, type 20.20.20.2/24 and click **OK**.
14. To configure another interface ge-5/0/0, repeat Step 2 through Step 9.
15. To configure more addresses like 30.30.30.1/24 and 30.30.30.1/24, repeat Step 2 through Step 4 and click **OK**.

To create an access profile:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Profile, click **Add new entry**.
4. In the Profile name box, type WebAuth.
5. Next to Client, click **Add new entry**.
6. In the Name box, type FWClient1.
7. Next to Firewall User, click **Configure** or **Edit**.
8. In the Password box, type pwd and click **OK**.

To add the WebAuth profile for firewall Web authentication:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Firewall Authentication, click **Configure** or **Edit**.
4. Next to Web authentication, click **Configure** or **Edit**.
5. In the Default profile box, type FWAuth.

To define a success banner for Telnet sessions:

1. Next to Telnet box, click **Configure** or **Edit**.
2. Next to Banner, click **Configure** or **Edit**.
3. In the Success box, type " WELCOME TO JUNIPER TELNET SESSION" and click **OK**.

To create security zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure** or **Edit** .
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type UT-ZONE.

6. Next to Host inbound traffic, click **Configure**.
7. Next to System services, click **Add new entry**.
8. From the Service name list, select **all** and click **OK.**
9. Next to Interfaces, click **Add new entry**.
10. In the Interface unit box, type ge-0/0/1.0 and click **OK**.
11. Next to Host inbound traffic, click **Configure** or **Edit**.
12. Next to Protocols, click **Add new entry**.
13. From the Protocol name list, select **all** and click **OK.**
14. To specify another interface fe-5/0/0.0 for the zone, repeat Step 9 and Step 10, and click **OK**.
15. To add another security zone T-ZONE, repeat Step 3 through Step 7 and click **OK**.

To assign a security policy to the zone:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policy, select the check box and click **Configure**.
4. Next to Policy, click **Add new entry**.
5. In the From zone name box, type UT-ZONE.
6. In the To zone name box, type T-ZONE.
7. Next to Policy, click **Add new entry**.
8. In the Policy name box, type Policy-W.
9. Select the **Match** check box.
10. Select the **Then** check box.
11. Next to Match, click **Configure** or **Edit**.
12. From the Source address choice list, select **Source address**.
13. Next to Source address, click **Add new entry**.
14. From the Value keyword list, select **any** and click **OK**.
15. From the Destination address choice list, select **Destination address**.
16. Next to Destination address, click **Add new entry.**
17. From the Value keyword list, select **any** and click **OK**.
18. From the Application Choice list, select **Application**.
19. Next to Application, click **Add new entry**.
20. In the Value keyword list, type any and click **OK**.
21. Next to Then, click **Configure** or **Edit**.
22. From the Action list, select **Permit** and click OK.
23. Next to Permit, click **Configure** or **Edit**.

24. Next to Firewall Authentication, click **Configure** or **Edit**.
25. From the Auth type list, select **Web authentication**.
26. Next to Web authentication, click **Configure** or **Edit**.
27. In the Client match box, type FW Client1 and click **OK**.

To activate the HTTP daemon on your device:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to System, click **Configure** or **Edit**.
3. Next to Services, select the check box and click **Configure** or **Edit**.
4. Next to Web management, click **Configure** or **Edit**.
5. Select the **Http** check box and click **OK**.
6. To authenticate the firewall user authentication, point the browser to the Web Authentication IP (20.20.20.1).
7. If you are finished configuring the device, commit the configuration.
8. To check the configuration, see "Verifying Firewall User Authentication" on page 333

CLI Configuration

1. Create IP addresses for the interfaces on the device.

```
user@host# set interfaces ge-0/0/1
user@host# set unit 0 family inet address 20.20.20.1/24 web authentication
   http
user@host# set unit 0 family inet address 20.20.20.2/24
user@host# set interfaces fe-5/0/0
user@host# set unit 0 family inet address 30.30.30.1/24
user@host# set unit 0 family inet address 30.30.30.2/24
```

2. Create an access profile, WEBAUTH, for FWClient1 and specify a password, pwd.

```
user@host# set access profile WEBAUTH client FWClient1 firewall-user password
   pwd
```

3. Add the above WEBAUTH profile for firewall Web authentication and define a success banner for Telnet sessions.

```
user@host# set access firewall-authentication web-authentication default-profile
   WEBAUTH banner success "WEB AUTH LOGIN SUCCESS"
```

4. Create security zones.

```
user@host# set security zones security-zone UT-ZONE host-inbound-traffic
   system-services all
user@host# set security zones security-zone UT-ZONE interfaces ge-0/0/1.0
   host-inbound-traffic protocols all
user@host# set security zones security-zone T-ZONE host-inbound-traffic
   system-services all
```

```
user@host# set security zones security-zone T-ZONE interfaces fe-5/0/0.0
    host-inbound-traffic protocols all
```

5. Assign a security policy, policy-W, to the zones.

```
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
    policy-W match source-address any
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
    policy-W match destination-address any
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
    policy-W match application any
user@host# set security policies from-zone UT-ZONE to-zone T-ZONE policy
    policy-W then permit firewall-authentication web-authentication client-match
    FWclient1
```

6. Activate the HTTP daemon on your device.

```
user@host# set system services web-management http
```

7. Firewall user FWClient1 does the following to get authenticated:
   a. Points the browser to the Web authentication IP (20.20.20.1) to get authenticated first.
   b. Starts traffic to access resources specified by policy, policy-W.
8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see “Verifying Firewall User Authentication” on page 333.

   The following screen appears after the firewall user is authenticated.

Related Topics

- Understanding Client Groups for Firewall Authentication on page 319
- Firewall User Authentication Overview on page 305

Understanding Client Groups for Firewall Authentication

To manage a number of firewall users, you can create user or client groups and store the information either on the local Juniper Networks device or on an external RADIUS or LDAP server.

Before You Begin
For background information, read “Firewall User Authentication Overview” on page 305.

A client group is a list of groups that the client belongs to. As with client-idle timeout, a client group is used only if the external authentication server does not return a value in its response (for example, LDAP servers do not return such information).

The RADIUS server sends the client's group information to the Juniper Networks device using Juniper VSA (46). The client-match portion of the policy accepts a string that can either be the username or groupname the client belongs to.

Example 1 shows client groups configured for a client. If a client group is not defined for the client, then the client group under the hierarchy access> profile>session-options is used.

Two example configurations are shown below. The first shows how to configure a local user called client1 for groups G1, G2, and G3 using J-Web and the CLI configuration editor. Within this example, client groups are configured for a client. If a client group is not defined for the client, then the client group under the hierarchy access> profile>session-options is used. The example configuration shows how to configure the default client group for all users in a profile called managers using J-Web and the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 320
- CLI Configuration on page 321
- J-Web Configuration on page 321
- CLI Configuration on page 321
- Related Topics on page 322

J-Web Configuration

To configure a local user called client1 for groups G1, G2, and G3 using the J-Web Configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Profile, click **Add new entry**.
4. Next to Profile, click Add new entry.
5. In the Profile name box, type Managers.
6. Next to Client, click Add new entry.
7. In the Name box, type Client1.
8. Next to Firewall User, click **Configure** or **Edit**.
9. In the Password box, type test and click OK.
10. Next to Client group, click Add new entry.
11. In the Value box, type [g1 g2 g3] and click OK.

12. If you are finished configuring the device, commit the configuration.
13. To check the configuration, see "Verifying Firewall User Authentication" on page 333.

CLI Configuration

To configure a local user called client1 for groups G1, G2, and G3, enter the following:

```
user@host# set access profile managers client client1 firewall-user password test
client-group [g1 g2 g3]
```

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

To check the configuration, see "Verifying Firewall User Authentication" on page 333.

J-Web Configuration

To configure a default client group for all users using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure**.
3. Next to Profile, click **Add new entry**.
4. In the Profile name box, type **Managers**.
5. Next to Session options, click **Configure** or **Edit**.
6. Next to Client group, click **Add new entry**.
7. In the Value box, type [m1 m2] and click **OK**.
8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see "Verifying Firewall User Authentication" on page 333.

CLI Configuration

To configure a default client group for all users in a profile called managers, enter the following:

```
user@host# set access profile managers session-options client-group [m1 m2]
```

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

To check the configuration, see "Verifying Firewall User Authentication" on page 333.

If a client belonging to the profile, managers, does not have any group information configured, then the client automatically inherits the groups m1 and m2.

The reason to have a single database for different types of clients (except admins) is based on the assumption that a single client can be of multiple types. For example, a firewall user client can also be an L2TP client.

Related Topics

- Configuring for External Authentication Servers on page 322

Configuring for External Authentication Servers

You can put several user accounts together to form a user group, which you can store on the local database or on a RADIUS, LDAP, or SecurID server. When you reference an authentication user group and an external authentication server in a policy, the traffic matching the policy provokes an authentication check.

In this example, the access profile called prof_1 is configured for external authentication. Two RADIUS servers and one LDAP server are configured in the access profile. However, the order of authentication specifies RADIUS server only, so if the RADIUS server authentication fails, then the firewall user fails to authenticate. The local database is not accessed.

NOTE: If the firewall clients are authenticated by the RADIUS server, then the group-membership VSA returned by the RADIUS server should contain alpha, beta, or gamma client-groups in the RADIUS server configuration or in the access profile, prof_1. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored.

Before You Begin
For background information, read “Firewall User Authentication Overview” on page 305.

To configure a server for external authentication, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 322
- CLI Configuration on page 325
- Related Topics on page 326

J-Web Configuration

To specify the RADIUS server for external authentication order using the J-Web configuration editor:

1. Select **ConfigureCLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Profile, click **Add new entry**.
4. In the Profile name box, type prof1.
5. Next to Authentication order, click **Add new entry**.
6. From the Value choice list, select **radius** and click **OK**.

To configure firewall user (ClientsA-E) and assign firewall users (ClientA and ClientB) to client groups alpha, beta, and gamma:

1. Next to Client, click **Add new entry**.
2. In the Name box, type **ClientA**.
3. Next to Client group, click **Add new entry**.
4. In the Value box, type alpha and click **OK**.
5. To specify another client group, in the Value box, click **Add new entry**, type beta in the Name box, and click **OK**.
6. To specify another client group, in the Value box, click **Add new entry**, type gamma in the Name box, and click **OK**.
7. Next to Firewall User, click **Configure** or **Edit**.
8. In the Password box, type pwd1 and click OK.
9. Next to Client, click Add new entry.
10. In the Name box, type ClientB and click OK.
11. In the Value box, type alpha and click OK.
12. To specify another client group, in the Value box, type beta and click **OK**.
13. Next to Firewall User, click **Configure** or **Edit**.
14. In the Password box, type pwd3 and click **OK**.
15. To specify another client, next to Client, click **Add new entry**.
16. In the Name box, type ClientC and click **OK**.
17. Next to Firewall User, click **Configure** or **Edit**.
18. In the Password box, type pwd4 and click **OK**.
19. To specify another client, next to Client, click **Add new entry**.
20. In the Name box, type ClientD and click **OK**.
21. Next to Firewall User, click **Configure** or **Edit**.
22. In the Password box, type pwd5 and click **OK**.
23. To specify another client, next to Client, click **Add new entry**.
24. In the Name box, type ClientE and click **OK**.

25. Next to Firewall User, click **Configure** or **Edit**.
26. In the Password box, type pwd2 and click OK.

To configure client groups in the session options:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Profile, click **Add new entry**.
4. In the Profile name box, type prof1.
5. Next to Session options, click **Configure**.
6. In the Value box, type **u1** and click OK.
7. To specify another client group, in the Value box, type alpha and click **OK**.
8. To specify another client group, in the Value box, type gamma and click **OK**.
9. In the Client idle timeout box, type 255.
10. In the Client session timeout box, type 4 and click **OK**.

To configure the IP address for the LDAP server and LDAP server options:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Ldap options, click **Configure** or **Edit**.
4. In the Base distinguished name box, type CN=Users,DC=screenos,DC=spg,DC=juniper,DC=net
5. From the Search type list, select **Search**.
6. Next to Search, click **Configure** or **Edit**.
7. In the Search filter box, type sAMAccountName= and click **OK**.
8. Select the **Admin search** check box and click **Configure** or **Edit**.
9. In the Distinguished name box, type cn=administrator,cn=users,dc=screenos,dc=spg,dc=juniper,dc=net.
10. In the Password box, type pwd10 and click **OK**.
11. Next to Ldap server, click **Add new entry**.
12. In the Name box, type 3.3.3.3 and click OK.

To configure the IP addresses for the two RADIUS servers:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure**.
3. Next to Radius server, click **Add new entry**.
4. In the Address box, type 4.4.4.4 and click **OK**.
5. In the Secret box, type any unreadable data.

6. In the Retry box, type 10 and click **OK**.
7. Next to Radius server, click **Add new entry**.
8. In the Address box, type 5.5.5.5 and click **OK**.
9. In the Secret box, type any unreadable data.
10. If you are finished configuring the device, commit the configuration.
11. To check the configuration, see "Verifying Firewall User Authentication" on page 333

CLI Configuration

To configure the device for external authentication using a RADIUS server follow these steps:

1. Specify the RADIUS server for external authentication order. This restricts firewall users to authenticate through the RADIUS server only. If the RADIUS server authentication fails and the default password (local database) option is not specified, the firewall user is locked out.

```
user@host# set access profile prof_1 authentication-order radius
```

2. Configure firewall user (ClientsA-E) and assign firewall users (ClientA and ClientB) to client groups alpha, beta, and gamma.

```
user@host# set access profile prof_1 client clientA client-group alpha
user@host# set access profile prof_1 client clientA client-group beta
user@host# set access profile prof_1 client clientA client-group gamma
user@host# set access profile prof_1 client clientA firewall-user password pwd1
user@host# set access profile prof_1 client clientB client-group alpha
user@host# set access profile prof_1 client clientB client-group beta
user@host# set access profile prof_1 client clientB firewall-user password pwd3
user@host# set access profile prof_1 client clientC firewall-user password pwd4
user@host# set access profile prof_1 client clientD firewall-user password pwd5
user@host# set access profile prof_1 client clientE firewall-user password pwd2
```

3. Configure client groups in the session options.

```
user@host# set access profile prof_1 session-options client-group u1
user@host# set access profile prof_1 session-options client-group alpha
user@host# set access profile prof_1 session-options client-group gamma
user@host# set access profile prof_1 session-options client-idle-timeout 255
user@host# set access profile prof_1 session-options client-session-timeout 4
```

4. Configure the IP address for the LDAP server and LDAP server options.

```
user@host# set access profile prof_1 ldap-options base-distinguished-name
CN=Users,DC=screenos,DC=spg,DC=juniper,DC=net
user@host# set access profile prof_1 ldap-options search search-filter
  sAMAccountName=
user@host# set access profile prof_1 ldap-options search admin-search
  distinguished-name
cn=administrator,cn=users,dc=screenos,dc=spg,dc=juniper,dc=net
```

```
user@host# set access profile prof_1 ldap-options search admin-search password
   pwd10
cn=administrator,cn=users,dc=screenos,dc=spg,dc=juniper,dc=net
user@host# set access profile prof_1 ldap-server 3.3.3.3
```

5. Configure the IP addresses for the two RADIUS servers.

```
user@host# set access profile prof_1 radius-server 4.4.4.4
user@host# set access profile prof_1 radius-server 4.4.4.4 secret
user@host# set access profile prof_1 radius-server 4.4.4.4 retry 10
user@host# set access profile prof_1 radius-server 5.5.5.5 secret
```

6. If you are finished configuring the device, commit the configuration.
7. To check the configuration, see "Verifying Firewall User Authentication" on page 333

Related Topics

- "Firewall User Authentication Overview" on page 305
- "Understanding Client Groups for Firewall Authentication" on page 319

Understanding SecurID User Authentication

SecurID is an authentication method that allows users to enter either static or dynamic passwords as their credentials. A dynamic password is a combination of a user's PIN and a randomly generated token that is valid for a short period of time, approximately one minute. A static password is set for the user on the SecurID server. For example, the SecurID server administrator might set a temporary static password for a user who lost his or her SecurID token.

Before You Begin
For background information, read "Firewall User Authentication Overview" on page 305.

When a user attempts to access a resource protected by a policy and SecurID is configured in the profile authentication-order parameter as either the only authentication mode or the first one to be used, the device forwards the user's credentials to the SecurID server for authentication. If the user enters valid values, the user is allowed access to the requested resource.

The SecurID server includes a feature that presents a user with a challenge if the user provides wrong credentials repeatedly. However, JUNOS Software does not support the challenge feature. Instead, the SecurID server administrator must resynchronize the RSA token for the user.

Related Topics

- Configuring the SecurID Server on page 327

- Configuring for External Authentication Servers on page 322

Configuring the SecurID Server

You can put several user accounts together to form a user group, which you can store on the local database or on a RADIUS, LDAP, or SecurID server. When you reference an authentication user group and an external authentication server in a policy, the traffic matching the policy provokes an authentication check.

Before You Begin
For background information, read "Understanding SecurID User Authentication" on page 326.

This example uses SecurID as the external authentication server.

The topic covers:

- Configuring SecurID as the External Authentication Server on page 327
- CLI Configuration on page 327
- Deleting the Node Secret File on page 328
- Related Topics on page 329

Configuring SecurID as the External Authentication Server

For SecurID, you configure information about the Juniper Networks device on the SecurID server and this information is exported to a file called sdconf.rec.

To install the sdconf.rec file on the device, you must use an out-of-band method such as FTP. Install the file in a directory whose files are not deleted regularly. Do not put it in a temporary directory. For example, you might install it in /var/db/secureid/server1/sdconf.rec.

The sdconf.rec file contains information that provides the Juniper Networks device with the address of the SecurID server. You do not need to configure this information explicitly when you configure the SecurID server to be used as the external authentication server.

CLI Configuration

1. For this example, specify that server1 is to be used as the SecurID server and that the configuration file for it resides on the device in the /var/db/securid/server1/sdconf.rec file.

   user@host> **set access securid-server name server1 config-file "**
   /var/db/securid/server1/sdconf.rec"

2. For prof_2 profile, configure SecurID as the server to be used for external authentication. This restricts firewall users to authenticate through the SecurID

server only. If the SecurID server authentication fails, the firewall user is locked out.

```
user@host# set access profile prof_2 authentication-order [securid]
```

To share a single SecurID server across multiple profiles, for each profile set the authentication-order parameter to include securid as the authentication mode.

3. Configure firewall user (ClientsA-E) and assign firewall users (ClientA and ClientB) to client groups alpha, beta, and gamma.

```
user@host# set access profile prof_2 client clientA client-group alpha
user@host# set access profile prof_2 client clientA client-group beta
user@host# set access profile prof_2 client clientA client-group gamma
user@host# set access profile prof_2 client clientA firewall-user password pwd1
user@host# set access profile prof_2 client clientB client-group alpha
user@host# set access profile prof_2 client clientB client-group beta
user@host# set access profile prof_2 client clientB firewall-user password pwd3
user@host# set access profile prof_2 client clientC firewall-user password pwd4
user@host# set access profile prof_2 client clientD firewall-user password pwd5
user@host# set access profile prof_2 client clientE firewall-user password pwd2
```

4. Configure client groups in the session options.

```
user@host# set access profile prof_2 session-options client-group u1
user@host# set access profile prof_2 session-options client-group alpha
user@host# set access profile prof_2 session-options client-group gamma
user@host# set access profile prof_2 session-options client-idle-timeout 255
user@host# set access profile prof_2 session-options client-session-timeout 4
```

Deleting the Node Secret File

When the Juniper Networks device initially communicates successfully with the SecurID server, a node secret file is created for it automatically. The file is created as a result of the handshake between the Juniper Networks device and the SecurID server after the software authenticates the first user successfully. All subsequent communication between the Juniper Networks device and the SecurID server relies on this secret as a representation of trust between the two nodes instead of repeating the handshake with each authentication request.

Under normal circumstances you should not delete the node secret file. In the rare case that you must do so, for example, to debug a serious problem, you can use the clear command to remove the file.

WARNING: If you delete the file, you must deselect a box on the SecurID server to indicate that the node secret file for the Juniper Networks device and the SecurID server no longer exists. Otherwise authentication attempts will fail.

To delete the node secret file, enter the following command. During subsequent user authentication, the device reestablishes a shared secret with the SecurID server and re-creates the node secret file.

```
user@host> clear network-access securid-node-secret-file
```

Related Topics

- Configuring for External Authentication Servers on page 322
- Firewall User Authentication Overview on page 305

Displaying the Authentication Table

The authentication table allows you to view a list of users and IP addresses that are currently authenticated by the Juniper Networks device.

This topic covers:

- J-Web Configuration on page 329
- CLI Configuration on page 329
- Related Topics on page 329

J-Web Configuration

To display the authentication table, select **Monitor > Security > Firewall Authentication > Authentication Table** .

CLI Configuration

To display the authentication table, enter the following commands:

- show security firewall-authentication users
- show security firewall-authentication users identifier <num>
- show security firewall-authentication users address <ipaddress>

  where:

  users—Shows the list of users and IP addresses currently authenticated by JUNOS Software.

  identifier—Shows more details on a specific row.

  address—Shows more details on a specific IP address.

Related Topics

- Firewall User Authentication Overview on page 305
- Understanding Authentication Schemes on page 307

Understanding Banner Customization

A banner is a message that appears on a monitor in different places depending on the type of login:

Figure 41: Banner Customization

- At the top of a browser screen after a user has successfully logged onto a Web authentication address as shown Figure 41 on page 330.
- Before or after a Telnet, an FTP, or an HTTP login prompt, success message, and fail message for users

All of the banners, except for the one for a console login, already have default messages. You can customize the messages that appear on the banners to better suit the network environment in which you use the device.

Related Topics

- Customizing a Banner on page 330

Customizing a Banner

NOTE: In this example, you change the message that appears in the browser to indicate that a user has successfully authenticated after successfully logging in through Web authentication. The new message is " Web Authentication is successful." If the authentication fails, then the new message reads " Authentication failed."

Before You Begin
For background information, read "Understanding Banner Customization" on page 330.

To customize a banner, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 331
- CLI Configuration on page 331
- Related Topics on page 332

J-Web Configuration

To specify the banner text for failed pass-through authentication through FTP using the J-Web configuration editor:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Firewall Authentication, click **Configure** or **Edit**.
4. Next to Pass through, click **Configure** or **Edit**.
5. In the Default profile box, type Prof1.
6. Next to Ftp, click **Configure** or **Edit**.
7. Next to Banner, click **Configure** or **Edit**.
8. In the Fail box, type " Authentication failed" and click **OK**.

To specify the banner text for successful Web authentication:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Access, click **Configure** or **Edit**.
3. Next to Firewall Authentication, click **Configure** or **Edit**.
4. Next to Web authentication, click **Configure** or **Edit**.
5. In the Default profile box, type Prof1.
6. Next to Banner, click **Configure** or **Edit**.
7. In the Success box, type " Web Authentication is successful" and click OK.
8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see "Verifying Firewall User Authentication" on page 333

CLI Configuration

To customize the banner text, do the following:

1. Specify the banner text for failed pass-through authentication through FTP.

```
user@host# set access firewall-authentication pass-through default-profile prof_1
user@host# set access firewall-authentication pass-through ftp banner fail "
   Authentication failed"
```

2. Specify the banner text successful Web authentication.

```
user@host# set access web-authentication default-profile prof_1
user@host# set access web-authentication banner success " Web Authentication
   is successful"
```

3. If you are finished configuring the device, commit the configuration.
4. To check the configuration, see "Verifying Firewall User Authentication" on page 333.

Related Topics

- Understanding Banner Customization on page 330

Configuring Firewall Authentication—Quick Configuration

You can use J-Web Quick Configuration to quickly configure banner pages for firewall authentication.

Before You Begin
For background information, read ■ Firewall User Authentication Overview on page 305 ■ Understanding Banner Customization on page 330 ■ Customizing a Banner on page 330

To configure banners for firewall authentication with Quick Configuration, follow these steps:

1. In the J-Web user interface, select **Configure > Authentication > Firewall Authentication**.
2. Enter information into the Quick Configuration page for Firewall Authentication banners as described in Table 51 on page 332.
3. From the Quick Configuration page, click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel.**

Table 51: Firewall Authentication

Field	Function	Action
Pass-Through Firewall Authentication Settings		
Default Profile	Policies use this profile to authenticate users. The profile may contain authentication servers or users.	Select the authentication profile to use if a profile is not specified in a policy.

Table 51: Firewall Authentication *(continued)*

Field	Function	Action
HTTP Banner	Customize banner text for users logging in using HTTP: Login: This message displays the login prompt. Fail: This message displays failed login. Success: This message displays a successful login.	Enter a string up to 250 characters in length.
FTP Banner	Customize banner text for users logging in using FTP: Login: This message displays the login prompt. Fail: This message displays failed login. Success: This message displays a successful login.	Enter a string up to 250 characters in length.
Telnet Banner	Customize banner text for users logging in using Telnet: Login: This message displays the login prompt. Fail: This message displays failed login. Success: This message displays a successful login.	Enter a string up to 250 characters in length.
Web Authentication Settings		
Default Profile	Policies use this profile to authenticate users. The profile may contain authentication servers or users.	Select the authentication profile to use if a profile is not specified in a policy.
Banner (Success)	Customize banner text for users logging in through the Web browser. Success: This message displays a successful login.	Enter a string up to 250 characters in length.

Verifying Firewall User Authentication

Purpose Displays firewall authentication user history.

Action Use the show security firewall-authentication CLI command to display information on authenticated firewall users. For more information, see the *JUNOS Software CLI Reference*.

```
user@host# show security firewall-authentication history
History of firewall authentication data:
Authentications: 2
Id Source Ip Date Time Duration Status User
```

```
1 99.99.99.1 2007-10-12 21:24:02 0:00:24 Failed troy
2 99.99.99.1 2007-10-12 21:24:48 0:00:22 Success voyager
user@host>  show security firewall-authentication history identifier 1
Username: troy
Source IP: 99.99.99.1
Authentication state: Failed
Authentication method: Pass-through using Telnet
Access start date: 2007-10-12
Access start time: 21:24:02
Duration of user access: 0:00:24
Policy name: lnx2-telnet-lnx1
Source zone: dl2
Destination zone: dl1
Access profile: wonder
Bytes sent by this user: 0
Bytes received by this user: 2660
Client-groups: Sunnyvale Bangalore
user@host>  show security firewall-authentication users
Firewall authentication data:
Total users in table: 1
Id Source Ip Src zone Dst zone Profile Age Status User
3 99.99.99.1     dl2   dl1   wonder   1 Failed TechPubs
user@host>  show security firewall-authentication users identifier 3
Username: TechPubs
Source IP: 99.99.99.1
Authentication state: Failed
Authentication method: Pass-through using Telnet
Age: 1
Access time remaining: 9
Source zone: dl2
Destination zone: dl1
Policy name: lnx2-telnet-lnx1
Access profile: wonder
Interface Name: ge-0/0/1.0
Bytes sent by this user: 0
Bytes received by this user: 1521
```

What it Means The output displays information about firewall users authenticating to the network. Verify the following information:

- Number of firewall users who successfully authenticated and firewall users who failed to log in.
- Details on each firewall user trying to authenticate.

Chapter 10

Infranet Authentication

This chapter includes the following information on how you can use SRX Series or J Series devices as JUNOS Enforcers in a Unified Access Control (UAC) deployment:

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Unified Access Control Overview on page 335
- Enabling Unified Access Control on SRX Series and J Series Devices on page 340
- Configuring the SRX Series and J Series Devices as a JUNOS Enforcer on page 342
- Configuring the JUNOS Enforcer Failover Options on page 350
- Enabling the JUNOS Enforcer Test-Only Mode on page 351

Unified Access Control Overview

A Unified Access Control (UAC) deployment uses the following components to secure a network and ensure that only qualified end users can access protected resources:

- **Infranet Controllers**—An Infranet Controller is a policy decision point in the network. It uses authentication information and policy rules to determine whether or not to provide access to individual resources on the network. You can deploy one or more Infranet Controllers in your network.
- **Infranet Enforcers**—An Infranet Enforcer is a policy enforcement point in the network. It receives policies from the Infranet Controller and uses the rules defined in those policies to determine whether or not to allow an endpoint access to a resource. You deploy the Infranet Enforcers in front of the servers and resources that you want to protect.
- **Infranet Agents**—An Infranet Agent is a client-side component that runs directly on network endpoints (such as users' computers). The agent checks that the endpoint complies to the security criteria specified in Host Checker policies and relays that compliance information to the Infranet Enforcer. The Infranet Enforcer then allows or denies the endpoint access based on the compliance results.

An SRX Series or J Series device can act as an Infranet Enforcer in a UAC network. Specifically, it acts as a Layer 3 enforcement point, controlling access by using IP-based policies pushed down from the Infranet Controller. When deployed in a UAC network, an SRX Series or J Series device is called a *JUNOS Enforcer*.

Figure 42: Integrating a JUNOS Security Device into a Unified Access Control Network

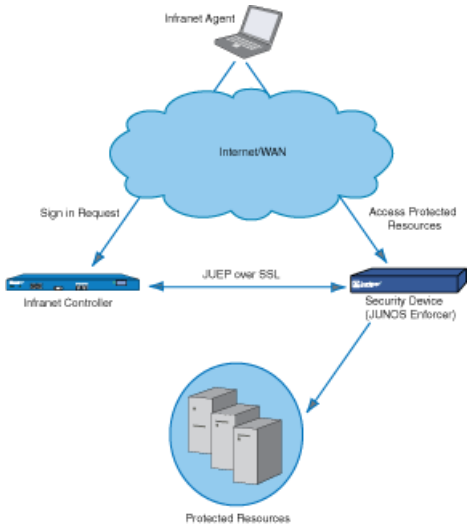

This topic includes the following information about deploying an SRX Series or J Series device as a JUNOS Enforcer in a UAC network:

- Communications Between the JUNOS Enforcer and the Infranet Controller on page 336
- JUNOS Enforcer Policy Enforcement on page 337
- Communications Between the JUNOS Enforcer and a Cluster of Infranet Controllers on page 338
- Communications Between the JUNOS Enforcer and the Infranet Agent on page 339
- Related Topics on page 339

Communications Between the JUNOS Enforcer and the Infranet Controller

When you configure an SRX Series or J Series device to connect to an Infranet Controller through the JUNOS CLI, the SRX Series or J Series device and the Infranet Controller establish secure communications as follows:

1. The Infranet Controller presents its server certificate to the SRX Series or J Series device. If configured to do so, the SRX Series or J Series device verifies the certificate. (Server certificate verification is not required; however, as an extra security measure you can verify the certificate to implement an additional layer of trust.)
2. The SRX Series or J Series device and the Infranet Controller perform mutual authentication using the proprietary challenge-response authentication. For security reasons, the password is not included in the message sent to the Infranet Controller.
3. After successfully authenticating the SRX Series or J Series device, the Infranet Controller sends it user authentication and resource access policy information.

The SRX Series and J Series devices uses this information to act as the JUNOS Enforcer in the UAC network.

4. Thereafter, the Infranet Controller and the JUNOS Enforcer can communicate freely with one another over the SSL connection. The communications are controlled by a proprietary protocol called *JUNOS UAC Enforcer Protocol (JUEP)*.

JUNOS Enforcer Policy Enforcement

Once the SRX Series or J Series device has successfully established itself as the JUNOS Enforcer, it secures traffic as follows:

1. First, the JUNOS Enforcer uses the appropriate JUNOS security policy to process the traffic. A *security policy* uses criteria such as the traffic's source IP address or the time of day that the traffic was received to determine whether or not the traffic should be allowed to pass. (For more information about security policies, see "Security Policies" on page 67.)

2. Once it determines that the traffic may pass based on the JUNOS security policy, the JUNOS Enforcer maps the traffic flow to an authentication table entry. The JUNOS Enforcer uses the source IP address of the first packet in the flow to create the mapping.

   An *authentication table entry* contains the source IP address and user role(s) of a user who has already successfully established a UAC session. A *user role* identifies a group of users based on criteria such as type (for instance, "Engineering" or "Marketing") or status (for instance, "Antivirus Running"). The JUNOS Enforcer determines whether to allow or deny the traffic to pass based on the authentication results stored in the appropriate authentication table entry. (For more information about authentication tables and user roles, see the *Unified Access Control Administration Guide*.)

   The Infranet Controller pushes authentication table entries to the JUNOS Enforcer when the devices first connect to one another (as explained in "Communications Between the JUNOS Enforcer and the Infranet Controller" on page 336) and as necessary throughout the session. For example, the Infranet Controller might push updated authentication table entries to the JUNOS Enforcer when the user's computer becomes noncompliant with endpoint security policies, when you change the configuration of a user's role, or when you disable all user accounts on the Infranet Controller in response to a security problem such as a virus on the network.

   If the JUNOS Enforcer drops a packet due to a missing authentication table entry, the device sends a message to the Infranet Controller, which in turn may provision a new authentication table entry and send it to the JUNOS Enforcer. This process is called *dynamic authentication table provisioning*.

   To display a summary of the authentication table entries configured from the Infranet Controller, use the show services unified-access-control authentication-table command.

3. Once it determines that the traffic may pass based on the authentication table entries, the JUNOS Enforcer maps the flow to a resource. The JUNOS Enforcer uses the destination IP address specified in the flow to create the mapping. Then

the device uses that resource as well as the user role specified in the authentication table entry to map the flow to a resource access policy.

A *resource access policy* specifies a particular resource to which you want to control access based on user role. For instance, you might create a resource access policy that only allows users who are members of the "Engineering" and "Antivirus Running" user roles access to the "Engineering-Only" server. Or you might create a resource access policy that allows members of the "No Antivirus Running" user role access to the "Remediation" server where antivirus software is available for download. (For more information about resource access policies, see the *Unified Access Control Administration Guide*.)

The Infranet Controller pushes resource access policies to the JUNOS Enforcer when the devices first connect to one another (as explained in "Communications Between the JUNOS Enforcer and the Infranet Controller" on page 336) and when you modify your resource access policy configurations on the Infranet Controller.

If the JUNOS Enforcer drops the packet due to a "deny" policy, the JUNOS Enforcer sends a message to the Infranet Controller, which in turn sends a message to the endpoint's Odyssey Access Client (if available). (The Infranet Controller does not send "deny" messages to the "agentless" client.)

To display a summary of UAC resource access policies configured from the Infranet Controller, use the show services unified-access-control policies command.

4. Once it determines that the traffic may pass based on the resource access policies, the JUNOS Enforcer processes the traffic using the remaining application services defined in the JUNOS policy. The JUNOS Enforcer runs the remaining services in the following order: Intrusion Detection and Prevention (IDP), URL filtering, and Application Layer Gateways (ALGs).

Communications Between the JUNOS Enforcer and a Cluster of Infranet Controllers

You can configure a JUNOS Enforcer to work with more than one Infranet Controller in a high availability configuration known as an Infranet Controller cluster. The JUNOS Enforcer communicates with only one Infranet Controller at a time; the other Infranet Controllers are used for failover. If the JUNOS Enforcer cannot connect to the first Infranet Controller you added to a cluster, it tries to connect to the failed Infranet Controller again. Then it fails over to the other Infranet Controllers in the cluster. It continues trying to connect to Infranet Controllers in the cluster until a connection occurs.

When the JUNOS Enforcer cannot establish a connection to an Infranet Enforcer, it preserves all its existing authentication table entries and Unified Access Control (UAC) policies and takes the timeout action that you specify. Timeout actions include:

- close—Close existing sessions and block any further traffic. This is the default option.
- no-change—Preserve existing sessions and require authentication for new sessions.
- open—Preserve existing sessions and allow new sessions access.

Once the JUNOS Enforcer is able to reestablish a connection to an Infranet Controller, the Infranet Controller compares the authentication table entries and UAC policies stored on the JUNOS Enforcer with the authentication table entries and policies stored on the Infranet Controller and reconciles the two as required.

The Infranet Controllers configured on a JUNOS Enforcer should all be members of the same Infranet Controller cluster.

Communications Between the JUNOS Enforcer and the Infranet Agent

An Infranet Agent helps you secure traffic on your network starting with the endpoints that initiate communications as follows:

1. The Infranet Agent, which runs directly on the endpoint, checks that the endpoint is compliant with your Unified Access Control (UAC) Host Checker policies. You can use a wide variety of criteria within a UAC Host Checker policy to determine compliance. For example, you can configure the Host Checker policy to confirm that the endpoint is running antivirus software or a firewall or that the endpoint is not running specific types of malware or processes.
2. The Infranet Agent transmits the compliance information to the JUNOS Enforcer.
3. The JUNOS Enforcer allows or denies the endpoint access to the resources on your network based on the Host Checker compliance results.

Because the Infranet Agent runs directly on the endpoint, you can use the Infranet Agent to check the endpoint for security compliance at any time. For instance, when a user tries to sign into the Infranet Controller, you can require the Infranet Agent to send compliance results immediately—the user will not even see the sign-in page until the Infranet Agent returns positive compliance results to the Infranet Controller. You can also configure the Infranet Agent to check for compliance after the user signs in or periodically during the user session. For more information about the Infranet Agent, see the *Unified Access Control Administration Guide*.

To integrate the Infranet Agent into a JUNOS-UAC deployment, no special configuration is required on the JUNOS Enforcer. You simply need to create security policies enabling access to the appropriate endpoints as you would for any other JUNOS-UAC deployment. (For more information, see “JUNOS Enforcer Policy Enforcement” on page 337.) If the endpoints running the Infranet Agent have appropriate access, they will automatically send their compliance results to the Infranet Controller and the Infranet Controller will update the authentication table entries accordingly and push them to the JUNOS Enforcer. The JUNOS Enforcer supports connections with the Odyssey Access Client and “agentless” Infranet Agents.

Related Topics

- Enabling Unified Access Control on SRX Series and J Series Devices on page 340
- Configuring the SRX Series and J Series Devices as a JUNOS Enforcer on page 342
- Configuring the JUNOS Enforcer Failover Options on page 350
- Enabling the JUNOS Enforcer Test-Only Mode on page 351

Enabling Unified Access Control on SRX Series and J Series Devices

JUNOS security policies enforce rules for transit traffic, defining what traffic can pass through the Juniper Networks device. The policies control traffic that enters from one zone (from-zone) and exits another (to-zone). To enable an SRX Series or J Series device as a JUNOS Enforcer in a Unified Access Control (UAC) deployment, you must identify the following:

- The source and destination zones through which UAC traffic will travel. It also needs the list of interfaces including which zones they are in. The infranet controller uses the destination zone to match its own IPsec routing policies configured on infranet controller.
- Define JUNOS security policies that encompass those zones, and enable UAC for those policies.
- Include the "identity" which is configured under Security IKE Gateway. This is a string such as "gateway1.juniper.net". where gateway1.juniper.net distinguishes between IKE gateways. (This is in order for JUNOS to know for which tunnel traffic is intended.)
- Include the preshared seed. This is used to generate the pre-shared key from the full identity of the remote user for phase 1 credentials.
- Include the RADIUS shared secret. This allows the infranet controller to accept RADIUS packets for XAuth authentication (extended authentication) from the JUNOS Infranet Enforcer.

Before You Begin

1. For background information, read "Unified Access Control Overview" on page 335.
2. Set up the interfaces through which UAC traffic should enter the SRX Series or J Series device. See "Understanding Security Zone Interfaces" on page 61 for background information and "Configuring Interfaces—Quick Configuration" on page 62 for instructions.
3. Group interfaces with identical security requirements into zones. See "Understanding Security Zones" on page 47 for background information and "Creating Security Zones" on page 49 for instructions.
4. Create security policies to control the traffic that passes through the security zones. See "Security Policies Overview" on page 69 for background information and "Configuring Policies" on page 80 for instructions.

This topic covers:

- Support and Limitations on page 341
- CLI Configuration on page 341
- Related Topics on page 342

Support and Limitations

When configuring IPsec between the infranet controller, the Odyssey Access Client, and the SRX or J Series device, you should note that the following are IKE proposal (or phase1 proposal) methods or protocol configurations that are supported from the infranet controller to the Odyssey Access Client.

- IKE proposal: authentication-method pre-shared-keys (you must use pre-shared-keys)
- IKE policy: mode aggressive (you must use aggressive mode)
  - pre-shared-key ascii-text < key > (only ascii text pre-shared-keys are supported)
- IKE gateway: dynamic
  - hostname < identity > (you must specify a unique identity among gateways)
  - ike-user-type group-ike-id (you must use group-ike-id)
  - xauth access-profile < profile > (you must specify xauth)

The following are IPsec proposal (or phase2 proposal) methods or protocol configurations that are supported from the infranet controller to the Odyssey Access Client.

- IPsec proposal: protocol esp (you must specify esp)
- IPsec VPN: establish-tunnels immediately (you must establish-tunnels immediately)

NOTE: Refer to the next chapter, "Internet Protocol Security (IPsec)" on page 355, for details on configuring VPN.

Only one IPsec VPN tunnel is supported per from-zone to to-zone security policy. This is a limitation on the Infranet Controller such that only one IPsec VPN tunnel per from-zone to to-zone security policy is supported. On JUNOS security policies, you can define multiple policies differentiated by different source-address and/or destination-address. However, the Infranet Controller cannot differentiate such configurations and could potentially identify the incorrect IKE gateway.

CLI Configuration

To configure enable UAC through a JUNOS security policy, use the application-services statement:

edit security policies from-zone *zone-name* to-zone *zone-name* policy match then permit application-services uac-policy

Related Topics

- Configuring the SRX Series and J Series Devices as a JUNOS Enforcer on page 342
- Configuring the SRX Series Device as a JUNOS Enforcer using IPsec — Example on page 343
- Configuring the JUNOS Enforcer Failover Options on page 350
- Enabling the JUNOS Enforcer Test-Only Mode on page 351

Configuring the SRX Series and J Series Devices as a JUNOS Enforcer

To configure a SRX Series or J Series devices to act as a JUNOS Enforcer in a Unified Access Control (UAC) deployment, you must specify an Infranet Controller to which the SRX Series or J Series device should connect.

Before You Begin

1. For background information, read "Unified Access Control Overview" on page 335.
2. Enable UAC through the relevant JUNOS security policies. See "Enabling Unified Access Control on SRX Series and J Series Devices" on page 340
3. (Optional) Import the Infranet Controller's server certificate onto the SRX Series or J Series device and create a profile for the certificate authority (CA) that signed the certificate. See "Using Digital Certificates" on page 410 for background information and "Loading CA and Local Certificates Manually" on page 420 for instructions.
4. Configure user authentication and authorization by setting up user roles, authentication and authorization servers, and authentication realms on the Infranet Controller. For more information, see the *Unified Access Control Administration Guide*.
5. Configure resource access policies on the Infranet Controller to specify which endpoints are allowed or denied access to protected resources. For more information, see the *Unified Access Control Administration Guide*.

This topic covers:

- CLI Configuration on page 342
- Configuring the SRX Series Device as a JUNOS Enforcer using IPsec — Example on page 343
- Related Topics on page 350

CLI Configuration

To configure an SRX Series or J Series device to act as a JUNOS Enforcer:

1. Specify the Infranet Controller(s) to which the SRX Series or J Series device should connect:
   - Specify the Infranet Controller's host name:

edit services unified-access-control infranet-controller *hostname*

- Specify the Infranet Controller's IP address:

  edit services unified-access-control infranet-controller *hostname* address *ip-address*

NOTE: When configuring access to multiple Infranet Controllers, you must define each separately. For example:

```
edit services unified-access-control infranet-controller IC1
edit services unified-access-control infranet-controller IC2
edit services unified-access-control infranet-controller IC3

edit services unified-access-control infranet-controller IC1 address 10.10.10.01
edit services unified-access-control infranet-controller IC2 address 10.10.10.02
edit services unified-access-control infranet-controller IC3 address 10.10.10.03
```

Make sure that all of the Infranet Controllers are members of the same cluster.

---

2. Specify the Infranet Controller port to which the SRX Series or J Series device should connect:

   edit services unified-access-control infranet-controller *hostname* port *port*

3. Specify the JUNOS interface to which the Infranet Controller should connect:

   edit services unified-access-control infranet-controller *hostname* interface *interface-name*

4. Specify the password that the SRX Series or J Series device should use to initiate secure communications with the Infranet Controller:

   edit services unified-access-control infranet-controller *hostname* password *password*

5. (Optional) Specify information about the certificate that the device should use for SSL communications with the Infranet Controller:

   - Specify the certificate that the device should use:

     edit services unified-access-control infranet-controller *hostname* server-certificate-subject *certificate-name*

   - Specify the certificate authority (CA) profile associated with the certificate:

     edit services unified-access-control infranet-controller *hostname* ca-profile *ca-profile*

Configuring the SRX Series Device as a JUNOS Enforcer using IPsec — Example

1. Set system and syslog information.

```
system {
host-name test_host;
   domain-name test.juniper.net;
   host-name test_host;
   root-authentication {
   encrypted-password "$1$uhqXoD0T$6h26f0xXExOqkPHQLvaTF0";
   }
   services {
   ftp;
      ssh;
      telnet;
      web-management {
      http {
         interface ge-0/0/0.0;
      }
   }
}
syslog {
user * {
   any emergency;
      }
      file messages {
      any critical;
      authorization info;
      }
      file interactive-commands {
      interactive-commands error;
   }
   }
   max-configurations-on-flash 5;
   max-configuration-rollbacks 5;
   license {
   autoupdate {
      url https://ae1.juniper.net/junos/key_retrieval;
   }
   }
   ntp {
   boot-server 1.2.3.4;
   server 1.2.3.4;
}
}
```

2. Configure the interfaces.

```
interfaces {
ge-0/0/0 {
   unit 0 {
      family inet {
         address 10.64.75.135/16;
      }
   }
}
ge-0/0/1 {
   unit 0 {
      family inet {
         address 10.100.54.1/16;
```

```
        }
      }
    }
    ge-0/0/2 {
      unit 0 {
        family inet {
          10.101.54.1/16;
        }
      }
    }
```

3. Configure routing options.

```
routing-options {
static {
  route 0.0.0.0/0 next-hop 10.64.0.1;
    route 10.11.0.0/16 next-hop 10.64.0.1;
    route 172.0.0.0/8 next-hop 10.64.0.1;
  route 10.64.0.0/16 next-hop 10.64.0.1;
}
}
```

4. Configure security options.

```
security {
ike {
  traceoptions {
    file ike;
    flag all;
  }
  proposal prop1 {
  authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
  encryption-algorithm 3des-cbc;
  }
  policy pol1 {
  mode aggressive;
    proposals prop1;
  pre-shared-key ascii-text "$9$YS4ZjmPQ6CuTz6Au0cSvWLxNbiHm";
  }
  gateway gateway1 {
  ike-policy pol1;
    dynamic {
    hostname gateway1.juniper.net;
      connections-limit 1000;
    ike-user-type group-ike-id;
    }
    external-interface ge-0/0/0;
  xauth access-profile infranet;
  }
  gateway gateway2 {
  ike-policy pol1;
    dynamic {
    hostname gateway2.juniper.net;
      connections-limit 1000;
```

```
        ike-user-type group-ike-id;
        }
        external-interface ge-0/0/0;
    xauth access-profile infranet;
}
}
```

5. Configure IPsec parameters.

```
ipsec {
proposal prop1 {
    protocol esp;
        authentication-algorithm hmac-sha1-96;
        encryption-algorithm 3des-cbc;
    lifetime-seconds 86400;
    }
    policy pol1 {
    proposals prop1;
    }
    vpn vpn1 {
    ike {
        gateway gateway1;
        ipsec-policy pol1;
        }
    establish-tunnels immediately;
    }
    vpn vpn2 {
    ike {
        gateway gateway2;
        ipsec-policy pol1;
        }
    establish-tunnels immediately;
}
}
```

6. Configure screen options.

```
screen {
ids-option untrust-screen {
    icmp {
        ping-death;
        }
        ip {
        source-route-option;
        tear-drop;
        }
        tcp {
        syn-flood {
            alarm-threshold 1024;
                attack-threshold 200;
                source-threshold 1024;
                destination-threshold 2048;
                queue-size 2000;
            timeout 20;
            }
        land;
```

```
    }
  }
  }
```

7. Configure zones.

```
zones {
security-zone trust {
   tcp-rst;
      host-inbound-traffic {
      system-services {
         all;
         }
         protocols {
         all;
      }
      }
      interfaces {
      ge-0/0/0.0;
   }
   }
   security-zone untrust {
   host-inbound-traffic {
      system-services {
         all;
         }
         protocols {
         all;
      }
      }
      interfaces {
      ge-0/0/1.0;
   }
   }
   security-zone zone101 {
   host-inbound-traffic {
      system-services {
         all;
         }
         protocols {
         all;
      }
      }
      interfaces {
      ge-0/0/2.0;
   }
}
}
```

8. Configure policies for UAC.

```
policies {
inactive: from-zone trust to-zone trust {
   policy default-permit {
      match {
         source-address any;
```

```
            destination-address any;
        application any;
        }
        then {
        permit;
    }
}
}
from-zone trust to-zone untrust {
inactive: policy default-permit {
    match {
        source-address any;
            destination-address any;
        application any;
        }
        then {
        permit;
    }
    }
    inactive: policy default-deny {
    match {
        source-address any;
            destination-address any;
        application any;
        }
        then {
        permit;
    }
    }
    policy pol1 {
    match {
        source-address any;
            destination-address any;
        application any;
        }
        then {
        permit {
            tunnel {
                ipsec-vpn vpn1;
                }
                application-services {
                uac-policy;
            }
            }
            log {
            session-init;
            session-close;
        }
    }
}
}
from-zone untrust to-zone trust {
policy pol1 {
    match {
        source-address any;
            destination-address any;
```

```
        application any;
        }
        then {
        permit;
          log {
          session-init;
          session-close;
        }
      }
    }
    }
    from-zone trust to-zone zone101 {
    policy pol1 {
      match {
        source-address any;
          destination-address any;
        application any;
        }
        then {
        permit {
          tunnel {
            ipsec-vpn vpn2;
            }
            application-services {
            uac-policy;
          }
          }
          log {
          session-init;
          session-close;
        }
      }
      }
      policy test {
      match {
        source-address any;
          destination-address any;
        application any;
        }
        then {
        permit;
      }
    }
    }
    default-policy {
    deny-all;
}
}
}
```

9. Configure access using radius server authentication.

```
access {
profile infranet {
  authentication-order radius;
    radius-server {
```

```
        10.64.160.120 secret "$9$KBoWX-YgJHqfVwqfTzCAvWL";
    }
}
}
```

10. Configure services for UAC.

```
services {
unified-access-control {
    inactive: infranet-controller IC27 {
        address 3.23.1.2;
            interface ge-0/0/0.0;
        password "$9$Wjl8X-Vb2GDkev4aGUHkuOB";
        }
        infranet-controller prabaIC {
        address 10.64.160.120;
            interface ge-0/0/0.0;
        password "$9$jdkmT69pRhrz3hrev7Nik.";
        }
        traceoptions {
        flag all;
    }
}
}
```

Related Topics

- Configuring the JUNOS Enforcer Failover Options on page 350
- Enabling the JUNOS Enforcer Test-Only Mode on page 351

Configuring the JUNOS Enforcer Failover Options

To configure Infranet Controller failover processing, you must configure the JUNOS Enforcer to connect to a cluster of Infranet Controllers. The JUNOS Enforcer communicates with one of these Infranet Controllers at a time and uses the others for failover processing.

Before You Begin

1. For background information, read "Unified Access Control Overview" on page 335.
2. Enable UAC through the relevant JUNOS security policies. See "Enabling Unified Access Control on SRX Series and J Series Devices" on page 340
3. Configure the SRX Series or J Series device as a JUNOS Enforcer. During the configuration, define a cluster of Infranet Controllers to which the JUNOS Enforcer should connect. See "Configuring the SRX Series and J Series Devices as a JUNOS Enforcer" on page 342.

This topic covers:

- CLI Configuration on page 351
- Related Topics on page 351

CLI Configuration

To configure failover processing:

1. Specify how often (in seconds) the JUNOS Enforcer should expect a heartbeat signal from the Infranet Controller indicating an active connection:

   edit services unified-access-control interval *seconds*

2. Specify the interval (in seconds) at which the JUNOS Enforcer should consider the current connection timed out:

   edit services unified-access-control timeout *seconds*

3. Specify how the JUNOS Enforcer should handle all current and subsequent traffic sessions when its connection to an Infranet Controller cluster times out:

   edit services unified-access-control timeout-action (close | no-change | open)

Related Topics

- Enabling the JUNOS Enforcer Test-Only Mode on page 351

Enabling the JUNOS Enforcer Test-Only Mode

When configured in test-only mode, the SRX Series or J Series device enables all Unified Access Control (UAC) traffic to go through regardless of the UAC policy settings. The device logs the UAC policy's access decisions without enforcing them so you can test the implementation without impeding traffic.

Before You Begin

1. For background information, read "Unified Access Control Overview" on page 335.
2. Enable UAC through the relevant JUNOS security policies. See "Enabling Unified Access Control on SRX Series and J Series Devices" on page 340
3. Configure the SRX Series and J Series devices as a JUNOS Enforcer. See "Configuring the SRX Series and J Series Devices as a JUNOS Enforcer" on page 342.
4. If you are connecting to a cluster of Infranet Controllers, enable failover options. See "Configuring the JUNOS Enforcer Failover Options" on page 350.

To activate or deactivate test-only mode, use the test-only-mode configuration statement:

edit services unified-access-control test-only-mode (true | false)

Part 6
Virtual Private Networks

- Internet Protocol Security (IPsec) on page 355
- Public Key Cryptography for Certificates on page 401
- Dynamic Virtual Private Networks (VPNs) on page 433
- NetScreen-Remote VPN Client on page 477

Chapter 11

Internet Protocol Security (IPsec)

IP Security (IPsec) is a suite of related protocols for cryptographically securing communications at the IP Packet Layer. IPsec also provides methods for the manual and automatic negotiation of security associations (SAs) and key distribution, all the attributes for which are gathered in a Domain of Interpretation (DOI). The IPsec DOI is a document containing definitions for all the security parameters required for the successful negotiation of a VPN tunnel—essentially, all the attributes required for SA and IKE negotiations. See *RFC 2407* and *RFC 2408* for more information.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369
- Configuring VPN Global Settings (Standard VPNs) on page 370
- Configuring IPSec VPN Global Settings in J-Web (Standard VPNs) on page 372
- Configuring an IKE Phase 1 Proposal (Standard and Dynamic VPNs) on page 373
- Configuring Auto Tunnel in J-Web (Standard VPN) on page 375
- Configuring an IKE Policy (Standard and Dynamic VPNs) on page 384
- Configuring an IKE Gateway (Standard and Dynamic VPNs) on page 385
- Configuring an IPsec Phase 2 Proposal (Standard and Dynamic VPNs) on page 387
- Configuring an IPsec Policy (Standard and Dynamic VPNs) on page 388
- Configuring IPsec AutoKey (Standard and Dynamic VPNs) on page 390
- Configuring the Manual Tunnel using J-Web on page 391
- Configuring Hub-and-Spoke VPNs on page 393

Virtual Private Networks (VPNs) Overview

A virtual private network (VPN) provides a means for securely communicating among remote computers across a public wide area network (WAN), such as the Internet.

A VPN connection can link two local area networks (LANs) (site-to-site VPN) or a remote dial-up user and a LAN. The traffic that flows between these two points passes

through shared resources such as routers, switches, and other network equipment that make up the public WAN. To secure VPN communication while passing through the WAN, the two participants create an IP Security (IPsec) tunnel.

NOTE: The term *tunnel* does not denote tunnel mode (see "Packet Processing in Tunnel Mode" on page 364). Instead, it refers to the IPsec connection.

This topic covers:

- Security Associations (SAs) on page 356
- IPsec Key Management on page 357
- IPsec Security Protocols on page 359
- IPsec Tunnel Negotiation on page 360
- Related Topics on page 364

Security Associations (SAs)

A security association (SA) is a unidirectional agreement between the VPN participants regarding the methods and parameters to use in securing a communication channel. Through the SA, an IPsec tunnel can provide the following security functions:

- Privacy (through encryption)
- Content integrity (through data authentication)
- Sender authentication and—if using certificates—nonrepudiation (through data origin authentication)

The security functions you employ depend on your needs. If you only need to authenticate the IP packet source and content integrity, you can authenticate the packet without applying any encryption. On the other hand, if you are only concerned with preserving privacy, you can encrypt the packet without applying any authentication mechanisms. Optionally, you can both encrypt and authenticate the packet. Most network security designers choose to encrypt, authenticate, and replay-protect their VPN traffic.

An IPsec tunnel consists of a pair of unidirectional SAs—one at each end of the tunnel—that specify the security parameter index (SPI), destination IP address, and security protocol (Authentication Header or Encapsulating Security Payload) employed. An SA groups together the following components for securing communications:

- Security algorithms and keys
- Protocol mode, either transport or tunnel. JUNOS devices always use tunnel mode. (For more information, see "Packet Processing in Tunnel Mode" on page 364.)
- Key-management method, either manual key or AutoKey IKE. (For more information, see "IPsec Key Management" on page 357.)
- SA lifetime

For inbound traffic, JUNOS Software looks up the SA by using the following triplet:

- Destination IP address
- Security protocol, either AH or ESP (see "IPsec Security Protocols" on page 359)
- Security parameter index (SPI) value

For outbound VPN traffic, the policy invokes the SA associated with the VPN tunnel.

IPsec Key Management

The distribution and management of keys are critical to using VPNs successfully. JUNOS Software supports IPsec technology for creating VPN tunnels with three kinds of key creation mechanisms:

- Manual key
- AutoKey IKE with a preshared key or a certificate
- Diffie-Hellman key

You can choose your key creation mechanism—also called authentication method—during phase 1 and phase 2 proposal configuration. For more information, see "IPsec Tunnel Negotiation" on page 360.

NOTE: Manual key creation and AutoKey IKE with certificates are not supported with the dynamic VPN feature at this time.

Manual Key

With manual keys, administrators at both ends of a tunnel configure all the security parameters. This is a viable technique for small, static networks where the distribution, maintenance, and tracking of keys are not difficult. However, safely distributing manual-key configurations across great distances poses security issues. Aside from passing the keys face-to-face, you cannot be completely sure that the keys have not been compromised while in transit. Also, whenever you want to change the key, you are faced with the same security issues as when you initially distributed it.

AutoKey IKE

When you need to create and manage numerous tunnels, you need a method that does not require you to configure every element manually. IPsec supports the automated generation and negotiation of keys and security associations using the Internet Key Exchange (IKE) protocol. JUNOS Software refers to such automated tunnel negotiation as AutoKey IKE and supports AutoKey IKE with preshared keys and AutoKey IKE with certificates.

AutoKey IKE with Preshared Keys

Using AutoKey IKE with preshared keys to authenticate the participants in an IKE session, each side must configure and securely exchange the preshared key in

advance. In this regard, the issue of secure key distribution is the same as that with manual keys. However, once distributed, an autokey, unlike a manual key, can automatically change its keys at predetermined intervals using the IKE protocol. Frequently changing keys greatly improves security, and automatically doing so greatly reduces key-management responsibilities. However, changing keys increases traffic overhead; therefore, doing so too often can reduce data transmission efficiency.

NOTE: A preshared key is a key for both encryption and decryption, which both participants must have before initiating communication.

AutoKey IKE with Certificates

When using certificates to authenticate the participants during an AutoKey IKE negotiation, each side generates a public-private key pair and acquires a certificate. As long as the issuing certificate authority (CA) is trusted by both sides, the participants can retrieve the peer's public key and verify the peer's signature. There is no need to keep track of the keys and SAs; IKE does it automatically.

Diffie-Hellman Exchange

A Diffie-Hellman (DH) exchange allows the participants to produce a shared secret value. The strength of the technique is that it allows the participants to create the secret value over an unsecured medium without passing the secret value through the wire. There are five Diffie-Hellman groups; JUNOS Software supports groups 1, 2, and 5. The size of the prime modulus used in each group's calculation differs as follows:

- **DH Group 1**—768-bit modulus
- **DH Group 2**—1024-bit modulus
- **DH Group 5**—1536-bit modulus

NOTE: The strength of DH Group 1 security has depreciated; therefore we do not recommend its use.

The larger the modulus, the more secure the generated key is considered to be; however, the larger the modulus, the longer the key-generation process takes. Because the modulus for each DH group is a different size, the participants must agree to use the same group.

NOTE: If you configure multiple (up to four) proposals for Phase 1 negotiations, use the same Diffie-Hellman group in all proposals. The same guideline applies to multiple proposals for Phase 2 negotiations.

IPsec Security Protocols

IPsec uses two protocols to secure communications at the IP layer:

- Authentication Header (AH)—A security protocol for authenticating the source of an IP packet and verifying the integrity of its content
- Encapsulating Security Payload (ESP)—A security protocol for encrypting the entire IP packet (and authenticating its content)

You can choose your security protocols—also called authentication and encryption algorithms—during phase 1 and phase 2 proposal configuration. For more information, see "IPsec Tunnel Negotiation" on page 360.

Authentication Header (AH) Protocol

The Authentication Header (AH) protocol provides a means to verify the authenticity/integrity of the content and origin of a packet. You can authenticate the packet by the checksum calculated through a Hash Message Authentication Code (HMAC) using a secret key and either MD5 or SHA-1 hash functions.

- **Message Digest 5 (MD5)**—An algorithm that produces a 128-bit hash (also called a *digital signature* or *message digest*) from a message of arbitrary length and a 16-byte key. The resulting hash is used, like a fingerprint of the input, to verify content and source authenticity and integrity.
- **Secure Hash Algorithm-1 (SHA-1)**—An algorithm that produces a 160-bit hash from a message of arbitrary length and a 20-byte key. It is generally regarded as more secure than MD5 because of the larger hashes it produces. Because the computational processing is done in the ASIC, the performance cost is negligible.

NOTE: For more information on MD5 hashing algorithms, see *RFCs 1321* and *RFC 2403*. For more information on SHA hashing algorithms, see *RFC 2404*. For more information on HMAC, see *RFC 2104*.

Encapsulating Security Payload (ESP) Protocol

The Encapsulating Security Payload (ESP) protocol provides a means to ensure privacy (encryption) and source authentication and content integrity (authentication). ESP in tunnel mode encapsulates the entire IP packet (header and payload) and then appends a new IP header to the now-encrypted packet. This new IP header contains the destination address needed to route the protected data through the network. (For more information about tunnel mode, see "Packet Processing in Tunnel Mode" on page 364.)

With ESP, you can both encrypt and authenticate, encrypt only, or authenticate only. For encryption, you can choose one of the following encryption algorithms:

- **Data Encryption Standard (DES)**—A cryptographic block algorithm with a 56-bit key.

- **Triple DES (3DES)**—A more powerful version of DES in which the original DES algorithm is applied in three rounds, using a 168-bit key. DES provides a significant performance savings but is considered unacceptable for many classified or sensitive material transfers.
- **Advanced Encryption Standard (AES)**—An emerging encryption standard which, when adopted by Internet infrastructures worldwide, will offer greater interoperability with other devices. JUNOS Software supports AES with 128-bit, 192-bit, and 256-bit keys.

For authentication, you can use either MD5 or SHA-1 algorithms.

NOTE: Even though is possible to select NULL for encryption, it has been demonstrated that IPsec might be vulnerable to attack under such circumstances. Therefore, we suggest that you choose an encryption algorithm for maximum security.

IPsec Tunnel Negotiation

To establish an AutoKey IKE IPsec tunnel, two phases of negotiation are required:

- In Phase 1, the participants establish a secure channel in which to negotiate the IPsec security associations (SAs).
- In Phase 2, the participants negotiate the IPsec SAs for encrypting and authenticating the ensuing exchanges of user data.

For a manual key IPsec tunnel, because all of the SA parameters have been previously defined, there is no need to negotiate which SAs to use. In essence, the tunnel has already been established. When traffic matches a policy using that manual key tunnel or when a route involves the tunnel, the Juniper Networks device simply encrypts and authenticates the data, as you determined, and forwards it to the destination gateway.

Phase 1 of IKE Tunnel Negotiation

Phase 1 of an AutoKey IKE tunnel negotiation consists of the exchange of proposals for how to authenticate and secure the channel. The participants exchange proposals for acceptable security services such as:

- Encryption algorithms (DES and 3DES) and authentication algorithms (MD5 and SHA-1). (For more information, see“IPsec Security Protocols” on page 359.)
- A Diffie-Hellman group. (For more information, see“Diffie-Hellman Exchange” on page 358.)
- Preshared Key or RSA/DSA certificates. (For more information, see “IPsec Key Management” on page 357.)

A successful Phase 1 negotiation concludes when both ends of the tunnel agree to accept at least one set of the Phase 1 security parameters proposed and then process them. Juniper Networks devices support up to four proposals for Phase 1 negotiations,

allowing you to define how restrictive a range of security parameters for key negotiation you will accept.

The predefined Phase 1 proposals that JUNOS Software provides are as follows:

- **Standard**—pre-g2-aes128-sha and pre-g2-3des-sha
- **Compatible**—pre-g2-3des-sha, pre-g2-3des-md5, pre-g2-des-sha, and pre-g2-des-md5
- **Basic**—pre-g1-des-sha and pre-g1-des-md5

You can also define custom Phase 1 proposals.

NOTE: If you are using the dynamic VPN feature, note that you must create a custom Phase 1 proposal. Predefined Phase 1 proposals are not available at this time.

Phase 1 exchanges can take place in either main or aggressive mode. You can choose your mode during IKE policy configuration.

Main Modes

In main mode, the initiator and recipient send three two-way exchanges (six messages total) to accomplish the following services:

- First exchange (messages 1 and 2—Propose and accept the encryption and authentication algorithms.
- Second exchange (messages 3 and 4—Execute a Diffie-Hellman exchange, and the initiator and recipient each provide a pseudo-random number.
- Third exchange (messages 5 and 6)—Send and verify their identities.

The information transmitted in the third exchange of messages is protected by the encryption algorithm established in the first two exchanges. Thus, the participants' identities are not transmitted in the clear.

Aggressive Mode

In aggressive mode, the initiator and recipient accomplish the same objectives, but in only two exchanges, with a total of three messages:

- First message—The initiator proposes the SA, initiates a Diffie-Hellman exchange, and sends a pseudo-random number and its IKE identity.
- Second message—The recipient accepts the SA; authenticates the initiator; and sends a pseudo-random number, its IKE identity, and, if using certificates, the recipient's certificate.
- Third message—The initiator authenticates the recipient, confirms the exchange, and, if using certificates, sends the initiator's certificate.

Because the participants' identities are exchanged in the clear (in the first two messages), aggressive mode does not provide identity protection.

NOTE: When a dialup VPN user negotiates an AutoKey IKE tunnel with a preshared key, aggressive mode must be used. Therefore, you must always use aggressive mode with the dynamic VPN feature. Note also that a dialup VPN user can use an e-mail address, a fully qualified domain name (FQDN), or an IP address as its IKE ID. A dynamic peer can use either an e-mail address or FQDN, but not an IP address.

Phase 2 of IKE Tunnel Negotiation

After the participants have established a secure and authenticated channel, they proceed through Phase 2, in which they negotiate the SAs to secure the data to be transmitted through the IPsec tunnel.

Similar to the process for Phase 1, the participants exchange proposals to determine which security parameters to employ in the SA. A Phase 2 proposal also includes a security protocol—either Encapsulating Security Payload (ESP) or Authentication Header (AH)—and selected encryption and authentication algorithms. The proposal can also specify a Diffie-Hellman group, if Perfect Forward Secrecy (PFS) is desired.

Regardless of the mode used in Phase 1, Phase 2 always operates in quick mode and involves the exchange of three messages.

Juniper Networks devices support up to four proposals for Phase 2 negotiations, allowing you to define how restrictive a range of tunnel parameters you will accept. The predefined Phase 2 proposals that JUNOS Software provides are as follows:

- **Standard—**g2-esp-3des-sha and g2-esp-aes128-sha
- **Compatible**—nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, and nopfs-esp-des-md5
- **Basic—**nopfs-esp-des-sha and nopfs-esp-des-md5

You can also define custom Phase 2 proposals.

NOTE: If you are using the dynamic VPN feature, note that you must create a custom Phase 2 proposal. Predefined Phase 2 proposals are not available at this time.

Proxy IDs

In Phase 2, the peers exchange proxy IDs. A proxy ID is a three-part tuple consisting of local IP address-remote IP address-service. The proxy ID for both peers must match, which means that the service specified in the proxy ID for both peers must be the same, and the local IP address specified for one peer must be the same as the remote IP address specified for the other peer.

You can specify the format of Proxy IDs during IPsec Autokey configuration.

Perfect Forward Secrecy

Perfect Forward Secrecy (PFS) is a method for deriving Phase 2 keys independent from and unrelated to the preceding keys. Alternatively, the Phase 1 proposal creates the key (the SKEYID_d key) from which all Phase 2 keys are derived. The SKEYID_d key can generate Phase 2 keys with a minimum of CPU processing. Unfortunately, if an unauthorized party gains access to the SKEYID_d key, all your encryption keys are compromised.

PFS addresses this security risk by forcing a new Diffie-Hellman key exchange to occur for each Phase 2 tunnel. Using PFS is thus more secure, although the rekeying procedure in Phase 2 might take slightly longer with PFS enabled.

You can enable PFS during IPsec policy configuration.

Replay Protection

A replay attack occurs when somebody intercepts a series of packets and uses them later either to flood the system, causing a denial-of-service (DoS), or to gain entry to the trusted network. JUNOS Software provides a replay protection feature, which enables devices to check every IPsec packet to see if it has been received previously. If packets arrive outside a specified sequence range, JUNOS Software rejects them. Use of this feature does not require negotiation because packets are always sent with sequence numbers. You simply have the option of checking or not checking the sequence numbers.

You can enable replay protection during IPsec Autokey configuration.

Distributed VPN in SRX Series Services Gateways

In the SRX3000 and SRX5000 lines, the IKE provides tunnel management for IPsec and authenticates end entities. The IKE performs a Diffie-Hellman key exchange to generate an IPsec tunnel between network devices. The IPsec tunnels generated by IKE are used to encrypt, decrypt, and authenticate user traffic between the network devices at the IP layer.

The VPN is created by distributing the IKE and IPsec workload among the multiple Security Processing Units (SPUs) of the platform. The IKE workload is distributed based on a key generated from the IKE packet's 4 tuples (source IP address, destination IP addresses, and UDP ports). The workload is distributed by assigning anchoring SPUs logically and mapping the logical SPUs to physical SPU based on the composition at that given time. This distribution prevents any change in the number and composition of SPUs in the device, which may happen due to hot swap or SPC failure. The SPU in a device communicates with the Routing Engine to create a distributed VPN.

In IPsec, the workload is distributed by the same algorithm that distributes the IKE. The Phase 2 SA for a given VPN tunnel termination points pair is exclusively owned by a particular SPU, and all IPsec packets belonging to this Phase 2 SA are forwarded to the anchoring SPU of that security association for IPsec processing.

Related Topics

- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

Understanding IKE and IPsec Packet Processing

An IPsec VPN tunnel consists of tunnel setup and applied security. During tunnel setup, the peers establish security associations (SAs), which define the parameters for securing traffic between themselves. (For more information, see "Virtual Private Networks (VPNs) Overview" on page 355.) After the tunnel is established, IPsec protects the traffic sent between the two tunnel endpoints by applying the security parameters defined by the SAs during tunnel setup. Within the JUNOS implementation, IPsec is applied in tunnel mode, which supports the Encapsulating Security Payload (ESP) and Authentication Header (AH) protocols.

This topic covers:

- Packet Processing in Tunnel Mode on page 364
- IKE Packet Processing on page 365
- IPsec Packet Processing on page 367
- Related Topics on page 369

Packet Processing in Tunnel Mode

IPsec operates in in one of two modes—transport or tunnel. When both ends of the tunnel are hosts, you can use either mode. When at least one of the endpoints of a tunnel is a security gateway, such as a JUNOS router or firewall, you must use tunnel mode. Juniper Networks devices always operate in tunnel mode for IPsec tunnels.

In tunnel mode, the entire original IP packet—payload and header—is encapsulated within another IP payload and a new header is appended to it, as shown in Figure 43 on page 364. The entire original packet can be encrypted, authenticated, or both. With the Authentication Header (AH) protocol, the AH and new headers are also authenticated. With the Encapsulating Security Payload (ESP) protocol, the ESP header can also be authenticated.

Figure 43: Tunnel Mode

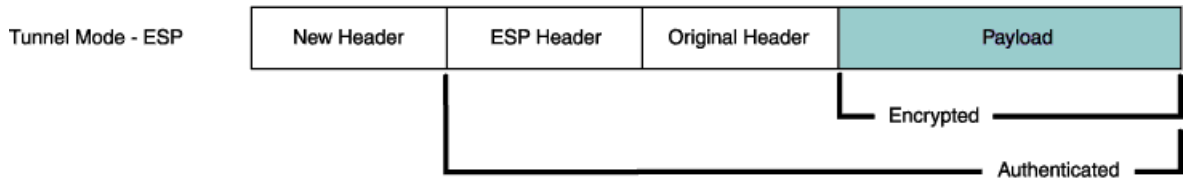

In a site-to-site VPN, the source and destination addresses used in the new header are the IP addresses of the outgoing interface. See Figure 44 on page 365.

Figure 44: Site-to-Site VPN in Tunnel Mode

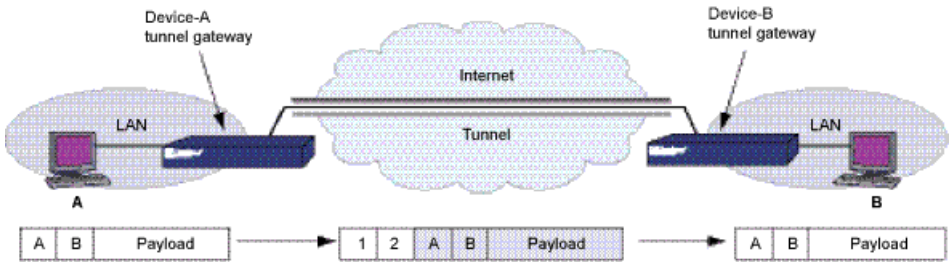

In a dial-up VPN, there is no tunnel gateway on the VPN dial-up client end of the tunnel; the tunnel extends directly to the client itself (see Figure 45 on page 365). In this case, on packets sent from the dial-up client, both the new header and the encapsulated original header have the same IP address: that of the client's computer. See Figure 45 on page 365.

NOTE: Some VPN clients such as the dynamic VPN client and Netscreen-Remote use a virtual inner IP address (also called a “sticky address”). Netscreen-Remote enables you to define the virtual IP address. The dynamic VPN client uses the virtual IP address assigned by the Radius server during the Xauth configuration exchange. In such cases, the virtual inner IP address is the source IP address in the original packet header of traffic originating from the client, and the IP address that the ISP dynamically assigns the dial-up client is the source IP address in the outer header.

Figure 45: Dial-up VPN in Tunnel Mode

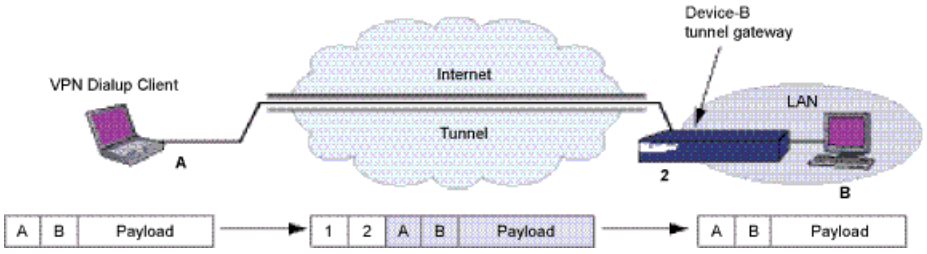

IKE Packet Processing

When a clear-text packet arrives on a Juniper Networks device that requires tunneling, and no active Phase 2 SA exists for that tunnel, JUNOS Software begins IKE negotiations and drops the packet. The source and destination addresses in the IP packet header are those of the local and remote IKE gateways, respectively. In the IP packet payload, there is a UDP segment encapsulating an ISAKMP (IKE) packet. The format for IKE packets is the same for Phase 1 and Phase 2. See Figure 46 on page 366.

Meanwhile, the source host has resent the dropped packet. Typically, by the time the second packet arrives, IKE negotiations are complete and JUNOS Software protects it—and all subsequent packets in the session—with IPsec before forwarding it.

Figure 46: IKE Packet for Phases 1 and 2

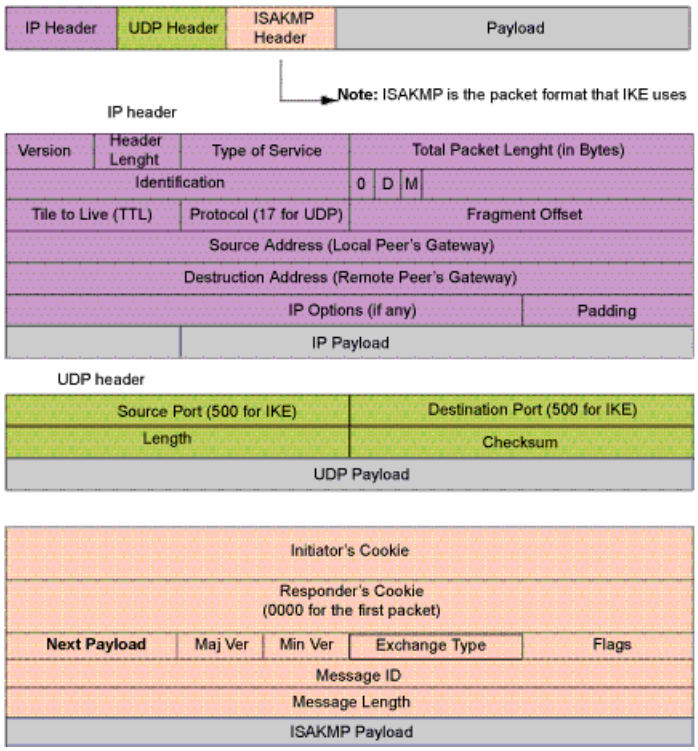

The Next Payload field contains a number indicating one of the following payload types:

- 0002—SA Negotiation Payload contains a definition for a Phase 1 or Phase 2 SA.
- 0004—Proposal Payload can be a Phase 1 or Phase 2 proposal.
- 0008—Transform Payload gets encapsulated in a proposal payload that gets encapsulated in an SA payload.
- 0010—Key Exchange (KE) Payload contains information necessary to perform a key exchange, such as a Diffie-Hellman public value.
- 0020—Identification (IDx) Payload.
  - In Phase 1, IDii indicates the initiator ID, and IDir indicates the responder ID.
  - In Phase 2, IDui indicates the user initiator, and IDur indicates the user responder.

  The IDs are IKE ID types such as FQDN, U-FQDN, IP address, and ASN.1_DN.
- 0040—Certificate (CERT) Payload.
- 0080—Certificate Request (CERT_REQ) Payload.
- 0100—Hash (HASH) Payload contains the digest output of a particular hash function.

- 0200—Signature (SIG) Payload contains a digital signature.
- 0400—Nonce (Nx) Payload contains some pseudo-random information necessary for the exchange).
- 0800—Notify Payload.
- 1000—ISAKMP Delete Payload.
- 2000—Vendor ID (VID) Payload can be included anywhere in Phase 1 negotiations. JUNOS Software uses it to mark support for NAT-T.

Each ISAKMP payload begins with the same generic header, as shown in Figure 47 on page 367.

Figure 47: Generic ISAKMP Payload Header

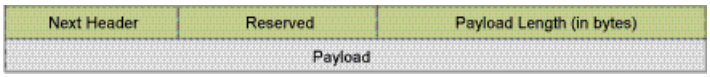

There can be multiple ISAKMP payloads chained together, with each subsequent payload type indicated by the value in the Next Header field. A value of **0000** indicates the last ISAKMP payload. See Figure 48 on page 367 for an example.

Figure 48: ISAKMP Header with Generic ISAKMP Payloads

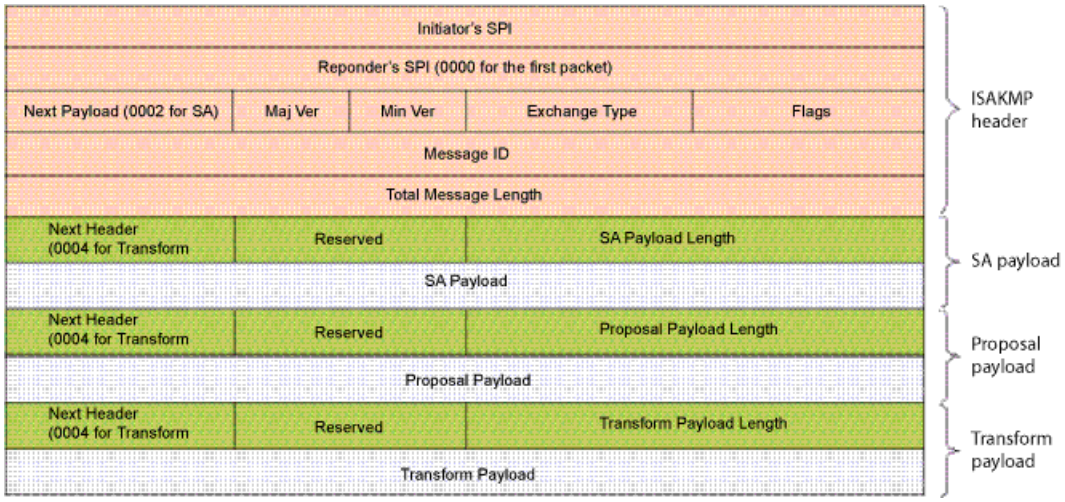

IPsec Packet Processing

After IKE negotiations complete and the two IKE gateways have established Phase 1 and Phase 2 security associations (SAs), the device applies IPsec protection to subsequent clear-text IP packets that hosts behind one IKE gateway send to hosts behind the other gateway (assuming that policies permit the traffic). If the Phase 2 SA specifies the Encapsulating Security Protocol (ESP) in tunnel mode, the packet looks like the one shown below. The device adds two additional headers to the original packet that the initiating host sends.

NOTE: For information about ESP, see “Encapsulating Security Payload (ESP) Protocol” on page 359. For information about tunnel mode, see “Packet Processing in Tunnel Mode” on page 364.

As shown in Figure 49 on page 368, the packet that the initiating host constructs includes the payload, the TCP header, and the inner IP header (IP1).

Figure 49: IPsec Packet—ESP in Tunnel Mode

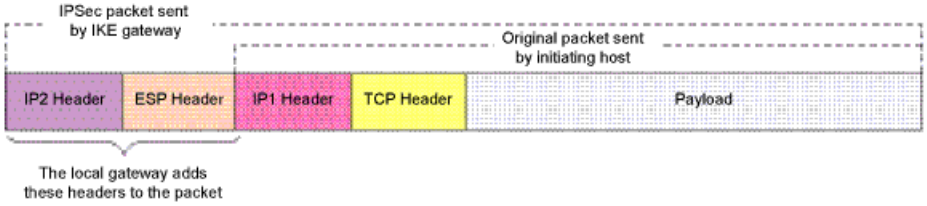

The outer IP header (IP2), which JUNOS Software adds, contains the IP address of the remote gateway as the destination IP address and the IP address of the local router as the source IP address. JUNOS Software also adds an ESP header between the outer and inner IP headers. The ESP header contains information that allows the remote peer to properly process the packet when it receives it. This is illustrated in Figure 50 on page 368.

Figure 50: Outer IP Header (IP2) and ESP Header

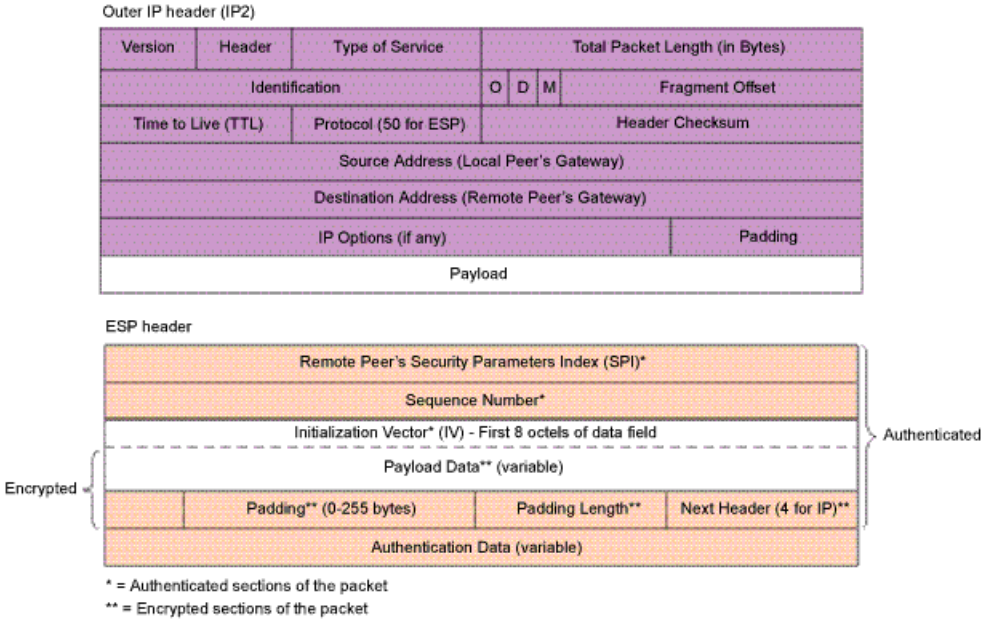

The Next Header field indicates the type of data in the payload field. In tunnel mode, this value is 4, indicating IP-in-IP. See Figure 51 on page 369.

Figure 51: Inner IP Header (IP1) and TCP Header

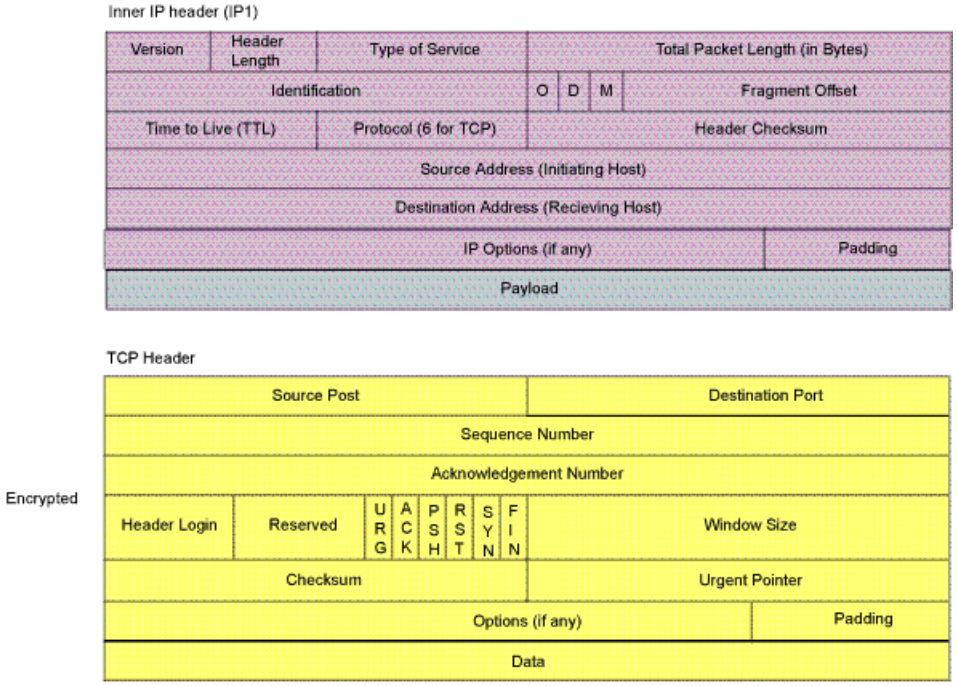

Related Topics

- Virtual Private Networks (VPNs) Overview on page 355
- Configuring an IPsec Tunnel—Overview on page 369

Configuring an IPsec Tunnel—Overview

IKE IPsec tunnel negotiation occurs in two phases. In Phase 1, participants establish a secure channel in which to negotiate the IPsec security association (SA). In Phase 2, participants negotiate the IPsec SA for authenticating traffic that will flow through the tunnel.

Before You Begin

For background information, read

- "Virtual Private Networks (VPNs) Overview" on page 355.
- "Understanding IKE and IPsec Packet Processing" on page 364.

Tunnel configuration can be confusing, and a good way to understand it is to keep in mind that just as there are two phases to tunnel negotiation, there are two phases to tunnel configuration. The following procedure lists the order in which you must configure an IPsec tunnel if you use J-Web. Although you need not follow this sequence when using the CLI configuration editor, we recommend that you do. If, for example,

you go out of sequence and configure a Phase 1 policy before you have configured a proposal, you cannot easily reference the proposal in the policy because it will not appear in the interface.

1. Phase 1
   a. Configure IKE Phase 1 proposals. For instructions, see:
      - Configuring an IKE Phase 1 Proposal (Standard and Dynamic VPNs) on page 373
   b. Configure IKE policies (and reference the proposals). For instructions, see:
      - Configuring an IKE Policy (Standard and Dynamic VPNs) on page 384
   c. Configure IKE gateway (and reference the policy). For instructions, see:
      - Configuring an IKE Gateway (Standard and Dynamic VPNs) on page 385
2. Phase 2
   a. Configure Phase 2 proposals. For instructions, see:
      - Configuring an IPsec Phase 2 Proposal (Standard and Dynamic VPNs) on page 387
   b. Configure policies (and reference proposals). For instructions, see:
      - Configuring an IPsec Policy (Standard and Dynamic VPNs) on page 388
   c. Configure IPsec Autokey IKE (and reference the policy and gateway). For instructions, see:
      - Configuring IPsec AutoKey (Standard and Dynamic VPNs) on page 390

In addition to configuring tunnel settings, you should also review your global VPN settings. For more information, see:

- Configuring VPN Global Settings (Standard VPNs) on page 370
- Configuring IPSec VPN Global Settings in J-Web (Standard VPNs) on page 372

Configuring VPN Global Settings (Standard VPNs)

Global VPN settings help you monitor and maintain the efficient operation of your VPN. The following example-based instructions show how to configure global IPsec VPN settings.

Before You Begin

For background information, read

- "Virtual Private Networks (VPNs) Overview" on page 355.
- "Understanding IKE and IPsec Packet Processing" on page 364.
- Configuring an IPsec Tunnel—Overview on page 369

Peers in a Security Association (SA) can become unsynchronized when one of the peers fails, for example, and reboots, causing it to send an incorrect SPI. You enable the device to detect such an event and resynchronize the peers by configuring the bad SPI response feature, and VPN monitoring.

In this example, you configure the device to detect and respond five times to a bad IPsec SPI before deleting the SA and initiating a new one. You also configure the device to monitor the VPN by sending Internet Control Message Protocol (ICMP) requests to the peer every 15 seconds, and to declare the peer unreachable after 15 unsuccessful pings.

To configure VPN global settings, use either the J-Web or the CLI configuration editor. (For information about configuring global settings using J-Web Configuration pages, see "Configuring IPSec VPN Global Settings in J-Web (Standard VPNs)" on page 372.)

This topic covers:

- J-Web Configuration on page 371
- CLI Configuration on page 372

J-Web Configuration

To configure global VPN settings in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ike, click **Configure** or **Edit**.
4. Check the **Respond Bad SPI** check box, then click **Configure**.
5. Enter a value in the **Max responses** field.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To apply the configuration, click **Commit**.
   - To cancel your entries and return to the main page, click **Cancel**.
7. Select **Configure > CLI Tools > Point and Click CLI**.
8. Next to Security, click **Configure** or **Edit**.
9. Next to Ipsec, click **Configure** or **Edit**.
10. Check the **VPN monitor options** check box, and then click **Configure**.
11. Enter a value in the **Interval** field.
12. Enter a value in the **Threshold** field.
13. Click one of the following buttons:
    - To apply the configuration and return to the main Configuration page, click **OK**.
    - To apply the configuration, click **Commit**.

- To cancel your entries and return to the main page, click **Cancel**.

CLI Configuration

To configure global VPN settings in the CLI editor:

```
user@host# set security ike respond-bad-spi 5
user@host# set security ipsec vpn-monitor-options interval 15 threshold 15
```

Configuring IPSec VPN Global Settings in J-Web (Standard VPNs)

You can use J-Web Configuration to quickly configure VPN global settings. (For information about configuring global settings using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Configuring VPN Global Settings (Standard VPNs)" on page 370.)

Before You Begin
For background information, read ■ "Virtual Private Networks (VPNs) Overview" on page 355. ■ "Understanding IKE and IPsec Packet Processing" on page 364. ■ Configuring an IPsec Tunnel—Overview on page 369

To access IPsec VPN using J-Web:

1. Select **Configure > IPSec VPN > Global Settings**.

   The details of the pane are described in :Table 52 on page 372
2. Click one:
   - To apply changes to the configuration, click **Save**.
   - To reset the configuration without saving changes, click .**Reset.**

Table 52: VPN Global Configuration Options

Field	Function	Action
IKE Global Settings		
Response Bad SPI	Enable response to invalid IPsec Security Parameter Index (SPI) values. If the SAs between two peers of an IPsec VPN become unsynchronized, the device resets the state of a peer so that the two peers are synchronized.	Click the check box if you want the device to respond to IPsec packets with bad SPI values.
Maximum Responses	Number of times to respond to invalid SPI values per gateway	Enter a value between 1 and 30. The default is 5. Disabled by default.
IPsec Global Settings		

Table 52: VPN Global Configuration Options *(continued)*

Field	Function	Action
VPN Monitor Options	Configure VPN monitoring options.	Click the check box if you want the device to monitor VPN liveliness.
Interval	Interval at which to send ICMP requests to the peer.	Enter a value from 1 through 36,000 seconds.
Threshold	The number of consecutive unsuccessful pings before the peer is declared unreachable.	Enter a value from 1 through 65,536.

Configuring an IKE Phase 1 Proposal (Standard and Dynamic VPNs)

When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway. The following example-based instructions show how to create the proposal portion of the IKE gateway.

Before You Begin

For background information about standard IPsec VPNs, read:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

For background information about dynamic IPsec VPNs, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

In Phase 1 proposal configuration, you must set the authentication method and authentication and encryption algorithms that will be used to open a secure channel between participants. In this example, you create an IKE proposal called ike_prop_1 and specify that peers use preshared keys for encryption and decryption, and that they use Diffie-Helman group 2 to produce the shared secret for the keys. You specify md5 as the authentication algorithm and 3DES cypher block chaining (CBC) for encryption. And you specify that after 300 seconds the participants renegotiate a new security association (SA).

NOTE: When configuring a Phase 1 proposal for the dynamic VPN feature, note that you must set the authentication method to preshared keys.

To configure Phase 1 proposals, use the J-Web or CLI configuration editor. (For information about configuring Phase 1 proposals using the J-Web Configuration pages, see "Configuring Auto Tunnel in J-Web (Standard VPN)" on page 375 or "Configuring an IKE Phase 1 Proposal— (Dynamic VPNs)" on page 451.)

This topic covers:

- J-Web Configuration on page 374
- CLI Configuration on page 374

J-Web Configuration

To configure a Phase 1 proposal in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ike, click **Configure** or **Edit**.
4. Next to Proposal, click **Add new entry**.
5. In the Name box, type **Ike_prop_1**.
6. From the Authentication algorithm list, select **md5**.
7. From the Authentication method list, select **pre-shared-keys.**
8. In the Description box, type **new Ike proposal**.
9. From the Dh-group list, select **group2**.
10. From the Encryption algorithm list, select **3des-cbc**.
11. In the Lifetime seconds box, type **300** and click **OK**.
12. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure a Phase 1 proposal using the CLI editor:

```
user@host# set security ike proposal ike_prop_1 description "new ike proposal"
user@host# set security ike proposal ike_prop_1 authentication-method
  pre-shared-keys
user@host# set security ike proposal ike_prop_1 dh-group group2
user@host# set security ike proposal ike_prop_1 authentication-algorithm md5
user@host# set security ike proposal ike_prop_1 encryption-algorithm 3des-cbc
user@host# set security ike proposal ike_prop_1 lifetime-seconds 300
```

Use the following command to display information about IKE proposals:

```
user@host# show security ike
```

Configuring Auto Tunnel in J-Web (Standard VPN)

Use the following procedure to configure the Auto Tunnel in J-Web.

Before You Begin
For background information, read ■ "Virtual Private Networks (VPNs) Overview" on page 355. ■ "Understanding IKE and IPsec Packet Processing" on page 364. ■ Configuring an IPsec Tunnel—Overview on page 369

To access the Auto Tunnel using J-Web:

1. In the J-Web user interface, select Configure > Auto Tunnel.The details of the display page are provided inTable 53 on page 375 and Table 54 on page 376.
2. Auto Tunnel options have the following suboptions:
   - Select Configure > Phase 1
   - Select Configure > Phase 2
3. Click one of the following:
   - Add—Adds a new Gateway. Enter information as specified in Table 55 on page 377 and Table 56 on page 381, for Phase 1 and Phase 2, respectively.
   - Edit—Modifies a selected Gateway.
   - Delete— Deletes a selected Gateway.

Table 53: : Display page for Auto Tunnel – Phase 1

Field	Function
Gateway	
Gateway Name	Name of the gateway to be searched.
Search	Text box for searching a gateway.
Name	Name of the destination peer gateway, specified as an alphanumeric string.
External Interface	Name of the interface to be used to send traffic to the IPsec VPN.
Remote Identity	Provides information about remote peer.
IKE Policy	
Name	Name of the policy.

Table 53: : Display page for Auto Tunnel – Phase 1 *(continued)*

Field	Function
Description	Description of the policy.
Mode	There are two modes: ■ Main mode has three two-way exchanges between the initiator and receiver. It is secure and preferred in the Auto Tunnel ■ Aggressive mode is faster than main mode. It is less secure and used mostly for dial-up VPN.
Authentication Method	Authentication Method configured.
Proposal	Name of the proposal configured to be used by this policy in phase 1
Proposal	
Name	Name of the proposal selected.
Authentication Algorithm	Hash algorithm configured or selected.
Authentication Method	Authentication method selected.
Encryption Algorithm	Supported Internet Key Exchange (IKE) proposals.

Table 54: Display page for Auto Tunnel – Phase 2

Field	Function
VPN name	Name of the VPN to be searched.
Search	Radio button to search a specific VPN listed.
Name	Name of the VPN.
Gateway	Name of the gateway.
IPSec Policy	Associate a policy with this IPsec tunnel.
Bind Interface	The tunnel interface to which the route-based VPN is bound.
Proxy Identity	The IPsec proxy identity.
VPN Monitoring	Name of the VPN monitoring option selected.
IPSec Policy	
Name	Name of the IPsec policy.
Description	Description of the policy.

Table 54: Display page for Auto Tunnel – Phase 2 *(continued)*

Field	Function
Perfect Forward Secrecy	The method the device uses to generate the encryption key. PFS generates each new encryption key independently from the previous key. ■ group1—Diffie-Hellman Group 1. ■ group2—Diffie-Hellman Group 2. ■ group5—Diffie-Hellman Group 5.
Proposal	Name of the proposal to be used by IPsec policy in Phase 2.
Proposal	
Name	Description of the Phase 2 proposal.
Authentication Algorithm	Hash algorithm that authenticates packet data. It can be one of the following: ■ hmac-md5-96—Produces a 128-bit digest. ■ Hmac-sha1-96—Produces a 160-bit digest
Protocol	The type of security protocol.
Encryption algorithm	Configures an IKE encryption algorithm. ■ des-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—AES 128-bit encryption algorithm. ■ aes-192-cbc—AES 192-bit encryption algorithm. ■ aes-256-cbc—AES 256-bit encryption algorithm.

Table 55: Phase 1(VPN Configuration) Add Options

Field	Function	Action
Gateway > IKE Gateway		
Name	Name of the gateway	Enter the name of the gateway.
Policy		Enter the name of policy you configured for Phase 1.
External Interface	Name of the interface to be used to send traffic to the IPsec VPN.	Specify the outgoing interface for IKE SAs. This interface is associated with a zone that acts as its carrier, providing firewall security for it.
Site to Site Tunnel	Configuration for VPN is of type site to site.	Click the radio button labeled Site to Site.

Table 55: Phase 1(VPN Configuration) Add Options *(continued)*

Field	Function	Action
Address/FQDN	Address or fully qualified domain name (FQDN) of the peer.	Provide information about the peer IP or domain name
Local ID		
Identify Type	There are four identify types: ■ IP Address ■ Host Name ■ Email Address ■ Distinguished Name	Select one of the identity type options.
Client Tunnel	The remote access dynamic VPN.	
Gatewayl > KE Gateway Options		
Local Identity	The local IKE identity to send in the exchange with the destination peer so that the destination peer can communicate with the local peer. If you do not configure a local identity, the device uses the IP address corresponding to the local endpoint. You can identify the local identity in any of the following ways: ■ IP Address—IPv4 IP address to identify the dynamic peer. ■ Hostname—Fully qualified domain name (FQDN) to identify the dynamic peer. ■ User at Hostname—E-mail address to identify the dynamic peer. ■ Distinguished Name—Name to identify the dynamic peer. The distinguished name appears in the subject line of the Public Key Infrastructure (PKI) certificate. For example: Organization: juniper, Organizational unit: slt, Common name: common.	Specify an IP address, hostname, user-at-hostname, or distinguished name.
Dead Peer Detection		
Always send	Instructs the device to send dead peer detection (DPD) requests regardless of whether there is outgoing IPsec traffic to the peer.	Click the check box.
Interval	The amount of time that the peer waits for traffic from its destination peer before sending a DPD request packet.	Enter the interval at which to send DPD messages. Range: 1 through 60 seconds.

Table 55: Phase 1(VPN Configuration) Add Options *(continued)*

Field	Function	Action
Threshold	The maximum number of unsuccessful DPD requests that can be sent before the peer is considered unavailable.	Enter the maximum number of unsuccessful DPD requests to be sent. Range: 1 through 5. Default: 5.
XAuth	Provides extended authentication (XAuth) in addition to IKE authentication for remote users trying to access a VPN tunnel.	Enter extended authentication (XAuth).
NAT-Traversal	Network Address Translation Traversal (NAT-T). NAT-T is enabled by default.	Click the check box to disable or enable.
IKE Policy > IKE Policy		
Name	Name of the IKE Policy.	Enter the policy.
Description	Description of the policy.	Enter the description of the policy.
Mode	Select a mode.	Use Main or Aggressive mode.
Proposal		
Predefined	Use one of the following types of predefined Phase 1 proposals: ■ Basic ■ Compatible ■ Standard	Click Predefined and select a proposal type.
User defined	Use a user-defined Phase 1 Proposal.	Click User Defined, select a proposal from the pop-up menu, and click Add.
Proposal List	Specfies one or more proposals that can be used during key negotiation: ■ Available P1 proposal ■ Selected P1 proposal	Click the Predefined Proposal radio button to select proposals preconfigured by JUNOS Software. Select User Defined Proposal if you want to use proposals that you have created.
IKE Policy > IKE Policy Options		
Radio Buttons	Select the preshared key of use of certificate for the VPN.	If a preshared key is selected, then configure the appropriate key in the form of ASCII text or hexadecimal.
Local Certificate	Use a particular certificate when the local device has multiple loaded certificates.	Enter a local certificate identifier.
Peer Certificate Type	Use a preferred type of certificate (PKCS7 or X509).	Select a certificate type.
Trusted CA		
Name	Name of the proposal.	Enter the name of the proposal.

Table 55: Phase 1(VPN Configuration) Add Options *(continued)*

Field	Function	Action
Authentication Algorithm	The Authentication Header (AH) algorithm the device uses to verify the authenticity and integrity of a packet. Supported algorithms include the following: ■ md5—Produces a 128-bit digest. ■ sha1—Produces a 160-bit digest. ■ sha-256—Produces a 256-bit digest. **NOTE:** The sha-256 authentication algorithm is not supported with the dynamic VPN feature.	Select a hash algorithm.
Authentication Method	The method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. Options include: ■ pre-shared-keys—Key for encryption and decryption that both participants must have before beginning tunnel negotiations. ■ rsa-key—Kinds of digital signatures, which are certificates that confirm the identity of the certificate holder.	Select an authentication method.
Description	Easy identification of the proposal.	Enter a brief description of the IKE proposal.
DH Group	The Diffie-Hellman exchange allows participants to produce a shared secret value over an unsecured medium without actually transmitting the value across the connection.	Select a group. If you configure multiple (up to four) proposals for Phase 1 negotiations, use the same Diffie-Hellman group in all proposals.
Encryption Algorithm	Supported Internet Key Exchange (IKE) proposals include the following: ■ 3des-cbc—3DES-CBC encryption algorithm. ■ aes-128-cbc—AES-CBC 128-bit encryption algorithm. ■ aes-192-cbc—AES-CBC 192-bit encryption algorithm. ■ aes-256-cbc—AES-CBC 256-bit encryption algorithm. ■ des-cbc—DES-CBC encryption algorithm.	Select an encryption algorithm.

Table 55: Phase 1(VPN Configuration) Add Options *(continued)*

Field	Function	Action
Lifetime seconds	The lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.	Select a lifetime for the IKE SA. Default: 3,600 seconds. Range: 180 through 86,400 seconds.

Table 56: Phase 2 (IPsec Autokey Configuration) Add Options

Field	Function	Action
Gateway > IPsec VPN		
VPN Name	Name of the remote gateway.	Enter a name.
Remote Gateway	Associates a policy with this IPsec tunnel.	Select a name.
IPsec Policy	The tunnel interface to which the route-based virtual private network (VPN) is bound.	Select a policy.
Bind to tunnel interface	Specifies when IKE is activated. ■ immediately—IKE is activated immediately after VPN configuration and configuration changes are committed. ■ on-traffic—IKE is activated only when data traffic flows and must be negotiated.	Select an interface.
Establish tunnels	Specifies when IKE is activated. ■ immediately—IKE is activated immediately after VPN configuration and changes are committed. ■ on-traffic—IKE is activated only when data traffic flows and must be negotiated.	Choose an option.
Disable anti replay	Disable the anti-replay checking feature of IPsec. By default, anti-replay checking is enabled.	Click the check box.
Gateway > IPsec VPN Options		
Enable VPN Monitor		
Destination IP	Associates a policy with this IPsec tunnel.	Enter an IP address.
Optimized	The tunnel interface to which the route-based virtual private network (VPN) is bound.	Click the check box.

Table 56: Phase 2 (IPsec Autokey Configuration) Add Options *(continued)*

Field	Function	Action
Source Interface	Specifies when IKE is activated. ■ immediately—IKE is activated immediately after VPN configuration and configuration changes are committed. ■ on-traffic—IKE is activated only when data traffic flows and must be negotiated.	Specify a source interface.
Use Proxy Identity		
Remote IP/Netmask	The remote IP address and subnet mask for proxy identity.	Enter an IP address and
Service	The service (port and protocol combination) to protect.	Select a service.
Do not fragment bit	Specifies how the device handles the Don't Fragment (DF) bit in the outer header. ■ clear—Clear (disable) the DF bit from the outer header. This is the default. ■ copy—Copy the DF bit to the outer header. ■ set—Set (enable) the DF bit in the outer header.	Choose an option.
Install interval	The maximum number of seconds to allow installation of a rekeyed outbound security association (SA) on the device.	Specify a value between 0 and 10 seconds.
IPsec Policy > IPsec Policy		
Name	Name of the remote gateway.	Enter a name.
Description	Associates a policy with this IPsec tunnel.	Enter a text description.
Perfect Forward Secrecy	The tunnel interface to which the route-based virtual private network (VPN) is bound.	Select a method.
Proposal	Specifies when IKE is activated. ■ immediately—IKE is activated immediately after VPN configuration and configuration changes are committed. ■ on-traffic—IKE is activated only when data traffic flows and must be negotiated.	

Table 56: Phase 2 (IPsec Autokey Configuration) Add Options *(continued)*

Field	Function	Action
Predefined	Disable the anti-replay checking feature of IPsec. By default, anti-replay checking is enabled.	Click Predefined, and select one of the following options: ■ basic ■ predefined ■ standard
User defined	A list of proposals you previously defined.	Click User Defined , select proposals from the pop-up menu, and then click Add.
Proposal List	Available Proposal List are: ■ Available P2 Proposal ■ Selected P2 Proposal	
Proposal > IPsec Proposal		
Name	Name of the Phase 2 proposal.	Enter a name.
Description	Description of the Phase 2 proposal.	Enter a text description.
Authentication Algorithm	Hash algorithm that authenticates packet data. It can be one of the following: ■ hmac-md5-96—Produces a 128-bit digest. ■ hmac-sha1-96—Produces a 160-bit digest.	Select a hash algorithm.
Encryption Algorithm	Configures an IKE encryption algorithm. ■ 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—AES 128-bit encryption algorithm. ■ aes-192-cbc—AES 192-bit encryption algorithm. ■ aes-256-cbc—AES 256-bit encryption algorithm.	Select an encryption algorithm.
Lifetime Kilobytes	The lifetime (in kilobytes) of an IPsec security association (SA). The SA is terminated when the specified number of kilobytes of traffic has passed.	Enter a value from 64 through 1,048,576 bytes.
Lifetime Seconds Protocol	The lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.	Enter a value from 180 through 86,400 seconds.

Configuring an IKE Policy (Standard and Dynamic VPNs)

When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway. The following example-based instructions show how to create the policy portion of the IKE gateway.

Before You Begin

For background information about standard IPsec VPNs, read:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

For background information about dynamic IPsec VPNs, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

During policy configuration, you must set the mode in which the Phase 1 channel will be negotiated, specify the type of key exchange to be used, and reference the Phase 1 proposal. In this example, you create a policy called ike_pol_1, specify that participants exchange proposals in aggressive mode, and reference the proposal called ike_prop_1. You specifiy that the preshared key be of type ASCII, and enter the key.

NOTE: When configuring an IKE policy for the dynamic VPN feature, note that you must set the mode to aggressive. Also note that you must use preshared keys rather than manual keys or certificates.

This topic covers:

- J-Web Configuration on page 384
- CLI Configuration on page 385

J-Web Configuration

To configure an IKE policy in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ike, click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry**.
5. In the Name box, type **Ike_pol_1**.
6. In the Description box, type **new Ike policy**.

7. From the Mode box, select **aggressive**.
8. Next to Pre shared key, click **Configure**.
9. From the key choice list, select **Ascii text**.
10. In the Ascii textbox, type **9UQiqf36A1RSTzRSreXxDik.Tzn/CuBI** and click **OK**.
11. Next to Proposals, click **Add new entry**.
12. In the Value keyword box, type **Ike_prop_1** and click **OK**.
13. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure and IKE policy using the CLI Editor:

```
user@host# set security ike policy ike_pol_1 mode aggressive
user@host# set security ike policy ike_pol_1 description "new ike policy"
user@host# set security ike policy ike_pol_1 proposals ike_prop_1
user@host# set security ike policy ike_pol_1 pre-shared-key ascii-text
  "$9$UQiqf36A1RSTzRSreXxDik.Tzn/CuBI"
```

Use the following command to display information about this IKE policy:

```
user@host# show security ike policy ike_pol_1
```

Configuring an IKE Gateway (Standard and Dynamic VPNs)

When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway. The following example-based instructions show how to create the IKE gateway.

Before You Begin

For background information about standard IPsec VPNs, read:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

For background information about dynamic IPsec VPNs, read:

- “Dynamic VPN Overview” on page 433.
- “Configuring a Dynamic VPN—Overview” on page 449.

When creating the gateway, you must reference the Phase 1 policy. In this example, you create an IKE gateway called ike_gateway_1, reference the policy ike_pol_1, and configure an IP address for the gateway. You configure dead peer detection (DPD) to send a DPD request packet when the device has not received traffic from a peer for 10 seconds, and to consider the peer unavailable after five sequences of waiting

10 seconds and sending a DPD request packet. You also specify ge-0/0/0 as the outgoing interface.

This topic covers:

- J-Web Configuration on page 386
- CLI Configuration on page 386

J-Web Configuration

To configure an IKE gateway in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ike, click **Configure** or **Edit**.
4. Next to Gateway, click **Add new entry**.
5. In the Gateway name box, type **Ike_gateway_1**.
6. Next to Dead peer detection, select the check box and click **Configure**.
7. In the Interval box, type **10**.
8. In the Threshold box, type **5** and click **OK**.
9. In the External interface box, **type ge-0/0/0**.
10. In the Ike policy box, type **ike_pol_1**.
11. From the Remote identifier list, select **Address**.
12. Next to Address, click **Add new entry**.
13. In the Value box, type **1.1.1.2**.
14. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure an IKE gateway using the CLI editor:

```
user@host# set security ike gateway ike_gateway_1 ike-policy ike_pol_1
user@host# set security ike gateway ike_gateway_1 address 1.1.1.2
user@host# set security ike gateway ike_gateway_1 dead-peer-detection interval
  10
user@host# set security ike gateway ike_gateway_1 dead-peer-detection threshold
  5
user@host# set security ike gateway ike_gateway_1 external-interface ge-0/0/0
```

Use the following command to display information about this IKE gateway:

```
user@host# show security ike gateway ike_gateway_1
```

Configuring an IPsec Phase 2 Proposal (Standard and Dynamic VPNs)

When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE). The following example-based instructions show how to create the initial proposal.

Before You Begin

For background information about standard IPsec VPNs, read:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

For background information about dynamic IPsec VPNs, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

In Phase 2 proposal configuration, you must create a proposal, specify a security protocol, and select authentication and encryption algorithms for the traffic that will flow through the tunnel. In this example, you create a proposal called ipsec_prop_1, specify ESP as the security protocol, and set hmac-md5-96 as the authentication algorithm and 3des-cbc as the encryption algorithm. You also specify that the security association (SA) terminate after 1,800 KB of data pass through it.

This topic covers:

- J-Web Configuration on page 387
- CLI Configuration on page 388

J-Web Configuration

To configure an IPsec Phase 2 proposal in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ipsec, click **Configure** or **Edit**.
4. Next to Proposal, click **Add new entry**.
5. In the Name box, type **Ipsec_prop_1**.
6. From the Authentication algorithm, select **hmac-md5-96**.
7. In the Description box, type **new ipsec proposal**.
8. From the Encryption algorithm, select **3des-cbc**.
9. In the Lifetime seconds box, type **1800** and click **OK**.

10. From the Protocol list, select **esp**.
11. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure an IPsec Phase 2 proposal using the CLI editor:

```
user@host# set security ipsec proposal ipsec_prop_1 description "new ipsec
  proposal"
user@host# set security ipsec proposal ipsec_prop_1 protocol esp
user@host# set security ipsec proposal ipsec_prop_1 authentication-algorithm
  hmac-md5-96
user@host# set security ipsec proposal ipsec_prop_1 encryption-algorithm 3des-cbc
user@host# set security ipsec proposal ipsec_prop_1 lifetime-seconds 1800
```

Use the following command to display information about this IKE proposal:

```
user@host# show security ipsec proposal ipsec_prop_1
```

Configuring an IPsec Policy (Standard and Dynamic VPNs)

When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE). The following example-based instructions show how to create the policy.

Before You Begin

For background information about standard IPsec VPNs, read:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

For background information about dynamic IPsec VPNs, read:

- “Dynamic VPN Overview” on page 433.
- “Configuring a Dynamic VPN—Overview” on page 449.

In Phase 2 IPsec policy configuration, you must create a policy and reference a Phase 2 proposal. In this example, you create a policy called ipsec_pol_1 and reference the proposal ipsec_prop_1. You also configure Perfect Forward Secrecy to use Diffie-Hellman Group 2 as the method the device uses to generate the encryption key.

This topic covers:

- J-Web Configuration on page 389
- CLI Configuration on page 389

J-Web Configuration

To configure an IPsec policy in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ipsec, click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry**.
5. In the Name box, type **Ipsec_pol_1**.
6. In the Description box, type **new Ipsec policy**.
7. Next to Perfect forward secrecy, click **Configure**.
8. From the Keys list, select **group2** and click **OK**.
9. Next to Proposals, click **Add new entry**.
10. In the Value keyword box, type **Ipsec_prop_1** and click **OK**.
11. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure an IPsec policy using the CLI editor:

```
user@host# set security ipsec policy ipsec_pol_1 description "new ipsec policy"
user@host# set security ipsec policy ipsec_pol_1 perfect-forward-secrecy keys
  group2
user@host# set security ipsec policy ipsec_pol_1 proposals ipsec_prop_1
```

Use the following command to display information about this IKE proposal:

```
user@host# show security ipsec policy ipsec_pol_1
```

Configuring IPsec AutoKey (Standard and Dynamic VPNs)

When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE). The following example-based instructions show how to configure the IPsec AutoKey.

Before You Begin

For background information about standard IPsec VPNs, read:

- Virtual Private Networks (VPNs) Overview on page 355
- Understanding IKE and IPsec Packet Processing on page 364
- Configuring an IPsec Tunnel—Overview on page 369

For background information about dynamic IPsec VPNs, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

In Phase 2 IPsec AutoKey configuration, you must create a VPN tunnel name, specify a gateway, and reference a Phase 2 policy. If you are using Route mode, you must bind the tunnel to an interface. In this example, you create a VPN tunnel named vpn_1 and bind it to interface st0.0, and you specify ike_gateway_1 as the gateway for the VPN tunnel and reference the IPsec policy ipsec_pol_1.

This topic covers:

- J-Web Configuration on page 390
- CLI Configuration on page 391

J-Web Configuration

To configure an IPsec Autokey in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Ipsec, click **Configure** or **Edit**.
4. Next to Vpn, click **Add new entry**.
5. In the Name box, type **vpn_1**.
6. In the Bind interface box, type **st0.0**.
7. From the Negotiation list, select **Ike**.
8. Next to Ike, click **Configure**.
9. In the Gateway box, type **Ike_gateway_1**.
10. In the Ipsec policy box, type **Ipsec_pol_1** and click **OK**.
11. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure an IPsec Autokey using the CLI editor:

```
user@host# set security ipsec vpn vpn_1 bind-interface st0.0
user@host# set security ipsec vpn vpn_1 ike gateway ike_gateway_1
user@host# set security ipsec vpn vpn_1 ike ipsec-policy ipsec_pol_1
```

Use the following command to display information about this IKE proposal:

```
user@host# show security ipsec vpn vpn_1
```

Configuring the Manual Tunnel using J-Web

Use the following procedure to configuring auto tunnel.

To access Manual Tunnel using J-Web:

1. In the J-Web user interface, select Configure > Manual Tunnel. The details of the display page are as shown in Table 57 on page 391.
2. Click one of the following:
   - Add—Adds a manual tunnel. Enter information as specified in Table 58 on page 391.
   - Edit— Modifies the selected manual tunnel.
   - Delete— Deletes a selected manual tunnel.

Table 57: Display Page for Manual Tunnel

Field	Function
Name	Name of the manual tunnel.
Gateway	Gateway selected.
Bind Interface	The tunnel interface to which the route-based virtual private network (VPN) is bound.
Df Bit	Don't Fragment (DF) bit in the outer header.

Table 58: VPN Manual Configuration Add Options

Field	Function	Action
IPsec Manual		
VPN Name	Name of the VPN Name for the IPsec tunnel..	Enter the VPN Name.

Table 58: VPN Manual Configuration Add Options *(continued)*

Field	Function	Action
Remote Gateway	Name of the remote gateway.	Enter the gateway.
	Types of protocols available for configuration. AH stands for authentication and ESP stands for Encryption algorithm ■ ESP ■ AH	Select a protocol for the proposal.
SIP	Security Parameter Index values ranging from (256..16639).	Select a policy.
Blind to tunnel interface	The tunnel interface to which the route-based virtual private network (VPN) is bound.	Select an interface.
Do not fragment bit	Specifies how the device handles the Don't Fragment (DF) bit in the outer header. ■ clear—Clear (disable) the DF bit from the outer header. This is the default. ■ copy—Copy the DF bit to the outer header. ■ Set—Sets the DF bot to the outer header.	Choose an option.
Enable VPN Monitor		
Destination IP	IP address of the destination peer.	Enter an IP address.
Optimized	Specifies that the device uses traffic patterns as evidence of peer liveliness. If enabled, ICMP requests are suppressed. This feature is disabled by default.	Click the check box.
Source Interface	The source interface for ICMP requests (VPN monitoring “hellos”). If no source interface is specified, the device automatically uses the local tunnel endpoint interface.	Specify a source interface.
Key Values		
Authentication		

Table 58: VPN Manual Configuration Add Options *(continued)*

Field	Function	Action
Algorithm	Hash algorithm that authenticates packet data. It can be one of the following: ■ hmac-md5-96—Produces a 128-bit digest. ■ hmac-sha1-96—Produces a 160-bit digest.	Select a hash algorithm.
ASCII Text	Pre-shared value of key	Enable the ASCII Text option and enter the key in the appropriate format.
Hexadecimal	Pre-shared value of key	Enable the hexadecimal option and enter the key in the appropriate format.
Encryption	Supported Internet Key Exchange (IKE) proposals include the following: ■ 3des-cbc—3DES-CBC encryption algorithm. ■ aes-128-cbc—AES-CBC 128-bit encryption algorithm. ■ aes-192-cbc—AES-CBC 192-bit encryption algorithm. ■ aes-256-cbc—AES-CBC 256-bit encryption algorithm. ■ des-cbc—DES-CBC encryption algorithm	Select an encryption algorithm.
ASCII Text	Preshared value of key	Enable the ASCII Text option and enter the key in the appropriate format.
Hexadecimal	Pre shared	Enable the hexadecimal option and enter the key in the appropriate format.

Configuring Hub-and-Spoke VPNs

If you create two VPN tunnels that terminate at a device, you can set up a pair of routes so that the device directs traffic exiting one tunnel to the other tunnel. If both tunnels are contained within a single zone, you do not need to create a policy to permit the traffic to pass from one tunnel to the other. You only need to define the routes. Such an arrangement is known as *hub-and-spoke VPN*.

You can also configure multiple VPNs and route traffic between any two tunnels. For more information on IPsec tunnels and global VPNs, see “Understanding IKE and IPsec Packet Processing” on page 364 and “Virtual Private Networks (VPNs) Overview” on page 355.

NOTE: SRX Series devices support only the route-based hub-and-spoke feature.

Figure 52: Multiple Tunnels in a Hub-and-Spoke VPN Configuration

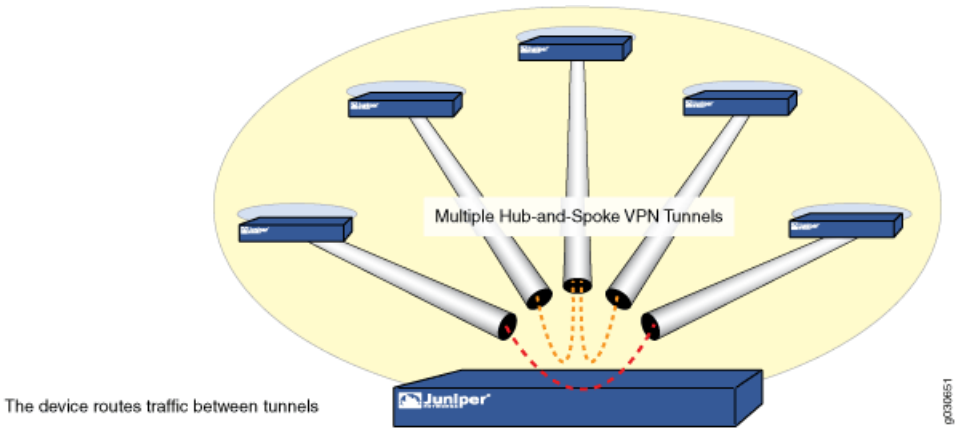

This topic covers:

- Creating Hub-and-Spoke VPNs on page 394

Creating Hub-and-Spoke VPNs

For the hub router to be able to distinguish between packets going to and coming from the spoke routers, you must configure it with two routing instances.

In this example, there are two spokes (First and Third) and the configuration is for route-based VPNs. Use the following to configure:

- Configuring the IPSec Tunnel on the Hub on page 394
- Configuring Spoke 1 on page 397
- Configuring Spoke 2 on page 398

Configuring the IPSec Tunnel on the Hub

The following procedure lists the order in which you must configure an IPsec tunnel by using the CLI configuration editor:

1. Phase 1
    a. When configuring Phase 1 of an IPsec tunnel using IKE, you must first configure proposals. In Phase 1 proposal configuration, set the authentication method and authentication and encryption algorithms that will be used to open a secure channel between participants. In this example, you create two proposals, first_ikeprop and third_ikeprop and specify that peers use preshared keys for encryption and decryption, and that they use Diffie-Hellman group 2 to produce the shared secret for the keys. You specify md5 as the authentication algorithm and 3DES cypher block chaining (CBC) for encryption.

NOTE: When configuring a Phase 1 proposal for the dynamic VPN feature, note that you must set the authentication method to preshared keys.

Configure IKE Phase 1 proposals:

```
user@host# set security ike proposal first_ikeprop authentication-method
  pre-shared-keys
user@host# set security ike proposal first_ikeprop dh-group group2
user@host# set security ike proposal first_ikeprop authentication-algorithm
  md5
user@host# set security ike proposal first_ikeprop encryption-algorithm
  3des-cbc

user@host# set security ike proposal third_ikeprop authentication-method
  pre-shared-keys
user@host# set security ike proposal third_ikeprop dh-group group2
user@host# set security ike proposal third_ikeprop authentication-algorithm
  md5
user@host# set security ike proposal third_ikeprop encryption-algorithm
  3des-cbc
```

b. When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies. During policy configuration, you must set the mode in which the Phase 1 channel will be negotiated, specify the type of key exchange to be used, and reference the Phase 1 proposal. In this example, you create two policies called first_ikepol and third_ikepol, specify that participants exchange proposals in aggressive mode, and reference the proposals called first_ikeprop and third_ikeprop. You specifiy that the preshared key be of type ASCII, and enter the key.

Configure IKE policies (and reference the proposals):

```
user@host# set security ike policy first_ikepol mode main
user@host# set security ike policy first_ikepol proposals first_ikeprop
user@host# set security ike policy first_ikepol pre-shared-key ascii-text
  "$9$xFU-b2ZUH5Qn4aQn/CB17-V"

user@host# set security ike policy third_ikepol mode main
user@host# set security ike policy third_ikepol proposals third_ikeprop
user@host# set security ike policy third_ikepol pre-shared-key ascii-text
  "$9$GvjkPFnCBIc5QIcylLXUjH"
```

c. When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway. When creating the gateway, you must reference the Phase 1 policies. In this example, you create IKE gateways called ike_gateway_first and ike_gateway_third, reference the policies first_ikepol and third_ikepol, and configure an IP address for the gateway.

Configure IKE gateway (and reference the policy):

```
user@host# set security ike gateway first ike-policy first_ikepol
user@host# set security ike gateway first address 4.4.4.2
```

```
user@host# set security ike gateway first external-interface ge-0/0/0.0

user@host# set security ike gateway third ike-policy third_ikepol
user@host# set security ike gateway third address 2.2.2.1
user@host# set security ike gateway third external-interface ge-0/0/3.0
```

2. Phase 2

   a. When configuring Phase 2 of an IPsec tunnel, you first configure proposals. In Phase 2 proposal configuration, you must create proposals for the two spokes, specify a security protocol, and select authentication and encryption algorithms for the traffic that will flow through the tunnel. In this example, you create proposals called first_ipsecprop and third_ipsecprop, set hmac-md5-96 as the authentication algorithm, and 3des-cbc as the encryption algorithm.

      Configure Phase 2 proposals:

      ```
      user@host# set security ipsec proposal first_ipsecprop protocol esp
      user@host# set security ipsec proposal first_ipsecprop
        authentication-algorithm hmac-md5-96
      user@host# set security ipsec proposal first_ipsecprop encryption-algorithm
        3des-cbc

      user@host# set security ipsec proposal third_ipsecprop protocol esp
      user@host# set security ipsec proposal third_ipsecprop
        authentication-algorithm hmac-md5-96
      user@host# set security ipsec proposal third_ipsecprop encryption-algorithm
        3des-cbc
      ```

   b. When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies. In Phase 2 IPsec policy configuration, you must create policies and reference the Phase 2 proposals. In this example, you create two policies called first_ipsecpol and third_ipsecpol and reference the proposals first_ipsecprop and third_ipsecprop.

      Configure policies (and reference proposals):

      ```
      user@host# set security ipsec policy first_ipsecpol perfect-forward-secrecy
        keys group1
      user@host# set security ipsec policy first_ipsecpol proposals first_ipsecprop

      user@host# set security ipsec policy third_ipsecpol perfect-forward-secrecy
        keys group1
      user@host# set security ipsec policy third_ipsecpol proposals third_ipsecprop
      ```

   c. When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE). In Phase 2 IPsec AutoKey configuration, you must create a VPN tunnel name, specify a gateway, and reference a Phase 2 policy. For Route mode, you must bind the tunnel to an interface. In this example, you create VPN tunnels named first_vpn and third_vpn and bind them to interface st0.0, and you specify ike_gateway_first and ike_gateway_third as the gateways for the VPN tunnel and reference the IPsec policies first_ipsecpol and third_ipsecpol.

Configure IPsec Autokey IKE (and reference the policy and gateway):

```
user@host# set security ipsec vpn first_vpn bind-interface st0.0
user@host# set security ipsec vpn first_vpn ike gateway first
user@host# set security ipsec vpn first_vpn ike ipsec-policy first_ipsecpol
user@host# set security ipsec vpn first_vpn establish-tunnels immediately

user@host# set security ipsec vpn third_vpn bind-interface st0.0
user@host# set security ipsec vpn third_vpn ike gateway third
user@host# set security ipsec vpn third_vpn ike ipsec-policy third_ipsecpol
user@host# set security ipsec vpn third_vpn establish-tunnels immediately
```

In addition to configuring tunnel settings, you also need to configure global policy, routing options and enable Next Hop Tunnel Binding (nhtb).

- Configure Security Policy

```
user@host# set security policies default-policy permit-all
```

- Configure Routing Options

```
user@host# set routing-options static route 1.1.1.0/24 next-hop st0.0
user@host# set routing-options static route 3.1.1.0/24 next-hop st0.0
```

- Enable Next Hop Tunnel Binding (nhtb)

```
user@host# set interfaces st0 unit 0 multipoint
```

Configuring Spoke 1

After you have configured the hub, you configure the spokes. Follow the same process to configure spoke First. When configuring Phase 1, you first configure proposals, then policies, and finally you configure the gateway.

1. Phase 1
   a. Configure IKE Phase 1 proposals:

```
user@host# set security ike proposal ike_prop authentication-method
  pre-shared-keys
user@host# set security ike proposal ike_prop dh-group group2
user@host# set security ike proposal ike_prop authentication-algorithm md5
user@host# set security ike proposal ike_prop encryption-algorithm 3des-cbc
```

   b. Configure IKE policies (and reference the proposals):

```
user@host# set security ike policy ike_pol mode main
user@host# set security ike policy ike_pol proposals ike_prop
user@host# set security ike policy ike_pol pre-shared-key ascii-text
  "$9$vA38xd24Zk.5bs.5QFAtM8X"
```

   c. Configure IKE gateway (and reference the policy):

```
user@host# set security ike gateway first ike-policy ike_pol
user@host# set security ike gateway first address 4.4.4.1
```

```
user@host# set security ike gateway first external-interface fe-2/0/0.0
```

2. Phase 2

   a. Configure Phase 2 proposals:

```
user@host# set security ipsec proposal ipsec_prop protocol esp
user@host# set security ipsec proposal ipsec_prop authentication-algorithm
  hmac-md5-96
user@host# set security ipsec proposal ipsec_prop encryption-algorithm
  3des-cbc
```

   b. Configure policies (and reference proposals):

```
user@host# set security ipsec policy ipsec_pol perfect-forward-secrecy keys
  group1
user@host# set security ipsec policy ipsec_pol proposals ipsec_prop
```

   c. Configure IPsec Autokey IKE (and reference the policy and gateway):

```
user@host# set security ipsec vpn first_vpn bind-interface st0.0
user@host# set security ipsec vpn first_vpn ike gateway gate
user@host# set security ipsec vpn first_vpn ike ipsec-policy ipsec_pol
```

3. Configure Security Policy

```
user@host# set security policies default-policy permit-all
```

4. Configure Routing Options

```
user@host# set routing-options static route 1.1.1.0/24 next-hop 7.7.7.1
```

Configuring Spoke 2

1. Phase 1

   a. Configure IKE Phase 1 proposals:

```
user@host# set security ike proposal ike_prop authentication-method
  pre-shared-keys
user@host# set security ike proposal ike_prop dh-group group2
user@host# set security ike proposal ike_prop authentication-algorithm md5
user@host# set security ike proposal ike_prop encryption-algorithm 3des-cbc
```

   b. Configure IKE policies (and reference the proposals):

```
user@host# set security ike policy ike_pol mode main
user@host# set security ike policy ike_pol proposals ike_prop
user@host# set security ike policy ike_pol pre-shared-key ascii-text
  "$9$JrUi.QF/0BEP5BEcyW8ZUj"
user@host# set security ike gateway gate ike-policy ike_pol
```

   c. Configure IKE gateway (and reference the policy):

```
user@host# set security ike gateway third address 2.2.2.2
```

```
user@host# set security ike gateway third external-interface ge-0/0/3.0
```

2. Phase 2

   a. Configure Phase 2 proposals:

```
user@host# set security ipsec proposal ipsec_prop protocol esp
user@host# set security ipsec proposal ipsec_prop authentication-algorithm
  hmac-md5-96
user@host# set security ipsec proposal ipsec_prop encryption-algorithm
  3des-cbc
```

   b. Configure policies (and reference proposals):

```
user@host# set security ipsec policy ipsec_pol perfect-forward-secrecy keys
  group1
user@host# set security ipsec policy ipsec_pol proposals ipsec_prop
```

   c. Configure IPsec Autokey IKE (and reference the policy and gateway):

```
user@host# set security ipsec vpn first_vpn bind-interface st0.0
user@host# set security ipsec vpn first_vpn ike gateway gate
user@host# set security ipsec vpn first_vpn ike ipsec-policy ipsec_pol
```

3. Configure Security Policy

```
user@host# set security policies default-policy permit-all
```

4. Configure Routing Options

```
user@host# set routing-options static route3.1.1.0/24 next-hop 7.7.7.1
```

Chapter 12

Public Key Cryptography for Certificates

The public-private key pairs used in public key cryptography play an important role in the use of digital certificates. The procedure for signing a certificate by a certificate authority (CA) and then verifying the signature, by the recipient, works as described in the topics in this chapter. The topics in this chapter also describe how to use self-signed certificates either automatically generated by the system or ones that you manually create.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding Public Key Cryptography on page 402
- Understanding Certificates on page 402
- Understanding Certificate Revocation Lists on page 405
- Understanding Public Key Infrastructure on page 405
- Understanding Self-Signed Certificates on page 408
- Understanding Automatically Generated Self-Signed Certificates on page 409
- Understanding Manually Generated Self-Signed Certificates on page 410
- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Configuring a Certificate Authority Profile on page 413
- Enrolling a CA Certificate Online on page 415
- Enrolling a Local Certificate Online on page 416
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Manually Loading a CRL onto the Device on page 423
- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425
- Using Automatically Generated Self-Signed Certificates on page 427
- Manually Generating Self-Signed Certificates on page 428

- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Understanding Public Key Cryptography

In public key cryptography, a public-private key pair is used to encrypt and decrypt data. Data encrypted with a public key, which the owner makes available to the public, can be decrypted with the corresponding private key only, which the owner keeps secret and protected. For example, if Alice wants to send Bob an encrypted message, Alice can encrypt it with Bob's public key and send it to him. Bob then decrypts the message with his private key.

The reverse process is also useful: encrypting data with a private key and decrypting it with the corresponding public key. This process is known as creating a digital signature. For example, if Alice wants to present her identity as the sender of a message, she can encrypt the message with her private key and send the message to Bob. Bob then decrypts the message with Alice's public key, thus verifying that Alice is indeed the sender.

Related Topics

- Understanding Certificates on page 402
- Understanding Self-Signed Certificates on page 408
- Understanding Certificate Revocation Lists on page 405
- Understanding Public Key Infrastructure on page 405
- Using Digital Certificates on page 410
- Using Automatically Generated Self-Signed Certificates on page 427
- Manually Generating Self-Signed Certificates on page 428

Understanding Certificates

A digital certificate is an electronic means for verifying your identity through a trusted third party, known as a certificate authority (CA). Alternatively, you can use a self-signed certificate to attest to your identity. For details on self-signed certificates, see “Understanding Self-Signed Certificates” on page 408.

Before You Begin
For background information, read “Understanding Public Key Cryptography” on page 402.

The CA server you use can be owned and operated by an independent CA or by your own organization, in which case you become your own CA. If you use an independent CA, you must contact them for the addresses of their CA and certificate revocation list (CRL) servers (for obtaining certificates and certificate revocation lists) and for

the information they require when submitting personal certificate requests. When you are your own CA, you determine this information yourself.

NOTE: The following CAs are supported: Entrust, Microsoft, and Verisign.

This topic covers:

- Certificate Signatures on page 403
- Certificate Verification on page 403
- Internet Key Exchange on page 404
- Related Topics on page 404

Certificate Signatures

The certificate authority (CA) that issues a certificate uses an MD5 or SHA-1 hash algorithm to generate a digest, then " signs" the certificate by encrypting the digest with its private key. The result is a digital signature. The CA then makes the digitally signed certificate available for download to the person who requested it. Figure 53 on page 404 illustrates this process.

Certificate Verification

The recipient of the certificate generates another digest by applying the same MD5 or SHA-1 hash algorithm to the certificate file, then uses the CA's public key to decrypt the digital signature. By comparing the decrypted digest with the digest just generated, the recipient is able to confirm the integrity of the CA's signature and, by extension, the integrity of the accompanying certificate. Figure 53 on page 404 illustrates this process.

NOTE: If the issuer of the end-entity (EE) certificate is not a root certificate, up to eight levels are verified (as explained in "Understanding Public Key Infrastructure" on page 405). Revocation status of each certificate in the verification chain is also verified. A certificate revocation status is considered " good" when its serial number is not in the CRL, which satisfies the refresh requirement per CA profile.

Figure 53: Digital Signature Verification

1. Using the MD5 or SHA-1 has algorithm, the CA makes digest A from the certificate.
2. Using the private key, the CA encrypts digest A. The result is digest B, the digital signature.
3. The CA sends the digitally signed certificate to the person who requested it.

1. Using either MD5 or SHA-1, the reciepient makes digest A from the certificate.
2. Use the CA's public key, the reciepient decrypts digest B.
3. The recipient compares digest A with digest B. If they match, the reciepient knows that the certificate has not been tampered with.

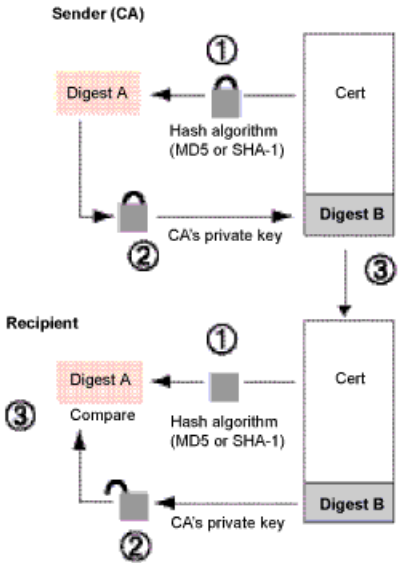

Internet Key Exchange

The procedure for digitally signing messages sent between two participants in an Internet Key Exchange (IKE) session is similar to digital certificate verification, with the following differences:

- Instead of making a digest from the CA certificate, the sender makes it from the data in the IP packet payload.
- Instead of using the CA's public-private key pair, the participants use the sender's public-private key pair.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Verifying Certificate Validity on page 424
- Understanding Certificate Revocation Lists on page 405

Understanding Certificate Revocation Lists

In the normal course of business, certificates are revoked for various reasons. You might wish to revoke a certificate if you suspect that it has been compromised, for example, or when a certificate holder leaves the company.

Before You Begin

For background information, read

- Understanding Public Key Cryptography on page 402
- Understanding Certificates on page 402

You can manage certificate revocations and validations in two ways:

- Locally— a limited solution.
- By referencing a CA's certificate revocation list (CRL). You can automatically access the CRL online at intervals you specify or at the default interval set by the CA.

In Phase 1 negotiations, participants check the CRL list to see if certificates received during an IKE exchange are still valid. If a CRL did not accompany a CA certificate and is not loaded on the device, the device tries to download it automatically from the CRL distribution point of the local certificate. If the device fails to connect to the URL in the certificate distribution point (CDP), it tries to retrieve the CRL from the URL configured in the CA profile.

If the certificate does not contain a certificate distribution point extension, and you cannot automatically retrieve the CRL through LDAP or HTTP, you can retrieve a CRL manually and load that in the device.

Related Topics

- Checking Certificate Validity Using CRLs on page 425
- Deleting a Loaded CRL on page 431
- Understanding Public Key Infrastructure on page 405

Understanding Public Key Infrastructure

Public key infrastructure (PKI) refers to the hierarchical structure of trust required for the successful implementation of public key cryptography. To verify the trustworthiness of a certificate, you must be able to track a path of certified CAs from

the one issuing your local certificate back to a root authority of a CA domain. See Figure 54 on page 406.

Before You Begin

For background information, read

- Understanding Public Key Cryptography on page 402
- Understanding Certificates on page 402
- Understanding Certificate Revocation Lists on page 405

This topic covers:

- PKI Hierarchy for a Single CA Domain or Across Domains on page 406
- PKI Management and Implementation on page 407
- Related Topics on page 407

PKI Hierarchy for a Single CA Domain or Across Domains

Figure 54 on page 406 shows the structure of a single-domain certificate authority.

Figure 54: PKI Hierarchy of Trust—CA Domain

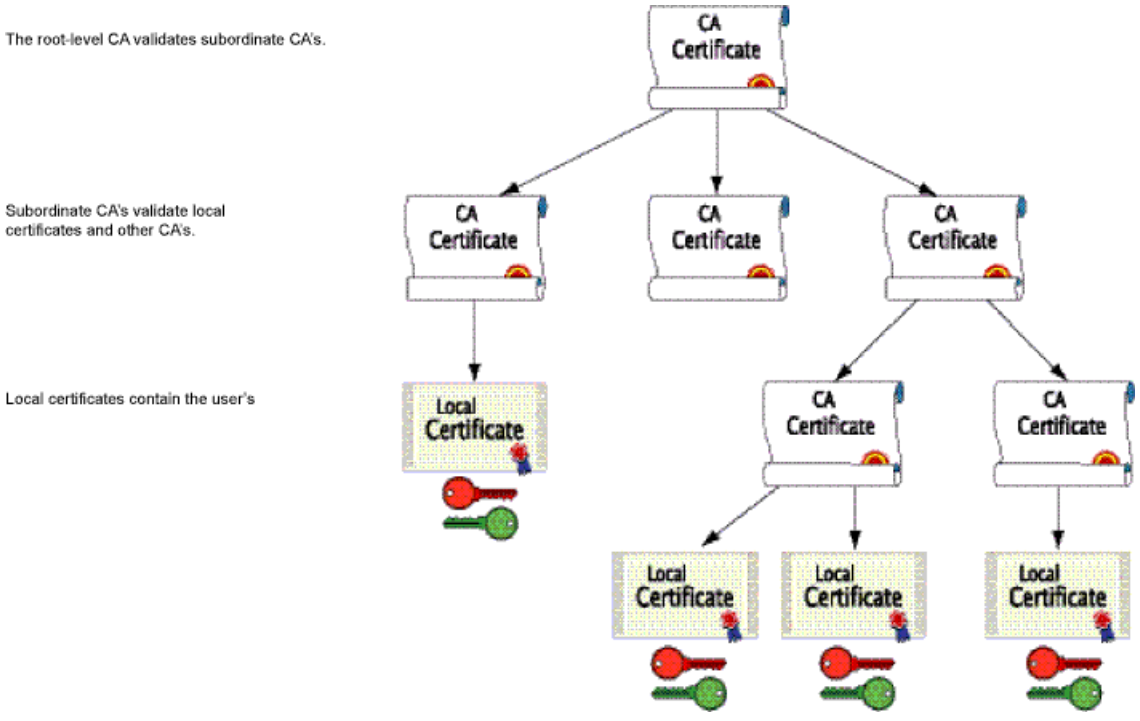

If certificates are used solely within an organization, that organization can have its own CA domain within which a company CA issues and validates certificates for its employees. If that organization later wants its employees to exchange their certificates

with those from another CA domain (for example, with employees at another organization that also has its own CA domain), the two CAs can develop cross-certification by agreeing to trust the authority of each other. In this case, the PKI structure does not extend vertically but does extend horizontally. See Figure 55 on page 407.

Figure 55: Cross-Certification

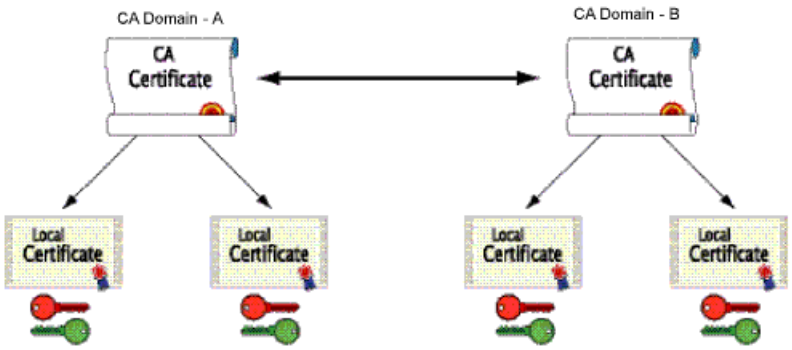

Users in the CA domain A can use their certificates and key pairs with users in CA domain B because the CA's have cross-certified each other.

PKI Management and Implementation

For convenience and practicality, PKI must be transparently managed and implemented. Toward this goal, JUNOS Software supports the following features:

- Generates a public-private key pair.
- Loads multiple local certificates from different CAs.
- Delivers a certificate when establishing an IPsec tunnel.
- Validates a certificate path upward through eight levels of CA authorities in the PKI hierarchy.
- Supports the Public-Key Cryptography #7 (PKCS-7) cryptographic standard. As a result, the device can accept X.509 certificates and CRLs packaged within a PKCS-7 envelope.

NOTE: JUNOS Software supports a PKCS-7 file size of up to 7 KB.

- Retrieves CRLs online retrieval through LDAP or HTTP.

Related Topics

- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418

- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Deleting a Loaded CRL on page 431

Understanding Self-Signed Certificates

A self-signed certificate is a certificate that is signed by its creator rather than by a Certificate Authority (CA).

Self-signed certificates allow for use of SSL-based (Secure Sockets Layer) services without requiring that the user or administrator undertake the considerable task of obtaining an identity certificate signed by a CA

Before You Begin

For background information, read

- Understanding Public Key Cryptography on page 402
- Understanding Certificates on page 402

NOTE: Self-signed certificates do not provide additional security as do those generated by CAs. This is because a client cannot verify that the server he or she has connected to is the one advertised in the certificate.

This topic covers:

- About Generating Self-Signed Certificates on page 408
- Related Topics on page 409

About Generating Self-Signed Certificates

JUNOS Software provides two methods for generating a self-signed certificate:

- Automatic generation

  In this case, the creator of the certificate is the Juniper Networks device. An automatically generated self-signed certificate is configured on the device by default.

  After the device is initialized, it checks for the presence of an automatically generated self-signed certificate. If it does not find one, the device generates one and saves it in the file system.

- Manual generation

  In this case, you create the self-signed certificate for the device.

At any time, you can use the CLI to generate a self-signed certificate. These certificates are also used to gain access to SSL services.

Self-signed certificates are valid for five years from the time they were generated.

Related Topics

- Understanding Automatically Generated Self-Signed Certificates on page 409
- Understanding Manually Generated Self-Signed Certificates on page 410
- Using Automatically Generated Self-Signed Certificates on page 427
- Manually Generating Self-Signed Certificates on page 428
- Deleting Certificates on page 430
- Understanding Public Key Cryptography on page 402

Understanding Automatically Generated Self-Signed Certificates

An automatically generated self-signed certificate allows for use of SSL-based (Secure Sockets Layer) services without requiring that the administrator obtain an identity certificate signed by a CA.

A self-signed certificate that is automatically generated by the device is similar to an SSH (Secure Shell) host key. It is stored in the file system, not as part of the configuration. It persists when the device is rebooted, and it is preserved when a request system snapshot command is issued

Before You Begin

For background information, read:

- Understanding Self-Signed Certificates on page 408
- Understanding Certificates on page 402

Related Topics

- Using Automatically Generated Self-Signed Certificates on page 427
- Understanding Manually Generated Self-Signed Certificates on page 410
- Manually Generating Self-Signed Certificates on page 428
- Deleting Certificates on page 430
- Understanding Public Key Cryptography on page 402

Understanding Manually Generated Self-Signed Certificates

A self-signed certificate that you manually generate allows for use of SSL-based (Secure Sockets Layer) services without requiring that you obtain an identity certificate signed by a CA. A manually generated self-signed certificate is one example of a PKI local certificate. As is true of all PKI local certificates, manually generated self-signed certificates are stored in the file system.

Before You Begin

For background information, read:

- Understanding Self-Signed Certificates on page 408
- Understanding Certificates on page 402
- Understanding Public Key Cryptography on page 402

You generate this kind of named local certificate using the command-line interface (CLI).

Related Topics

- Manually Generating Self-Signed Certificates on page 428
- Understanding Automatically Generated Self-Signed Certificates on page 409
- Using Automatically Generated Self-Signed Certificates on page 427
- Deleting Certificates on page 430

Using Digital Certificates

Digital certificates authenticate your identity when establishing secure virtual private network (VPN) connections.

To use a digital certificate to authenticate your identity when establishing a secure VPN connection, you must first do the following:

- Obtain a certificate authority (CA) certificate from which you intend to obtain a personal certificate, and then load the CA certificate in the device.

  The CA certificate can contain a certificate revocation list (CRL) to identify invalid certificates.

- Obtain a local certificate (also known as a personal certificate) from the CA whose CA certificate you have previously loaded, and then load the local certificate in the device. The local, or end-entity (EE), certificate establishes the identity of the Juniper Networks device with each tunnel connection.

You can obtain CA and local certificates manually, or online using Simple Certificate Enrollment Protocol (SCEP). Certificates are verifiable and renewable, and you can delete them when they are no longer needed.

Before You Begin

For background information, read

- Understanding Public Key Cryptography on page 402
- Understanding Certificates on page 402
- Understanding Certificate Revocation Lists on page 405
- Understanding Public Key Infrastructure on page 405

This topic covers:

- Obtaining Digital Certificates Online on page 411
- Obtaining Digital Certificates Manually on page 411
- Verifying the Validity of a Certificate on page 412
- Deleting a Certificate on page 412

Obtaining Digital Certificates Online

Simple Certificate Enrollment Protocol (SCEP) uses the online method to request digital certificates. To obtain a certificate online, do the following:

1. Generate a key pair in the device. See "Generating a Public-Private Key Pair" on page 412.
2. Create a CA profile containing information specific to a CA. You can have multiple CA profiles on the device. For example, you might have one profile for Microsoft and one for Entrust. See "Configuring a Certificate Authority Profile" on page 413.
3. Enroll the CA certificate onto the device. See "Enrolling a CA Certificate Online" on page 415.
4. Obtain a local certificate (also known as a personal certificate) online from the CA whose CA certificate you have previously loaded. See "Enrolling a Local Certificate Online" on page 416.
5. Configure automatic re-enrollment. See "Understanding SecurID User Authentication" on page 326.

Obtaining Digital Certificates Manually

To obtain digital certificates manually, do the following:

1. Generate a key pair in the device. See "Generating a Public-Private Key Pair" on page 412.
2. Create a CA profile containing information specific to a CA. You can have multiple CA profiles on the device. For example, you might have one profile for Microsoft

and one for Entrust. See “Configuring a Certificate Authority Profile” on page 413.

3. Generate a certificate request using the key pair, and manually copy that request and paste it into the appropriate field at the CA website to obtain a personal certificate (also known as a local certificate). See “Generating a Local Certificate Request Manually” on page 418.
4. Load the certificate onto the device. See “Loading CA and Local Certificates Manually” on page 420.
5. Configure automatic re-enrollment. See “Understanding SecurID User Authentication” on page 326.
6. If necessary, load the certificate's CRL on the device. See “Manually Loading a CRL onto the Device” on page 423.

Verifying the Validity of a Certificate

You can verify the validity of a certificate in one of the following ways:

- To verify manually, see “Verifying Certificate Validity” on page 424.
- To verify manually with CRLs, see “Checking Certificate Validity Using CRLs” on page 425.

Deleting a Certificate

To delete a certificate or a CRL, see “ Deleting Certificates” on page 177 and “ Deleting a Loaded CRL” on page 179.

Generating a Public-Private Key Pair

When you generate a public-private key pair, the device automatically saves the key pair in a file in the certificate store, where it is subsequently used in certificate request commands.

Before You Begin
For background information, read ■ Understanding Public Key Cryptography on page 402 ■ Understanding Certificates on page 402 ■ Understanding Certificate Revocation Lists on page 405 ■ Using Digital Certificates on page 410

If the device renews a great number of certificates at once, thus using up keys rapidly, it might run out of pregenerated keys and have to generate them promptly for each new request. In this case, the generation of keys might affect the performance of the device, especially in a high-availability environment where the performance of the device might slow down for a number of minutes.

You must have root-level privileges to generate a public-private key pair. When you generate a public-private key pair on the device, the generated key pair is saved as *certificate-id*.priv.

This topic covers:

- CLI Operation on page 413
- Related Topics on page 413

CLI Operation

1. To generate a public-private key pair named, for example, ca-ipsec, with a key size of 1024 bits, enter the following command:

   user@host> **request security pki generate-key-pair certificate-id ca-ipsec**

NOTE: The default RSA key size is 1024 bits. If you are using SCEP, JUNOS Software supports RSA only.

2. Go on to "Configuring a Certificate Authority Profile" on page 413.

Related Topics

- Using Digital Certificates on page 410
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Configuring a Certificate Authority Profile

A certificate authority (CA) profile configuration contains information specific to a CA. You can have multiple CA profiles on the device. For example, you might have one profile for Microsoft and one for Entrust. Each profile is associated with a CA certificate. If you want to load a new CA certificate without removing the older one, you must create a new CA profile (for example, Microsoft-2008).

Before You Begin
For background information, read ■ Public Key Cryptography for Certificates on page 401 ■ Understanding Certificates on page 402 ■ Understanding Certificate Revocation Lists on page 405 ■ Using Digital Certificates on page 410

This topic covers:

- CLI Configuration on page 414
- Related Topics on page 414

CLI Configuration

To configure a CA profile:

1. Create a CA profile. For example, the following command creates a CA profile called ca-profile-ipsec with CA identity microsoft-2008, specifies that the CRL be refreshed every 48 hours, location to retrieve the CRL from to http://www.my-ca.com:

   user@host# **set security pki ca-profile ca-profile-ipsec ca-identity microsoft-2008 revocation-check crl refresh-interval 48 url http://www.my-ca.com**

2. Specify the number of times a device resends a certificate request for online enrollment when attempts to enroll in Step 1 fail. For example, the following command sets the enrollment retry to 20 times:

   user@host# **set security pki ca-profile ca-profile-ipsec enrollment retry 20**

   The default value for retry is 10.

3. Specify the time interval in seconds between attempts to automatically enroll the CA certificate online. For example, the following command specifies automatic certificate polling every 30 minutes:

   user@host# **set security pki ca-profile ca-profile-ipsec enrollment retry-interval 1800**

   If you configure retry only without configuring a retry interval, then the default retry interval is 900 seconds (15 minutes). If you do not configure retry or a retry interval, then there is no polling.

Related Topics

- Generating a Public-Private Key Pair on page 412
- Enrolling a CA Certificate Online on page 415
- Enrolling a Local Certificate Online on page 416
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Enrolling a CA Certificate Online

With Simple Certificate Enrollment Protocol (SCEP), you can configure your Juniper Networks device to obtain a CA certificate online and start the online enrollment for the specified certificate ID. The CA public key verifies certificates from remote peers.

Before You Begin

1. Generate a public and private key pair. See "Generating a Public-Private Key Pair" on page 412.
2. Configure a CA profile. See "Configuring a Certificate Authority Profile" on page 413.
3. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

This topic covers:

- CLI Operation on page 415
- Related Topics on page 416

CLI Operation

1. Use the following command to get the CA certificate online using SCEP. The attributes required to reach the CA server are obtained from the defined CA profile.

   user@host> **request security pki ca-certificate enroll ca-profile** *ca-profile_name*

   The command is processed synchronously to provide the fingerprint of the received CA certificate as shown below:

   ```
   Fingerprint:
   e6:fa:d6:da:e8:8d:d3:00:e8:59:12:e1:2c:b9:3c:c0:9d:6c:8f:8d (sha1)
   82:e2:dc:ea:48:4c:08:9a:fd:b5:24:b0:db:c3:ba:59 (md5)
   Do you want to load the above CA certificate ? [yes,no]
   ```

   You must confirm that the correct certificate is loaded. The CA certificate is loaded only when you type yes at the CLI prompt. For more information on the certificate, such as the bit length of the key pair, use the command show security pki ca-certificate described in the *JUNOS Software CLI Reference*.

2. Go on to "Enrolling a Local Certificate Online" on page 416.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425

Enrolling a Local Certificate Online

With Simple Certificate Enrollment Protocol (SCEP), you can configure your Juniper Networks device to obtain a local certificate online and start the online enrollment for the specified certificate ID.

Before You Begin

1. Generate a public and private key pair. See "Generating a Public-Private Key Pair" on page 412.
2. Configure a CA profile. See "Configuring a Certificate Authority Profile" on page 413.
3. Enroll a CA certificate. See "Enrolling a CA Certificate Online" on page 415.
4. For background information, read:
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

This topic covers:

- CLI Configuration on page 416
- Related Topics on page 417

CLI Configuration

To configure the device for online enrollment:

1. Specify the CA profile—for example, wincs-5—and specify the CA location for your device to send the SCEP-based certificate enrollment requests. To specify the CA location by naming the CA URL, include the url statement. For example:

   user@host# **set security pki ca-profile wincs-5 enrollment url http://10.155.8.1/certsrv/mscep/mscep.dll**

2. Using the request security pki local-certificate enroll command, start the online enrollment for the specified certificate ID. You must specify the CA profile name

(for example, wincs-5), the certificate ID (for example, qqq), and the following information:

NOTE: SCEP sends a PKCS-10 format certificate request enveloped in PKCS-7 format.

- Specify the challenge CA password for certificate enrollment and revocation—for example, aaa. If the CA does not provide the challenge password, then choose your own password.
- Specify at least one of the following values:
  - Enter the domain name to identify the certificate owner in IKE negotiations—for example, qqq.juniper.net.
  - Specify the identity of the certificate owner for IKE negotiation with the e-mail statement—for example, qqq@juniper.net.
  - Enter an IP address if the Service router is configured for a static IP address—for example, 10.10.10.10.
- Specify the subject name in the distinguished name format in quotation marks, including the domain component (DC), common name (CN), organizational unit name (OU), organization name (O), locality (L), state (ST), and country (C).

  For example:

```
user@host> request security pki local-certificate enroll ca-profile wincs-5
certificate-id qqq challenge-password aaa domain-name qqq.juniper.net
email qqq@juniper.net ip-address 10.10.10.10 subject DC=juniper,
CN=router3, OU=marketing, O=juniper, L=sunnyvale, ST=california, C=us
```

The device certificate is obtained and the online enrollment begins for the certificate ID. The command is processed asynchronously.

3. Go on to "Re-enrolling Local Certificates Automatically" on page 421

The device certificate is obtained and the online enrollment begins for the certificate ID. The command is processed asynchronously.

1. Go on to "Understanding SecurID User Authentication" on page 326.

Related Topics

- Using Digital Certificates on page 410
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Verifying Certificate Validity on page 424

- Checking Certificate Validity Using CRLs on page 425
- Deleting a Loaded CRL on page 431

Generating a Local Certificate Request Manually

When you create a local certificate request, the device generates a CA certificate in PKCS-10 format from a key pair you previously generated using the same certificate ID.

Before You Begin

1. Generate a public and private key. See "Generating a Public-Private Key Pair" on page 412.
2. Create a CA profile. See "Configuring a Certificate Authority Profile" on page 413.
3. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

A subject name is associated with the local certificate request in the form of a common name (CN), organizational unit (OU), organization (O), locality (L), state (ST), country (C), and domain component (DC). Additionally, a subject alternative name is associated in the following form:

- IP address
- E-mail address
- Fully qualified domain name (FQDN)

NOTE: Some CAs do not support an e-mail address as the domain name in a certificate. If you do not include an e-mail address in the local certificate request, you cannot use an e-mail address as the local IKE ID when configuring the device as a dynamic peer. Instead, you can use a fully qualified domain name (if it is in the local certificate), or you can leave the local ID field empty. If you do not specify a local ID for a dynamic peer, enter the *hostname.domain-name* of that peer on the device at the other end of the IPsec tunnel in the peer ID field.

This topic covers:

- CLI Operation on page 419
- Related Topics on page 419

CLI Operation

1. To generate a certificate request using the certificate ID (ca-ipsec) of a public-private key pair you previously generated and specifying the domain name juniper.net and the associated common name abc, enter the following command:

   user@host> **request security pki generate-certificate-request certificate-id ca-ipsec domain-name juniper.net subject CN=abc**

   The following certificate request is displayed in PEM format.

   ```
   Generated certificate request
   -----BEGIN CERTIFICATE REQUEST-----
   MIHxMIGcAgEAMA4xDDAKBgNVBAMTA2htMTBcMA0GCSqGSIb3DQEBAQUAA0sAMEgC
   QQCbhaiWzmctH0ZDldCn+mSNM62kyiSgc4cmN68U/j9El09/DgGoMNy2y+RYA1xU
   sr4B0NedGrZZJx5L1sIYjHr/AgMBAAGgKTAnBgkqhkiG9w0BCQ4xGjAYMBYGA1Ud
   EQQPMA2CC2p1bmlwZXIubmV0MA0GCSqGSIb3DQEBBQUAA0EAleLR6Hp2ity8Dugs
   MW4HI6SxfwMc2eYM5Nj2UhwpEEpsce77dUBZriKdehAgli7vwNsHGIuhHjEaFzfO
   hpM3tA==
   -----END CERTIFICATE REQUEST-----
   Fingerprint:
   9e:d5:7d:44:e8:e7:b6:d7:4b:58:d4:4e:2b:fb:c6:b2:4b:b7:8b:82 (sha1)
   b0:8d:c7:6d:41:d5:58:61:dc:a0:3e:4e:d6:39:02:d7 (md5)
   ```

   Copy the generated certificate request and paste it into the appropriate field at the CA website to obtain a local certificate. Refer to the CA server documentation to determine where to paste the certificate-request.

   When PKCS-10 content is displayed, the SHA-1 hash and MD5 hash of the PKCS-10 file is also displayed. For more information on the certificate, such as the bit length of the key pair, use the command show security pki certificate-request described in the *JUNOS Software CLI Reference*.

2. Go on to "Loading CA and Local Certificates Manually" on page 420.

Related Topics

- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Loading CA and Local Certificates Manually

After you download certificates from a certificate authority (CA), you transfer them to the device (for example, using FTP), then load them.

Before You Begin

1. Generate a public-private key pair. See "Generating a Public-Private Key Pair" on page 412.
2. Create a CA profile. See "Configuring a Certificate Authority Profile" on page 413.
3. Generate a certificate request. See "Generating a Local Certificate Request Manually" on page 418 .
4. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

You can load the following certificates files onto a device running JUNOS Software:

- A local or end-entity (EE) certificate that identifies your local device. This certificate is your public key.
- A CA certificate that contains the CA's public key.
- A CRL that lists any certificates revoked by the CA.

NOTE: You can load multiple EE certificates onto the device.

This topic covers:

- CLI Operation on page 420
- Related Topics on page 421

CLI Operation

In this example, you have downloaded the following certificates and saved them to the /var/tmp/ directory on the device:

- local.cert
- ca.cert

1. To load the local certificate called local.cert from the /var/tmp directory on the device, enter the following command:

```
user@host> request security pki local-certificate load certificate-id local.cert
  filename /var/tmp/local.cert
```

2. To load the CA certificate called ca.cert from the /var/tmp directory on the device, enter the following command. The CA profile is called ca-profile-ipsec.

```
user@host> request security pki ca-certificate load ca-profile ca-profile-ipsec
  filename /var/tmp/ca.cert
```

3. Go on to "Re-enrolling Local Certificates Automatically" on page 421.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Deleting Certificates on page 430
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425
- Deleting a Loaded CRL on page 431

Re-enrolling Local Certificates Automatically

You can enable the device to automatically renew certificates that was acquired by online enrollment or loaded manually. This feature saves you from having to remember to renew certificates on the device before they expire, and helps maintain valid certificates at all times.

Automatic certificate renewal is disabled by default. You can configure the device to automatically send out a request to renew a certificate before it expires. You can set the time when you want the device to send out the certificate renewal request in number of days and minutes before the expiration date. By setting different times for each certificate, you prevent the device from having to renew all certificates at the same time.

Before You Begin

1. Obtain a certificate either online or manually. See "Obtaining Digital Certificates Online" on page 411 or "Obtaining Digital Certificates Manually" on page 411.
2. For background information, read:
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

For this feature to work, the device must be able to reach the SCEP server, and the certificate must be present on the device during the renewal process. Furthermore,

for this feature to work, you must also ensure that the CA issuing the certificate can return the same DN (domain name). The CA must not modify the subject name and Alternate Subject Name extension in the new certificate.

You can enable and disable automatic SCEP certificate renewal for all SCEP certificates or on a per-certificate basis.

This topic covers:

- CLI Configuration on page 422
- Related Topics on page 422

CLI Configuration

To enable and configure certificate re-enrollment use the set security pki auto-re-enrollment command with the following information:

- Certificate ID of the CA certificate—for example, sm1.
- Name of the CA profile associated with the certificate—for example, aaa.
- Challenge password for CA certificate enrollment and revocation. This password must be the same one configured previously for the CA—for example, abc.
- Trigger time for the re-enrollment. This value sets the certificate re-enrollment time as a percentage of the time left before expiration. For example, to start re-enrollment when 10 percent of the certificate time remains, specify 10.
- During automatic re-enrollment, by default the Juniper Networks device uses the existing key pair. To generate a new key pair, specify re-generate-key-pair.

For example:

```
user@host# set security pki auto-re-enrollment certificate-id sm1 ca-profile-name
    aaa challenge-password abc re-enroll-trigger-time-percentage 10
    re-generate-keypair
```

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Manually Loading a CRL onto the Device

You can load a certificate revocation list (CRL) manually, or you can have the device load it automatically when you verify certificate validity. To load a CRL manually, you obtain the CRL from a CA and transfer it to the device (for example, using FTP).

Before You Begin

1. Generate a public and private key pair. See "Generating a Public-Private Key Pair" on page 412.
2. Generate a certificate request. See "Generating a Local Certificate Request Manually" on page 418.
3. Configure a CA profile. See "Configuring a Certificate Authority Profile" on page 413.
4. Load your certificate onto the device. See "Loading CA and Local Certificates Manually" on page 420.
5. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

This topic covers:

- CLI Operation on page 423
- Related Topics on page 423

CLI Operation

With the following command, you load a CRL certificate called revoke.crl from the /var/tmp directory on the device. The CA profile is called ca-profile-ipsec. (Maximum file size is 5 MB.)

```
user@host> request security pki crl load ca-profile ca-profile-ipsec filename
  /var/tmp/revoke.crl
```

NOTE: JUNOS Software supports loading of CA certificates in X509, PKCS-7, DER, or PEM formats.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420

- Verifying Certificate Validity on page 424
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Verifying Certificate Validity

The CRL is updated automatically, but you must verify certificates manually to find out if a certificate has been revoked, or if the CA certificate used to create a local certificate is no longer present on the device.

When you verify certificates manually, the device uses the CA certificate to verify the local certificate. If the local certificate is valid, and if revocation-check is enabled in the CA profile, the device verifies that the CRL is loaded and valid. If not, the device downloads the new CRL.

Before You Begin

1. Obtain a certificate either online or manually. See "Obtaining Digital Certificates Online" on page 411 or "Obtaining Digital Certificates Manually" on page 411.
2. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

You verify certificates from the CLI in operational mode.

This topic covers:

- CLI Operation on page 424
- Related Topics on page 425

CLI Operation

Use the following command to verify the validity of a local certificate called local.cert:

```
user@host> request security pki local-certificate verify certificate-id local.cert
```

Use the following command to verify the validity of a CA certificate called ca-cert:

```
user@host> request security pki ca-certificate verify certificate-id ca-cert
```

NOTE: The associated private key and the signature are also verified.

For more information on the certificate, use the show commands (**show security pki ca-certificate** and **show security pki certificate-request**) described in the *JUNOS Software CLI Reference*.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Checking Certificate Validity Using CRLs

In Phase 1 negotiations, participants check the CRL list to see if certificates received during an IKE exchange are still valid. If a CRL did not accompany a CA certificate and is not loaded on the device, JUNOS Software tries to retrieve the CRL through the LDAP or HTTP CRL location defined within the CA certificate itself. If no URL address is defined in the CA certificate, the device uses the URL of the server that you define for that CA certificate. If you do not define a CRL URL for a particular CA certificate, the device gets the CRL from the URL in the CA profile configuration.

Before You Begin

1. Obtain a certificate either online or manually. See "Obtaining Digital Certificates Online" on page 411 or "Obtaining Digital Certificates Manually" on page 411.
2. For background information, read:
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

NOTE: The CRL distribution point extension (.cdp) in an X509 certificate can be added to either an HTTP URL or an LDAP URL.

This topic covers:

- J-Web Configuration on page 426
- CLI Configuration on page 426
- Related Topics on page 426

J-Web Configuration

To configure a certificate authority profile.

1. Select Configure>CLI Tools>Point and Click CLI
2. Next to Security, click **Configure** or **Edit**.
3. Next to PKI, select the check box and click Configure.
4. Next to Ca profile, click Add new entry.
5. In the Ca profile name box, type **my_profile**.
6. In the Ca identity box, type **sm1**.
7. Next to Revocation check, click Configure.
8. Next to Crl, click Configure.
9. Next to Url, click Add new entry.
10. In the Url string, type **http://abc** and click OK.
11. If you are finished configuring the router, Commit the configuration.
12. To check the configuration, see Verifying the Validity of a Certificate“Verifying the Validity of a Certificate” on page 412

CLI Configuration

With the following command, you direct the device to check the validity of the CA profile called my_profile and, if a CRL did not accompany a CA certificate and is not loaded on the device, to retrieve the CRL from the URL http://abc.

```
user@host# set security pki ca-profile my_profile revocation-check crl url http://abc
```

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Verifying Certificate Validity on page 424
- Deleting Certificates on page 430
- Deleting a Loaded CRL on page 431

Using Automatically Generated Self-Signed Certificates

After the device is initialized, it checks for the presence of a self-signed certificate. If a self-signed certificate is not present, the device automatically generates one.

Before You Begin

For background information, read:

- Understanding Self-Signed Certificates on page 408
- Understanding Manually Generated Self-Signed Certificates on page 410
- Understanding Certificates on page 402

This topic covers:

- J-Web Configuration on page 427
- CLI Configuration on page 427
- Related Topics on page 428

J-Web Configuration

To use the automatically generated self-signed certificate. The following tasks specify that the automatically generated self-signed certificate is used for Web management HTTPS services.

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to System, click **Configure** or **Edit**.
3. Next to Services, select the check box and click **Configure** or **Edit**.
4. Next to Web management, click **Configure** or **Edit**.
5. Next to Https, select the check box and click **Configure** or **Edit**.
6. From the Certificate choice list, select System generated certificate and click OK.
7. If you are finished configuring the router, commit the configuration.
8. To check the configuration, see Verifying the Validity of a Certificate“Verifying the Validity of a Certificate” on page 412.

CLI Configuration

You can add the following statement to your configuration if you want to use the automatically generated self-signed certificate to provide access to HTTPS services:

```
system {
 services {
  web-management {
   http {
    interface [ ... ];
```

```
   } https {
    system-generated-certificate;
    interface [ ... ];
   }
  }
 }
}
```

The device uses the following distinguished name for the automatically generated certificate:

" CN=<device serial number>, CN=system generated, CN=self-signed"

Use the following command to specify that the automatically generated self-signed certificate is to be used for Web management HTTPS services:

user@**host# set system services web-management https system-generated-certificate**

Use the following operational command to delete the automatically generated self-signed certificate:

user@host# clear security pki local-certificate system-generated

After you delete the system-generated self-signed certificate, the device automatically generates a new one and saves it in the file system.

Related Topics

- Understanding Manually Generated Self-Signed Certificates on page 410
- Manually Generating Self-Signed Certificates on page 428
- Understanding Public Key Infrastructure on page 405

Manually Generating Self-Signed Certificates

You can use the CLI to manually generate a self-signed certificate. For a manually generated self-signed certificate, you specify the distinguished name (DN) when you create it. (For an automatically generated self-signed certificate, the system supplies the DN, identifying itself as the creator.)

Before You Begin

For background information, read:

- Understanding Self-Signed Certificates on page 408
- Understanding Automatically Generated Self-Signed Certificates on page 409
- Understanding Certificates on page 402

This topic covers:

- J-Web Configuration on page 429
- CLI Configuration on page 429
- Related Topics on page 429

J-Web Configuration

To direct the router to use a manually generated self-signed certificate. The following tasks are used to direct the router to use a manually generated self-signed certificate called self-cert for Web management.

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to System, click **Configure** or **Edit**.
3. Next to Services, select the check box and click **Configure** or **Edit**.
4. Next to Web management, click **Configure** or **Edit**.
5. Next to Https, select the check box and click **Configure** or **Edit**.
6. From the Certificate choice list, select Pki local certificate and click OK.
7. If you are finished configuring the router, commit the configuration.
8. To check the configuration, see Verifying the Validity of a Certificate“Verifying the Validity of a Certificate” on page 412.

CLI Configuration

Use the following CLI command to manually generate a self-signed certificate created and signed by the user whose e-mail address is mholmes:

```
user@host# request security pki local-certificate generate-self-signed certificate-id
  self-cert subject cn=abc domain-name Juniper.net ip-address 1.2.3.4 email
  mholmes@juniper.net
```

Use the following CLI command to direct the device to use a manually generated self-signed certificate called self-cert for Web management:

```
user@host# set system services web-management https pki-local-certificate self-cert
```

Related Topics

- Understanding Manually Generated Self-Signed Certificates on page 410
- Manually Generating Self-Signed Certificates on page 428
- Understanding Public Key Infrastructure on page 405

Deleting Certificates

You can delete a local or trusted CA certificate that is automatically or manually generated.

Before You Begin

1. Obtain a certificate either online or manually. See "Obtaining Digital Certificates Online" on page 411 or "Obtaining Digital Certificates Manually" on page 411.
2. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

This topic covers:

- CLI Operation on page 430
- Related Topics on page 431

CLI Operation

Use the following command to delete a local certificate:

user@host> **clear security pki local certificate certificate-id** (*certificate-id* | **all** | **system-generated**)

Specify a certificate ID to delete a local certificate with a specific ID, use **all** to delete all local certificates, or specify system-generated to delete the automatically generated self-signed certificate.

When you delete an automatically generated self-signed certificate, the device generates a new one.

To delete a CA certificate:

user@host> **clear security pki ca-certificate ca-profile** (*ca-profile-name* **| all**)

Specify a CA profile to delete a specific CA certificate, or use **all** to delete all CA certificates present in the persistent store.

NOTE: You are asked for confirmation before a CA certificate can be deleted.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Checking Certificate Validity Using CRLs on page 425
- Deleting a Loaded CRL on page 431
- Understanding Automatically Generated Self-Signed Certificates on page 409

Deleting a Loaded CRL

You can choose to delete a loaded CRL if you no longer need to use it to manage certificate revocations and validation.

Before You Begin

1. Obtain a certificate either online or manually. See "Obtaining Digital Certificates Online" on page 411 or "Obtaining Digital Certificates Manually" on page 411.
2. For background information, read
   - Public Key Cryptography for Certificates on page 401
   - Understanding Certificates on page 402
   - Understanding Certificate Revocation Lists on page 405
   - Using Digital Certificates on page 410

This topic covers:

- CLI Operation on page 431
- Related Topics on page 432

CLI Operation

Use the following command to delete a loaded certificate revocation list:

user@host> **clear security pki crl ca-profile** (ca-profile | **all**)

Specify a CA profile to delete a CRL associated with the CA identified by the profile, or use **all** to delete all CRLs.

Related Topics

- Using Digital Certificates on page 410
- Generating a Public-Private Key Pair on page 412
- Generating a Local Certificate Request Manually on page 418
- Loading CA and Local Certificates Manually on page 420
- Re-enrolling Local Certificates Automatically on page 421
- Checking Certificate Validity Using CRLs on page 425
- Deleting Certificates on page 430

Chapter 13

Dynamic Virtual Private Networks (VPNs)

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Dynamic VPN Overview on page 433
- Access Manager Client-Side Reference on page 442
- Configuring a Dynamic VPN—Overview on page 449
- Configuring an IKE Phase 1 Proposal— (Dynamic VPNs) on page 451
- Configuring an IKE Policy—(Dynamic VPNs) on page 454
- Configuring an IKE Gateway—Quick Configuration (Dynamic VPNs) on page 456
- Configuring an IPsec Phase 2 Proposal—Quick Configuration (Dynamic VPNs) on page 460
- Configuring an IPsec Policy—Quick Configuration (Dynamic VPNs) on page 463
- Configuring an IPsec Autokey—Quick Configuration (Dynamic VPNs) on page 465
- Creating a Client Configuration—Quick Configuration (Dynamic VPNs) on page 469
- Creating a Client Configuration (Dynamic VPNs) on page 472
- Configuring Global Client Download Settings—Quick Configuration (Dynamic VPNs) on page 473
- Configuring Global Client Download Settings (Dynamic VPNs) on page 475

Dynamic VPN Overview

Virtual private network (VPN) tunnels enable users to securely access assets such as e-mail servers and application servers that reside behind a firewall. End-to-site VPN tunnels are particularly helpful to remote users such as telecommuters because a single tunnel enables access to all of the resources on a network—the users do not need to configure individual access settings to each application and server.

Before You Begin

For background information, read

- "Virtual Private Networks (VPNs) Overview" on page 355.
- "Understanding IKE and IPsec Packet Processing" on page 364.

Figure 56: Using a VPN Tunnel to Enable Remote Access to a Corporate Network

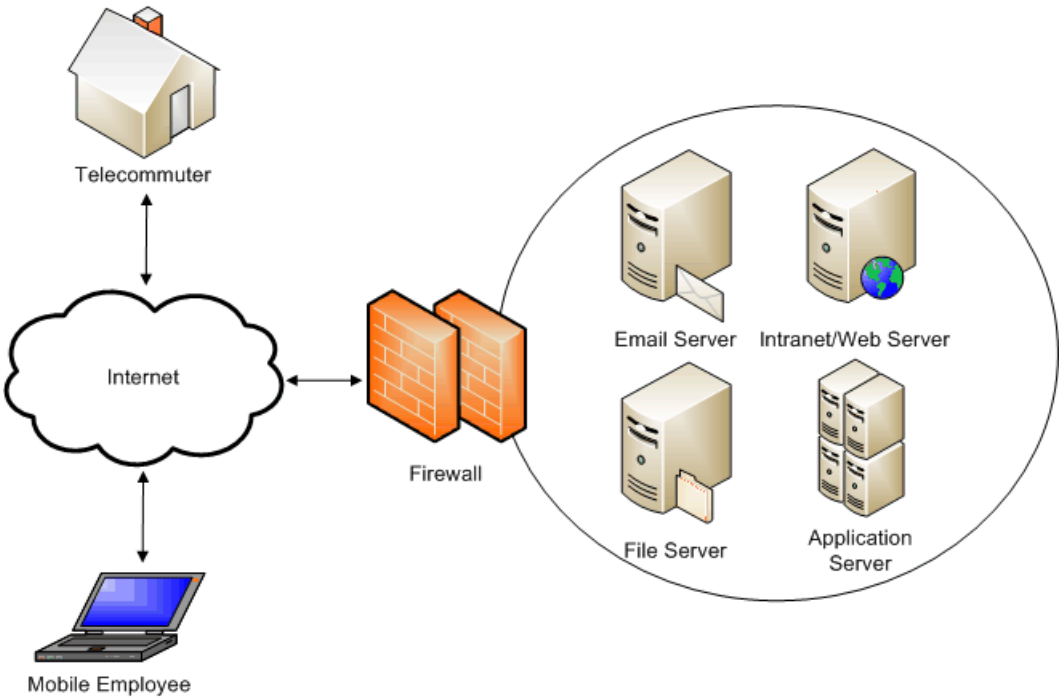

The dynamic VPN feature further simplifies remote access by enabling users to establish Internet Protocol Security (IPsec) VPN tunnels without having to manually configure VPN settings on their PCs or laptops. Instead, authenticated users can simply download the Access Manager Web client to their computers. This Layer 3 remote access client uses client-side configuration settings that it receives from the server to create and manage a secure end-to-site VPN tunnel to the server.

From the user's perspective, creating a secure VPN tunnel is very simple. The first time a user needs to establish a VPN tunnel, they simply navigate to https://<*serverhost*>/dynamic-vpn and enter their username and password in the login page that appears. Assuming that the user authenticates successfully and has administrator privileges, the Juniper Networks device (also called the Remote Access Server) installs Access Manager on the user's computer and provides a VPN configuration that is specific to the user. The Access Manager client provides the user with a simple GUI for launching the client configuration; the client configuration does all the work of establishing and negotiating the IPsec VPN tunnel for the user. Once installed, Access Manager and the client configuration are available for future IPsec VPN sessions.

The following sections describe how the Remote Access Server, Access Manager, and the client configuration work together to make the user experience simple:

- Connecting to the Remote Access Server for the First Time (Pre-IKE Phase) on page 435
- Connecting to the Remote Access Server for Subsequent Sessions (Pre-IKE Phase) on page 438

- Establishing an IPsec VPN Tunnel (IKE Phase) on page 439
- Related Topics on page 442

Connecting to the Remote Access Server for the First Time (Pre-IKE Phase)

In order to establish a secure VPN tunnel from the user's computer to the Remote Access Server, the user must first authenticate into the server and download the client-side files as follows:

1. **The user accesses the server's URL.** The user navigates to the https://<*serverhost*>/dynamic-vpn URL through a Web browser. This URL directs the user to the dynamic VPN login page on the Remote Access Server.
2. **The user signs into the server.** The user enters the appropriate username and password into the login page, and the Remote Access Server sends them to the authentication server for validation.
3. **The server retrieves a client configuration.** Once the server determines that the user has successfully authenticated, the server determines which client configuration to use when creating a secure VPN tunnel. The configuration includes an IKE ID for the user (such as johndoe.yourcompany.com), a Phase 1 security key, and a generated token to establish eligibility for future client downloads.
4. **The server downloads the setup client to the user's computer.** The server downloads the setup client (along with the client version information, client initialization parameters, and client VPN configuration parameters) to the user's computer:
   - If the user is using Internet Explorer with Active-X enabled, the Remote Access Server downloads an Active-X setup client to the user's computer.
   - Otherwise, if the user is using a Web browser with Java enabled, the Remote Access Server downloads an Java setup client to the user's computer.
   - If the user does not have Active-X or Java enabled, the server presents a download page to the user, enabling the user to manually download the setup client.
5. **The setup client checks that the user has administrator privileges.** Once the server has successfully downloaded the setup client to the user's computer, the setup client checks that the user has the proper rights to install a new client. (Administrator privileges are required only to install the client, but not to upgrade it.)
6. **The setup client installs Access Manager.** The setup client installs Access Manager on the user's computer. The user will be prompted to restart the computer to finish the installation.

Once the Access Manager client is successfully launched, the user can initiate a secure VPN connection to the Remote Access Server from Access Manager. For more information, see "Establishing an IPsec VPN Tunnel (IKE Phase)" on page 439.

NOTE:

- You can also download the latest version of the Access Manager client from the Juniper Networks Support site.
- The user can connect to the Remote Access Server and initiate a client download before you have finished configuring the dynamic VPN feature. In this case, the user can still authenticate into the server, but will not be able to establish a secure VPN tunnel.

Figure 57: Access Manager User Login

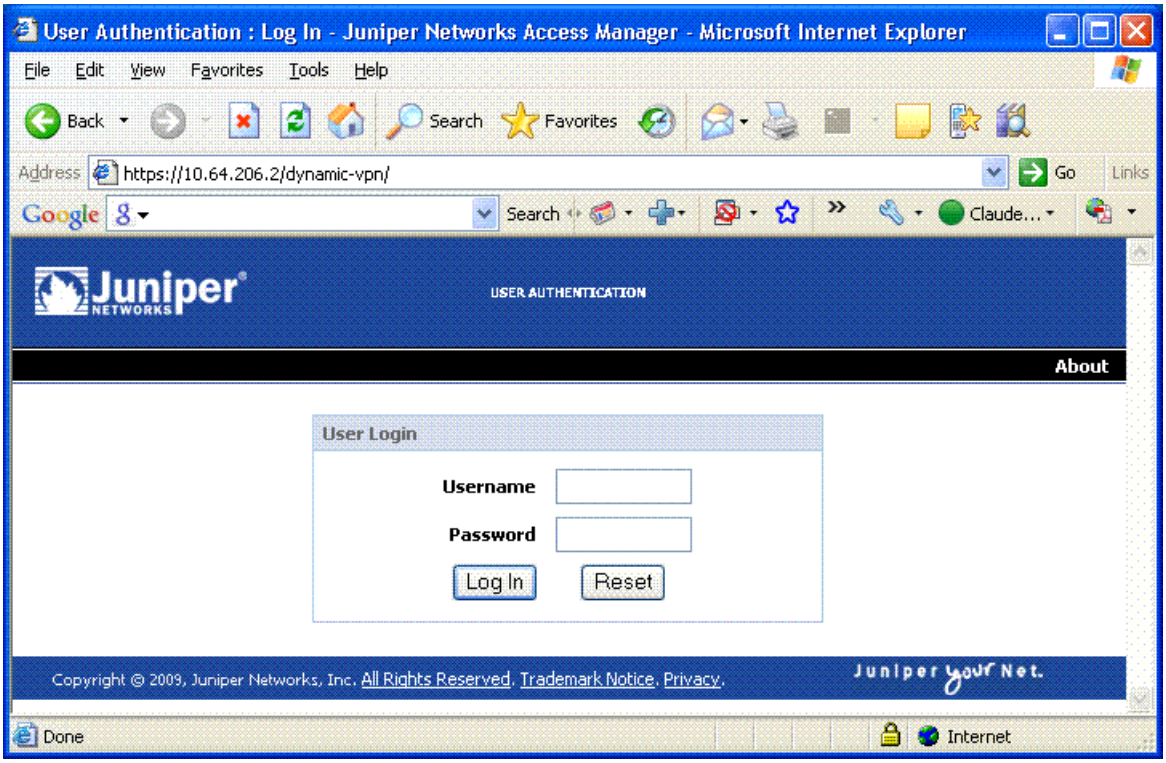

Figure 58: Access Manager Setup Client Installation Prompt

Figure 59: Access Manager Manual Installation Prompt

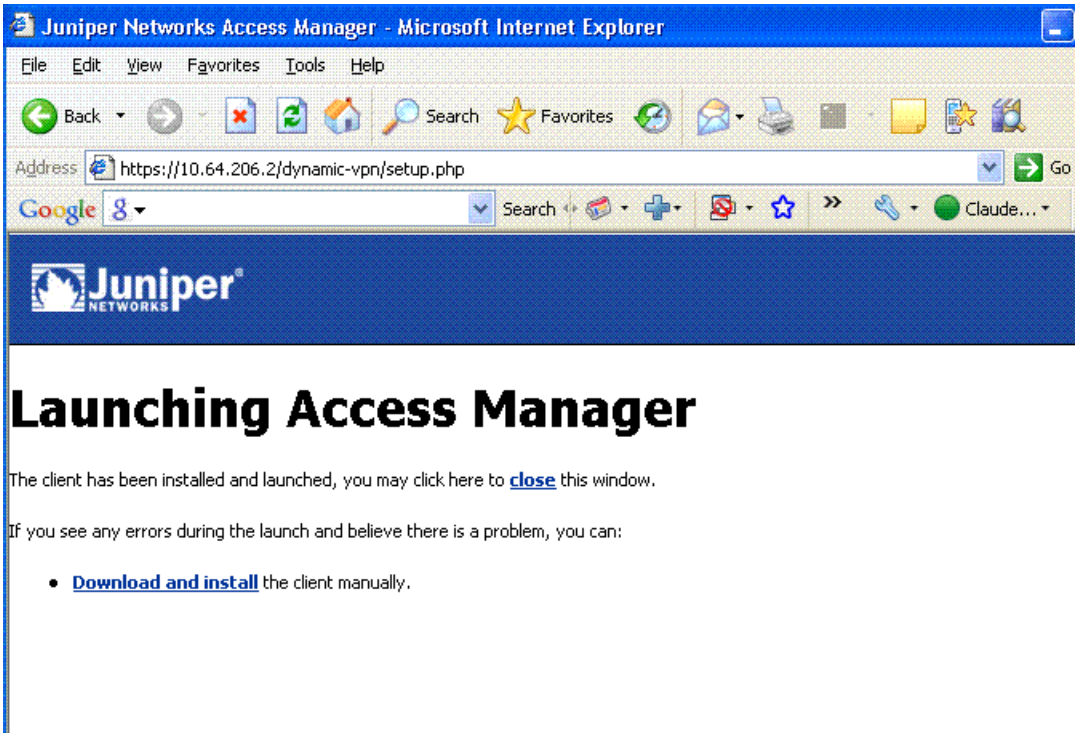

Figure 60: Downloading Access Manager

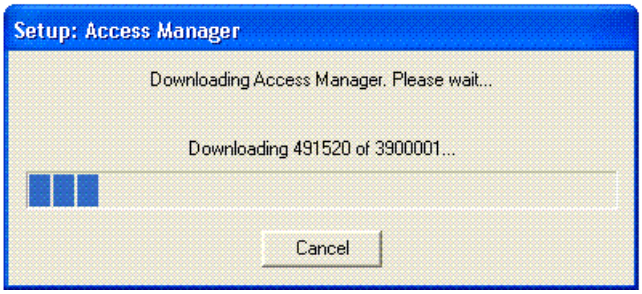

Figure 61: Installing Access Manager

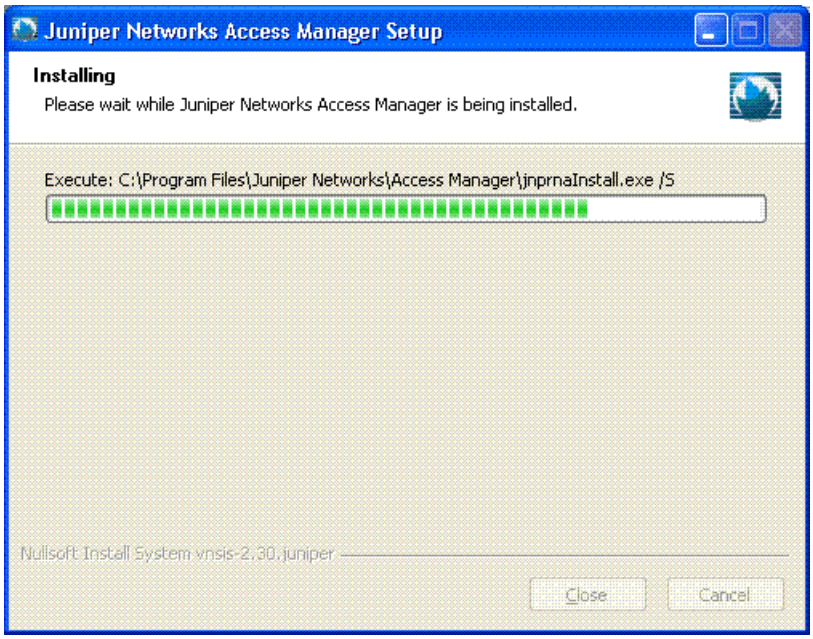

Figure 62: Access Manager Reboot Prompt

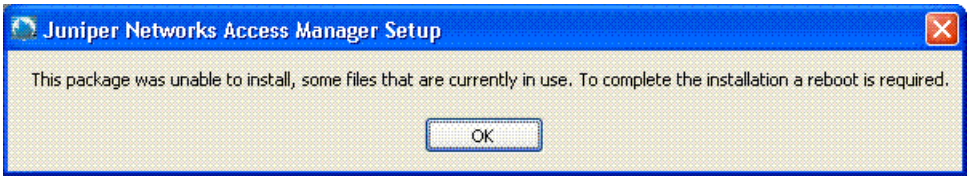

Connecting to the Remote Access Server for Subsequent Sessions (Pre-IKE Phase)

If the user has already downloaded Access Manager, the process for signing into the Remote Access Server for subsequent sessions is as follows:

1. **The user accesses the server.** The user launches the Access Manager client that is preinstalled on the computer (recommended). Alternatively, the user can access the server through the dynamic VPN URL (https://*<serverhost>*/dynamic-vpn). If the user chooses to use this method, however, the server downloads Access Manager to the user's desktop—even if the client already exists. As part of this process, the Remote Access Server prompts the user for a username and password, checks that the user has the proper install privileges, generates a new connection token, and downloads the setup client, as explained in "Connecting to the Remote Access Server for the First Time (Pre-IKE Phase)" on page 435.
2. **The client determines if an upgrade is required.** The client checks the client configuration version installed on the user's computer. If a more recent version of the client is available, the client code either automatically upgrades the client software (if you have enabled the Force Upgrade option) or gives the user the option of upgrading (if you have not enabled the Force Upgrade option.

The user can initiate a secure VPN connection to the Remote Access Server from Access Manager once it is successfully launched. For more information, see "Establishing an IPsec VPN Tunnel (IKE Phase)" on page 439.

Establishing an IPsec VPN Tunnel (IKE Phase)

Once Access Manager is installed, the user can use it to initiate a secure VPN tunnel to the Remote Access Server as follows:

1. **The user launches Access Manager.** The user can launch Access Manager by using either of the following methods:
   - Choose **All Programs > Juniper Networks > Access Manager** from the Windows Start menu.
   - Select the Access Manager icon in the system tray at the lower right corner of the Windows screen.

   When the user launches the client, the Access Manager dialog box appears.
2. **The user creates a connection to the server, if necessary.** If no connections are available in the Access Manager dialog box, the user must specify a connection server:
   a. From the File menu, choose **Setup Connection**.
   b. In the New Connection dialog box that appears, enter the hostname of the Remote Access Server and the appropriate username.
   c. Click **OK**. The specified connection appears in the Access Manager dialog box.
3. **The user starts the connection.** In the Access Manager dialog box, the user selects which server connection to initiate by using one of the following methods:
   - Select one of the connections, right-click, and choose **Connect**.
   - Select one of the connections, and from the File menu, choose **Start Connection**.
4. **The server checks for a valid license.** When the user initiates a connection to a Remote Access Server, the server checks that a seat license is currently available for the user's session.
5. **The user signs into the server.** The user enters the appropriate username and password into the login page, and the Remote Access Server sends them to the authentication server for validation.

NOTE: The username and password entered here are used to validate the user's eligibility to establish the VPN session. These credentials are separate from those used to validate the user's eligibility to download the client (as explained in "Connecting to the Remote Access Server for the First Time (Pre-IKE Phase)" on page 435).

6. **The client initiates the VPN session.** Once the user has successfully authenticated, the client sends a preshared key to the Remote Access Server. (The client initially received this key as part of the initial client configuration download.) The client and server use an AutoKey IKE exchange to create security associations (SAs) and establish a secure VPN tunnel (as explained in "Internet Protocol Security (IPsec)" on page 355).

Figure 63: Access Manager System Tray Icon

Figure 64: Access Manager Dialog Box

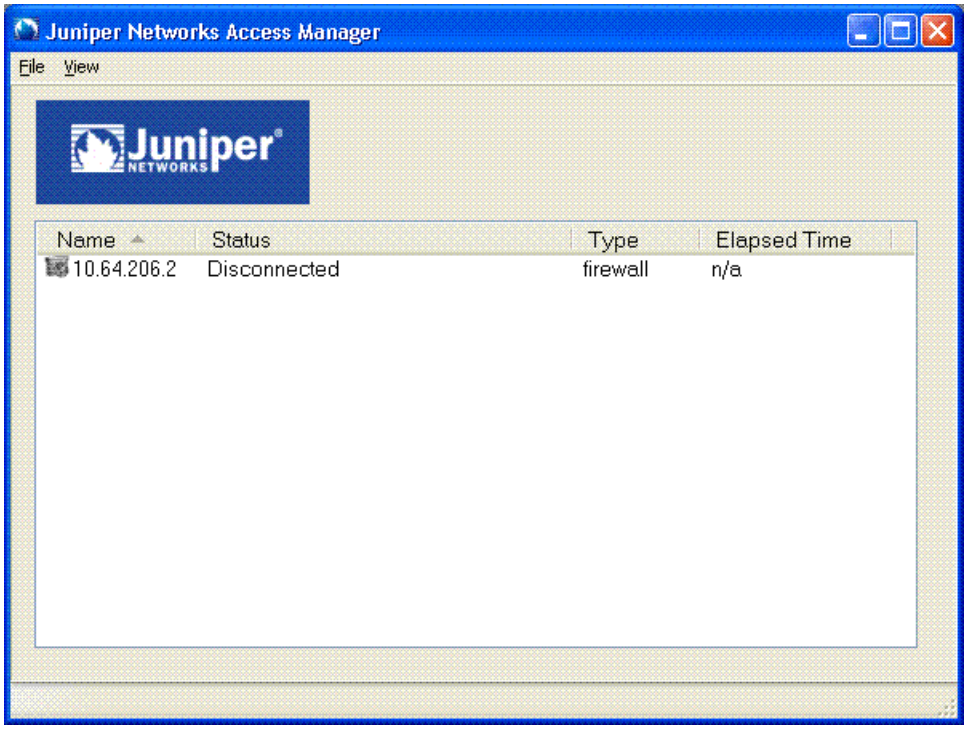

Figure 65: Creating a Firewall Connection in Access Manager

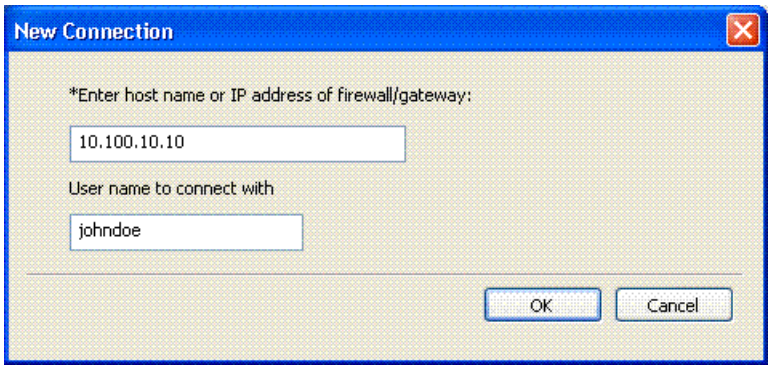

Figure 66: Starting a Firewall Connection in Access Manager

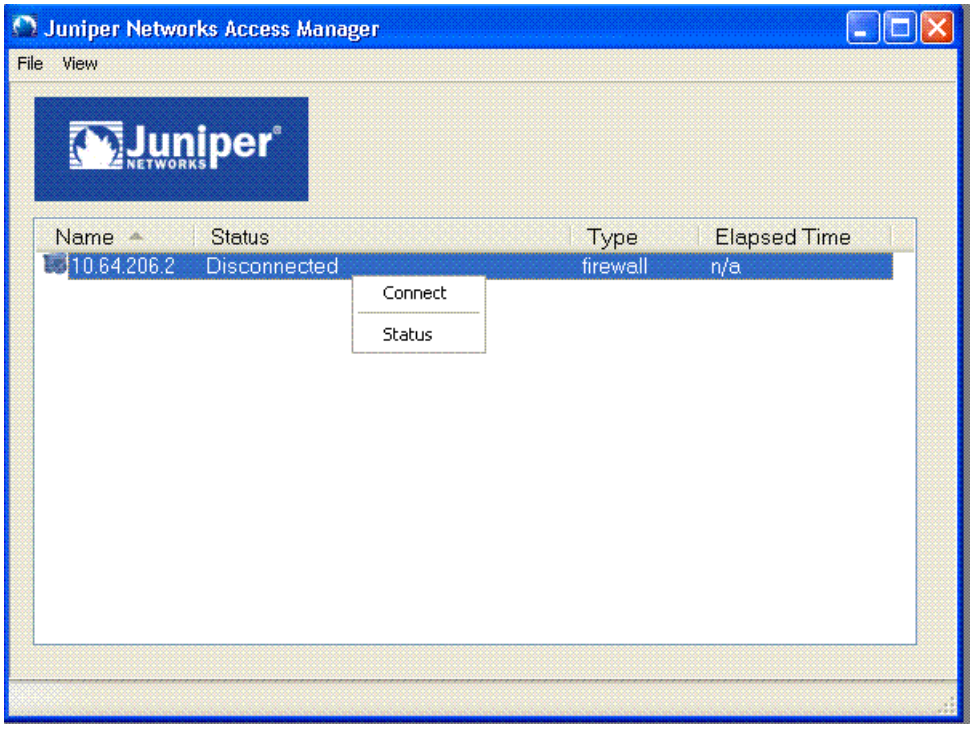

Figure 67: Successful Firewall Connection in Access Manager

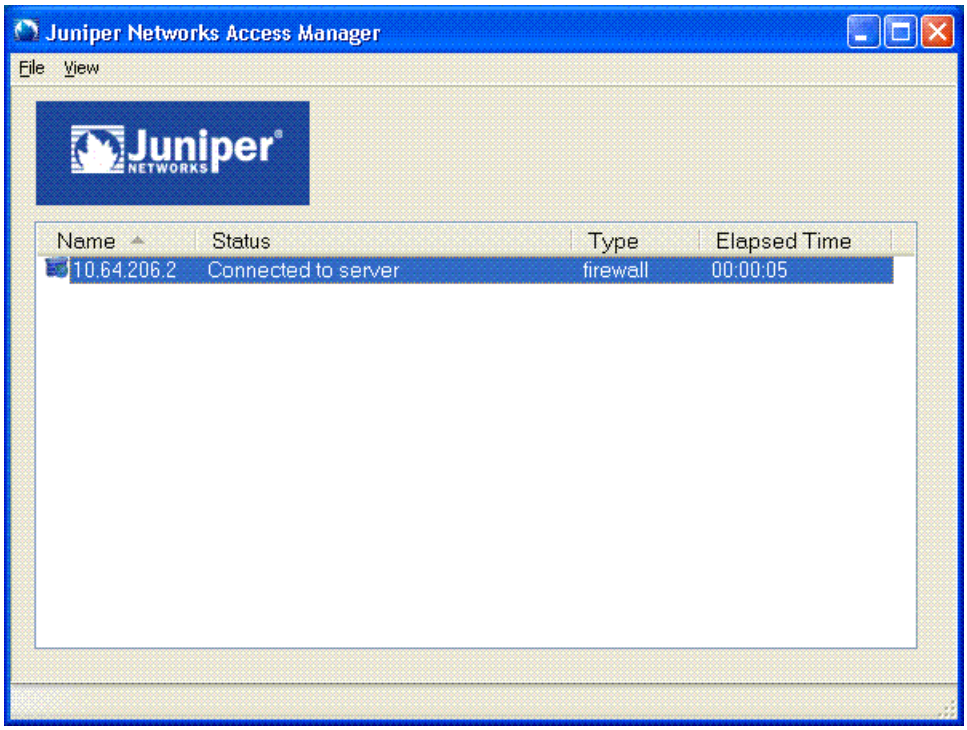

Related Topics

"Configuring a Dynamic VPN—Overview" on page 449

"Access Manager Client-Side Reference" on page 442

Access Manager Client-Side Reference

The following sections outline information that can help end-users and administrators when installing, running, and troubleshooting Access Manager on end-users' computers.

- Client-Side System Requirements on page 442
- Client-Side Files on page 443
- Client-Side Registry Changes on page 445
- Client-Side Error Messages on page 446
- Client-Side Logging and Connection Information on page 448

Client-Side System Requirements

The user can install Access Manager on any Windows XP or Windows Vista machine with an Internet connection. The user must have administrator privileges to install the client, but not to run it.

Access Manager can run simultaneously on the same computer with other Juniper Networks clients, including the Odyssey Access Client (OAC), Network Connect client, Windows Secure Application Manager (WSAM) client, Host Checker client, and WX client.

Client-Side Files

This section lists the directories where Access Manager installs files on a user's computer, the files it installs, and the files that remain after the user uninstalls the client.

Table 59: Access Manager Client-Side Files

Installation Directory	Files Installed in Directory	Files Remaining After Uninstall
%COMMONFILES%\Juniper Networks\Connection Manager	■ ConnectionManagerService.dll ■ install.log ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini	install.log
%COMMONFILES%\Juniper Networks\ConnectionStore	■ ConnectionStoreService.dll ■ dcfDOM.dll ■ install.log ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini	install.log
%COMMONFILES%\Juniper Networks\IPSecMgr	■ install.log ■ ipsecmgr.dll ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini	install.log
PROGRAMFILES%\Juniper Networks\Juniper Access Manager	■ AccessServiceComponent.x86.exe ■ ConnectionMgrComponent.x86.exe ■ ConnectionStoreComponent.x86.exe ■ install.log ■ IPSecMgrComponent.x86.exe ■ JamGUIComponent.x86.exe ■ JamInstaller.dep ■ jnprnaInstall.exe ■ TunnelManagerComponent.x86.exe ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini ■ vpnAccessMethodComponent.x86.exe	install.log Log file location: C:\Documents and Settings\All Users\Application Data\Juniper Networks\Logging

Table 59: Access Manager Client-Side Files *(continued)*

Installation Directory	Files Installed in Directory	Files Remaining After Uninstall
%COMMONFILES%\Juniper Networks\JamUI	■ install.log ■ jamCommand.exe ■ jamTray.exe ■ jamUI.exe ■ jamUIResource_EN.dll ■ uiPlugin.dll ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini	install.log
%COMMONFILES%\Juniper Networks\JUNS	■ access.ini ■ dsAccessService.exe ■ dsInstallerService.dll ■ dsLogService.dll ■ install.log ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini	install.log
%COMMONFILES%\Juniper Networks\JNPRNA	■ install.log ■ jnprna.cat ■ jnprna.inf ■ jnprna.sys ■ jnprnaapi.dll ■ jnprnaNetInstall.dll ■ jnprnaNetInstall.log ■ jnprna_m.cat ■ jnprna_m.inf ■ jnprva.cat ■ jnprva.inf ■ jnprva.sys ■ jnprvamgr.cat ■ jnprvamgr.dll ■ jnprvamgr.inf ■ jnprvamgr.sys ■ nsStatsDump.exe ■ uninst.exe ■ versionInfo.ini ■ %WINDIR%\system32\drivers\jnprna.sys ■ %WINDIR%\system32\drivers\jnprva.sys ■ %WINDIR%\system32\drivers\jnprvamgr.sys	■ install.log ■ jnprnaNetInstall.log

Table 59: Access Manager Client-Side Files *(continued)*

Installation Directory	Files Installed in Directory	Files Remaining After Uninstall
COMMONFILES%\Juniper Networks\Tunnel Manager	■ dsTMClient.dll ■ dsTMService.dll ■ dsTunnelManager.dll ■ install.log ■ TM.dep ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini	install.log
%COMMONFILES%\Juniper Networks\vpnAccessMethod	■ install.log ■ Uninstall.exe ■ Uninstall.exe.manifest ■ versionInfo.ini ■ vpnAccessMethod.dll ■ vpnAccessMethod_EN.dll	install.log

Client-Side Registry Changes

This section lists the Windows Registry changes that the Access Manager client and components make to your users' computers when creating dynamic VPN tunnels.

Table 60: Access Manager Client-Side Registry Changes

Registry Key Location	Registry Key Changes
HKEY_LOCAL_MACHINE\SOFTWARE\Juniper Networks\Common Files	■ jnprnaapi" = "C:\\Program Files\\Common Files\\Juniper Networks\\JNPRNA\\jnprnaapi.dll ■ jnprvamgr" = "C:\\Program Files\\Common Files\\Juniper Networks\\JNPRNA\\jnprvamgr.dll ■ nsStatsDump" = "C:\\Program Files\\Common Files\\Juniper Networks\\JNPRNA\\nsStatsDump.exe ■ dsLogService" = "C:\\Program Files\\Common Files\\Juniper Networks\\JUNS\\dsLogService.dll ■ dsTMClient" = "C:\\Program Files\\Common Files\\Juniper Networks\\Tunnel Manager\\dsTMClient.dll ■ dsTunnelManager" = "C:\\Program Files\\Common Files\\Juniper Networks\\Tunnel Manager\\dsTunnelManager.dll
HKEY_LOCAL_MACHINE\SOFTWARE\Juniper Networks\Logging	■ LogFileName" = "C:\\Documents and Settings\\All Users\\Application Data\\Juniper Networks\\Logging\\debuglog.log ■ "Level" = "3" ■ "LogSizeInMB" = "10"

Table 60: Access Manager Client-Side Registry Changes *(continued)*

Registry Key Location	Registry Key Changes
HKCU\Software\Juniper Networks\Access Manager\	(Content varies. Contains client configuration data downloaded from the server.)

Client-Side Error Messages

This section lists possible errors that end users might see when installing or running Access Manager.

The following table lists the possible errors that the end user might see, the possible causes for the messages, and suggested actions.

Table 61: Dynamic VPN Client-Side Errors

Error Message	Possible Causes	Suggested User Action
Component instance already in use	Internal error	Try to reconnect to the firewall.
Memory allocation failure	Internal error	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
Failed to load connection store	Internal error	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
Cannot get connection information for firewall	Internal error. Could not retrieve connection information for the specified firewall.	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
Authentication failure: Unknown HTTP response code	Internal error. Could not decipher the HTTP response.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Authentication failure: Incorrect username or password	The user entered an invalid username or password.	Reenter your credentials.
Authentication failure: Firewall is out of licenses	All available licenses are currently being used for other dynamic VPN sessions or no licenses are installed for the feature.	Try to reconnect to the firewall once a license has been freed by another user. If the problem persists, contact your system administrator.
Authentication failure: No configuration available	No configuration is currently available for the specified user account.	Contact your system administrator.
Cannot create IPsec route entry	Internal error. Failed to read route entry from the connection store.	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
Failed to read route entry from connection store	Internal error. Failed to read the route entry from the connection store.	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.

Table 61: Dynamic VPN Client-Side Errors *(continued)*

Error Message	Possible Causes	Suggested User Action
Failed to add route entry to policy	Internal error. Failed to add the route entry to the connection store.	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
Failed to initialize IPsec Manager	Internal error. Failed to initialize the IPsec Manager.	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
IPsec authentication failed	Phase 1 negotiations, extended authentication (Xauth), or Phase 2 negotiations failed.	Try to reconnect to the firewall.
IPsec configuration failed	Internal error or policy configuration error. The Tunnel Manager was unable to configure the local IP settings.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
IKE negotiations failed	The components cannot agree on security parameters during the IKE exchange. The administrator probably needs to reconfigure the Phase 1 proposal.	Contact your system administrator.
Failed to initialize authentication	Failed to authenticate when connecting to the firewall, possibly because the specified hostname did not resolve against the distinguished name server (DNS).	Try to reconnect to the firewall.
Failed to connect to server	The TCP connection to the Web server failed during authentication, possibly due to network connectivity issues.	Try to reconnect to the firewall.
Failed to send initial HTTP request	Web server authentication failed, possibly due to network connectivity issues.	Try to reconnect to the firewall.
Failed to get HTTP response	Web server authentication failed.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Firewall refused authentication request	Web server authentication failed.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Client failed to provide login page	Web server authentication failed.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Server failed to send authentication request	Web server authentication failed.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Server failed to respond to authentication request	The client sent the user's credentials to the Web server, but the server failed to respond in a useful manner.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Authentication negotiation failed	Web server authentication failed.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Failed to get configuration from firewall	Web server authentication failed.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
The user cancelled authentication.	User canceled authentication	Try to reconnect to the firewall and reenter your credentials.

Table 61: Dynamic VPN Client-Side Errors *(continued)*

Error Message	Possible Causes	Suggested User Action
Failed to enter username or password	Authentication request timed out.	Try to reconnect to the firewall and reenter your credentials.
Server failed to request username and password	The client failed to display the user interface asking the user for credentials.	Exit and restart Access Manager. If the problem persists, contact your system administrator.
Your client state is preventing the connection	The user's client is in an inoperable state.	Try to reconnect to the firewall. If the problem persists, exit and restart Access Manager.
Cannot open connection store	The client could not contact the connection store.	Exit and restart Access Manager. If the problem persists, reinstall Access Manager.
Cannot process configuration provided by firewall	The script provided by the firewall was in someway unusable. The configuration might need to be updated on the server.	Try to reconnect to the firewall. If the problem persists, contact your system administrator.
Access Manager is not running	The Access Manager service is not running.	Exit and restart Access Manager.
Please select a connection	The user chose **Start Connection** without selecting a connection first.	Select the firewall you want to connect to and then choose **Start Connection**.
Are you sure you want to delete the selected connection?	The user chose **Delete Connection**.	Specify whether or not you want to delete the selection connection profile.
Cannot add new connection. Service is not running.	The Access Manager service is not running; therefore it cannot create a new connection profile.	Exit and restart Access Manager. If the problem persists, reinstall Access Manager.
Cannot add new connection	The Access Manager failed to add the new connection profile.	Try again. If the problem persists, exit and restart Access Manager.
Connection name is already in use	Unable to add connection profile because the specified connection name already exists.	Specify a unique name for the connection profile.
Please reinstall Access Manager	Files could not be found when trying to finish the operation.	Reinstall Access Manager.
Invalid server certificate	Certificate validation failed.	Check client-side logs to determine why the certificate failed.
Initializing service...	Initializing one of the client's core components. If the component does not initialize, the client cannot function.	Wait for the service to finish initializing.

Client-Side Logging and Connection Information

Access Manager provides the following logging and connection information that you can use to troubleshoot client-side issues:

- **Client-side logs:** To view client-side logs, open Access Manager and choose **Save logs and diagnostics** from the File menu. Select a location on your computer to save the zipped log files and click **Save**.
- **Detailed logs:** To create more detailed client-side logs, open Access Manager and choose **Enable Detailed Logging** from the File menu.
- **Firewall connection information:** To view connection information for a given firewall, open Access Manager, select the firewall, right-click, and choose **Status**.

Configuring a Dynamic VPN—Overview

The dynamic VPN feature secures traffic through your network by passing it through IPsec VPN tunnels. To configure an IPsec VPN tunnel, you must specify Phase 1 settings (which enable participants to establish a secure channel in which to negotiate the IPsec security association (SA)), and Phase 2 settings (which enable participants to negotiate the IPsec SA that authenticates traffic flowing through the tunnel). This section describes the order in which you must configure these tunnel negotiation settings as well as other tasks you must complete in order to enable the tunnels on your network.

The dynamic VPN feature is disabled by default on the device. You must enable and configure it before you can use it.

Before You Begin

For background information, read

- "Dynamic VPN Overview" on page 433.
- "Virtual Private Networks (VPNs) Overview" on page 355.
- "Understanding IKE and IPsec Packet Processing" on page 364.

To configure the dynamic VPN feature, you must do the following:

1. Define an outgoing interface by using the edit interfaces configuration statement. Use this interface to pass IKE security associations (SAs) through the device. (You will need to select this interface when configuring your IKE gateway.) For more information about interfaces, see the *JUNOS Software Interfaces and Routing Configuration Guide*.
2. Create security policies by using the edit security policies configuration statement. Use these policies to define which traffic can pass through your network. (After you create your VPN configuration, you will need to add it to this policy.) For more information about security policies, see "Security Policies Overview" on page 69
3. Create at least one access profile by using the edit access profile configuration statement. Use the access profile(s) to control the authentication of users who want to download Access Manager and users who want to establish dynamic VPN tunnels to your firewall. (You will need to select these access profiles when configuring your IKE gateway and dynamic VPN global options. Note that you can use the same access profile to authenticate users in both cases, or you can

use separate access profiles to authenticate downloads and VPN sessions.) For more information about access profiles, see “Understanding Authentication Schemes” on page 307.

4. Create an IKE gateway to include in your VPN configuration:

   a. Create one or more IKE Phase 1 proposals by using the edit security ike proposal configuration statement. (You will need to select this proposal when configuring your IKE policy.) For more detailed configuration instructions, see:

      - Configuring an IKE Phase 1 Proposal (Standard and Dynamic VPNs) on page 373
      - Configuring an IKE Phase 1 Proposal— (Dynamic VPNs) on page 451

   b. Create one or more IKE policies by using the edit security ike policy configuration statement. (You will need to select this policy when configuring your IKE gateway.) For more detailed configuration instructions, see:

      - Configuring an IKE Policy (Standard and Dynamic VPNs) on page 384
      - Configuring an IKE Policy—(Dynamic VPNs) on page 454

   c. Create an IKE gateway configuration by using the edit security ike gateway configuration statement. (You will need to select this gateway when configuring your IPsec AutoKey.) For more detailed configuration instructions, see:

      - Configuring an IKE Gateway (Standard and Dynamic VPNs) on page 385
      - Configuring an IKE Gateway—Quick Configuration (Dynamic VPNs) on page 456

5. Create an IPsec AutoKey to include in your VPN configuration:

   a. Create one or more IPsec Phase 2 proposals by using the edit security ipsec proposal configuration statement. (You will need to select this proposal when configuring your IPsec policy.) For more detailed configuration instructions, see:

      - Configuring an IPsec Phase 2 Proposal (Standard and Dynamic VPNs) on page 387
      - Configuring an IPsec Phase 2 Proposal—Quick Configuration (Dynamic VPNs) on page 460

   b. Create one or more IPsec policies by using the edit security ipsec policy configuration statement. (You will need to select this policy when configuring your IPsec AutoKey.) For more detailed configuration instructions, see:

      - Configuring an IPsec Policy (Standard and Dynamic VPNs) on page 388
      - Configuring an IPsec Policy—Quick Configuration (Dynamic VPNs) on page 463

   c. Create an IKE AutoKey configuration by using the edit security ipsec autokey configuration statement. (You will need to select this IKE AutoKey configuration when configuring your VPN client configuration.) For more detailed configuration instructions, see:

- Configuring IPsec AutoKey (Standard and Dynamic VPNs) on page 390
- Configuring an IPsec Autokey—Quick Configuration (Dynamic VPNs) on page 465

6. Create a client VPN configuration by using the edit security dynamic-vpn clients configuration statement. The settings are downloaded as part of the client to your users' computers and are used to establish the dynamic VPN tunnels between the clients and the server. For more detailed configuration instructions, see:
   - Creating a Client Configuration—Quick Configuration (Dynamic VPNs) on page 469
   - Creating a Client Configuration (Dynamic VPNs) on page 472
7. Update your security policy (or policies) to include your client VPN configuration by using the edit security from-zone *zone-name* to-zone *zone-name* policy then permit tunnel ipsec-vpn *vpn-name* configuration statement. For more information about policies, see "Security Policies Overview" on page 69.
8. Specify global settings for client downloads by using the edit security dynamic-vpn access-profile configuration statement and the edit security dynamic-vpn force-upgrade configuration statement. For more detailed configuration instructions, see:
   - Configuring Global Client Download Settings—Quick Configuration (Dynamic VPNs) on page 473
   - Configuring Global Client Download Settings (Dynamic VPNs) on page 475

Configuring an IKE Phase 1 Proposal— (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly configure an IKE Phase 1 proposal. (For information about configuring a Phase 1 proposal using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Configuring an IKE Phase 1 Proposal (Standard and Dynamic VPNs)" on page 373.)

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

Figure 68 on page 452 shows the Configuration page where you can select an existing proposal, or click **Add** to create a new one.

Figure 68: IKE Phase 1 Proposal Quick Configuration Page – Adding a Proposal

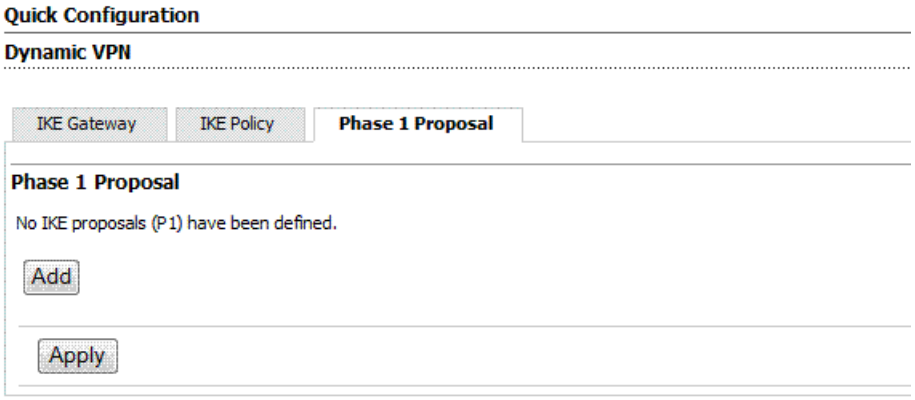

Figure 69 on page 452 shows the Configuration page where you create a new proposal.

Figure 69: IKE Phase 1 Proposal Quick Configuration Page – Configuring a Proposal

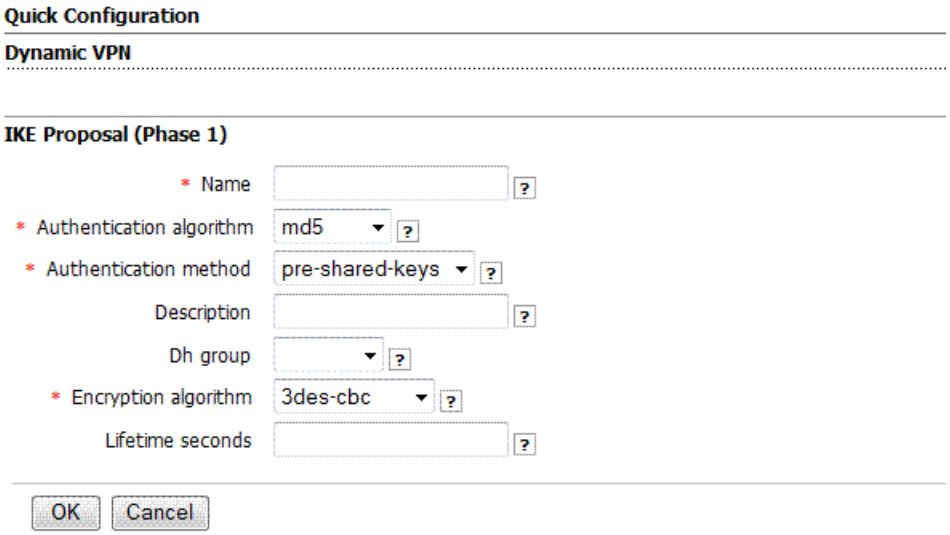

To configure a Phase 1 Proposal with Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > IKE**.
2. Select the **Phase 1 Proposal** tab if it is not selected.
3. To modify an existing proposal, click the appropriate link in the **Name** column to go to the proposal's configuration page. Or, select the proposal from among those listed and click one of the following buttons:
   - To apply the configuration, click **Apply**.
   - To delete the configuration, click **Delete.**
4. To configure a new Phase 1 proposal, click **Add**.

5. Fill in the options as described in Table 62 on page 453.
6. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration and return to the main Configuration page, click **Cancel**.

Table 62: Phase 1 Proposal Configuration Options

Field	Function	Action
IKE Proposal (Phase 1)		
Name	Name to identify the proposal.	Enter a name.
Authentication algorithm	Authentication Header (AH) algorithm the device uses to verify the authenticity and integrity of a packet. Supported algorithms include the following: ■ **md5**—Produces a 128-bit digest. ■ **sha1**—Produces a 160-bit digest. ■ **sha-256**—Produces a 256-bit digest.	Select an authentication algorithm.
Authentication method	Method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. The dynamic VPN feature only uses preshared keys for authentication. With this method, both participants must have the key before beginning tunnel negotiations.	No action is required. The device displays this information for informational purposes only.
Description	Description of the proposal.	Enter a brief description of the Phase 1 proposal.
Dh group	Allow participants to produce a shared secret value over an unsecured medium without actually transmitting the value across the connection.	Select a Diffie-Hellman group. If you configure multiple (up to four) proposals for Phase 1 negotiations, use the same Diffie-Hellman group in all proposals.
Encryption algorithm	Supported Internet Key Exchange (IKE) proposals include the following: ■ **3des-cbc**—3DES-CBC encryption algorithm ■ **aes-128-cbc**—AES-CBC 128-bit encryption algorithm ■ **aes-192-cbc**—AES-CBC 192-bit encryption algorithm ■ **aes-256-cbc**—AES-CBC 256-bit encryption algorithm ■ **des-cbc**—DES-CBC encryption algorithm	Select an encryption algorithm.
Lifetime seconds	Lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is either replaced by a new SA and security parameter index (SPI) or the SA is terminated.	Select a lifetime for the IKE security association (SA). Range: 180 through 86,400 seconds. Default: 3,600 seconds.

Configuring an IKE Policy—(Dynamic VPNs)

You can use J-Web Configuration to configure an IKE policy. (For information about configuring an IKE policy using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see “Configuring an IKE Policy (Standard and Dynamic VPNs)” on page 384.)

Before You Begin

For background information, read:

- “Dynamic VPN Overview” on page 433.
- “Configuring a Dynamic VPN—Overview” on page 449.

Figure 70 on page 454 shows the Quick Configuration page where you can select an existing proposal, or click **Add** to create a new one.

Figure 70: IKE Policy Quick Configuration Page – Adding a New Policy

Figure 71 on page 455 shows the Quick Configuration page where you create a new IKE policy.

Figure 71: IKE Policy Quick Configuration Page – Configuring a Policy

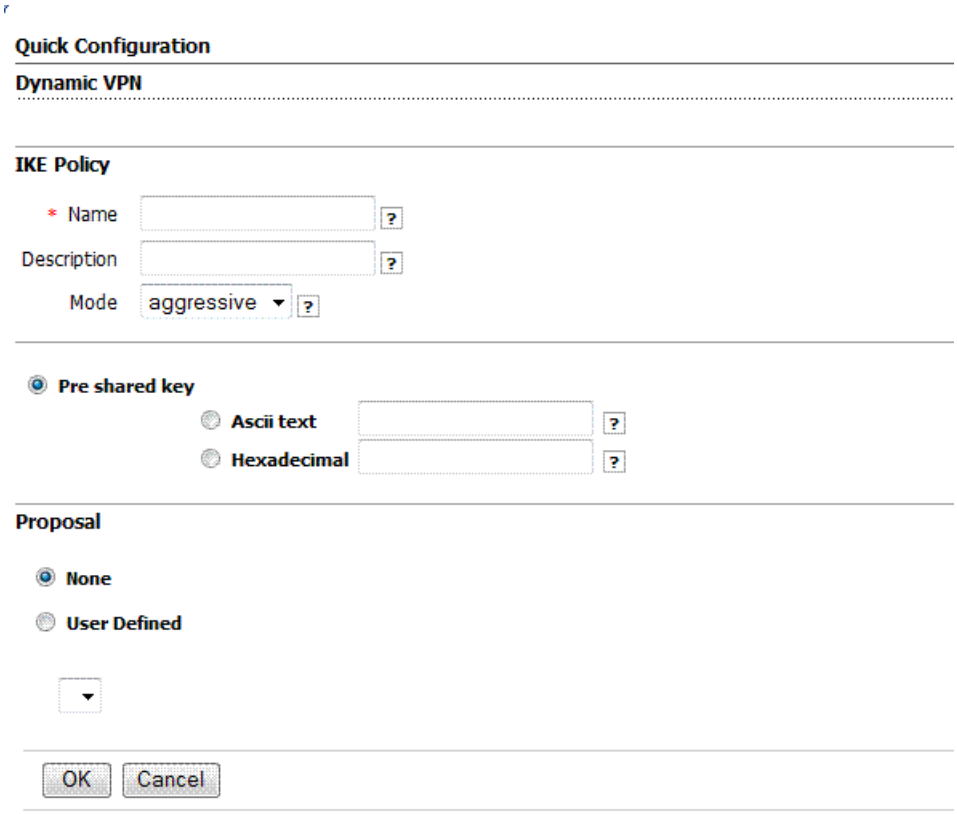

To configure an IKE policy with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > IKE**.
2. Select **IKE Policy** tab if it is not selected.
3. To modify an existing policy, click the appropriate link in the **Name** column to go to the policy's configuration page. Or, select the policy from among those listed and click one of the following buttons:
   - To apply the configuration, click **Apply**.
   - To delete the configuration, click **Delete.**
4. To configure a new IKE policy, click **Add**.
5. Fill in the options as described in Table 63 on page 456.
6. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration, click **Cancel**.

Table 63: IKE Policy, Authentication, and Proposal Options

Field	Function	Action
IKE Policy		
Name	Name to identify the policy.	Enter a name.
Description	Description of the policy.	Enter a brief description of the policy.
Mode	Specifies how participants should exchange encryption and authentication information during Phase 1 tunnel negotiations. The dynamic VPN feature only uses aggressive mode, which transfers the information between participants in two exchanges.	No action is required. The device displays this information for informational purposes only.
Pre-shared Key		
Pre-shared key	Use one of the following preshared key types: ■ ASCII text ■ Hexadecimal	Click **Pre shared key**, click the type of key, and enter the key in the appropriate format.
Proposal		
None	Do not use proposals	Click **None**.
User Defined	Use up to four Phase 1 proposals that you previously defined. If you include multiple Phase 1 proposals in the IKE policy, use the same Diffie-Hellman group in all of the proposals.	Click **User Defined**, select a proposal (or proposals) from the pop-up menu, and click **Add**.

Configuring an IKE Gateway—Quick Configuration (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly configure an IKE Gateway. (For information about configuring an IKE gateway using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Configuring an IKE Gateway (Standard and Dynamic VPNs)" on page 385.)

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

Figure 72 on page 457 shows the Quick Configuration page where you can select an existing gateway, or click **Add** to create a new one.

Figure 72: IKE Gateway Quick Configuration Page – Adding a Gateway

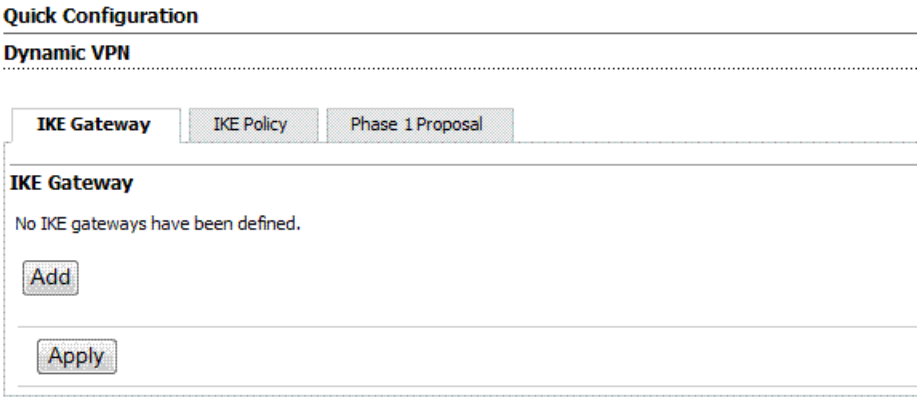

Figure 73 on page 458 shows the Quick Configuration page where you create a new IKE gateway.

Figure 73: IKE Gateway Quick Configuration Page – Configuring a Gateway

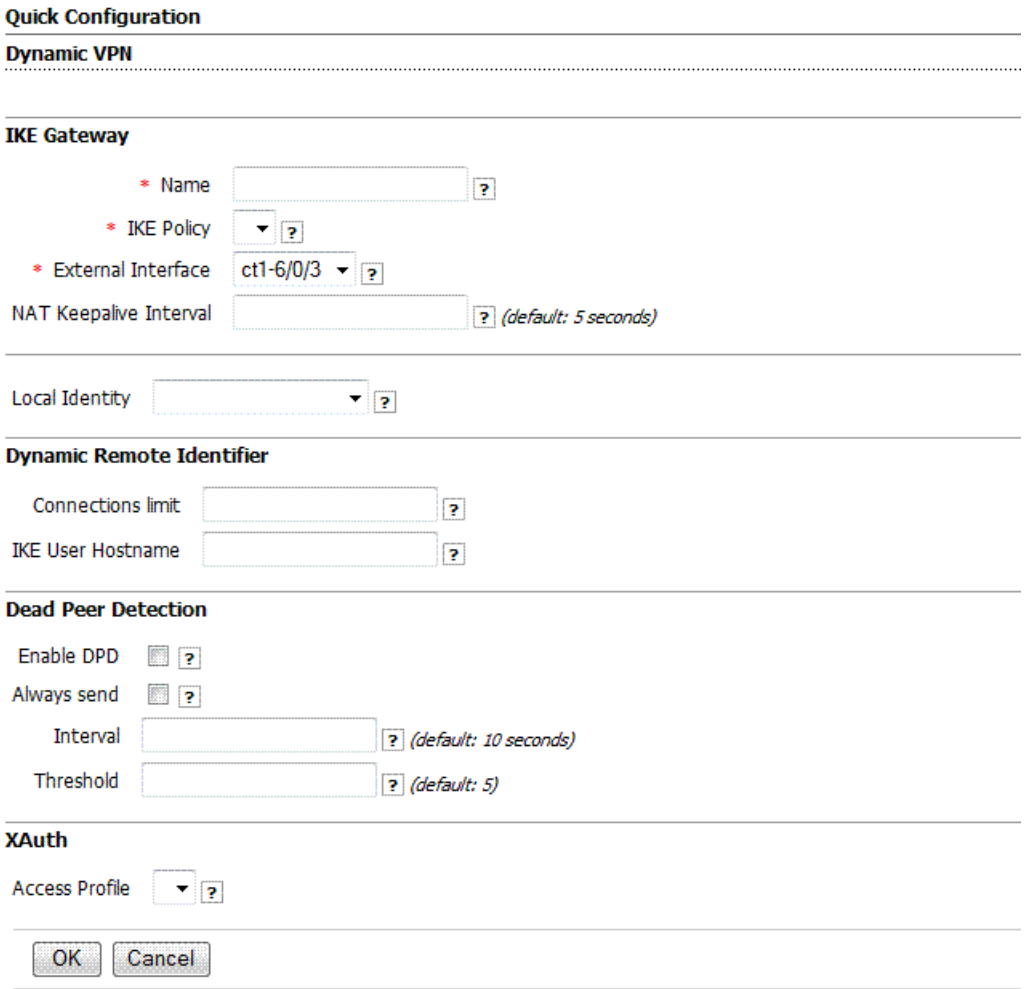

To configure an IKE gateway with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > IKE**.
2. Select the **IKE Gateway** tab if it is not selected.
3. To modify an existing IKE gateway, click the appropriate link in the **Name** column to go to the gateway's configuration page. Or, select the gateway from among those listed and click one of the following buttons:
   - To apply the configuration, click **Apply**.
   - To delete the configuration, click **Delete.**

NOTE: The list of IKE gateways displayed on this page includes both standard VPN gateways and dynamic VPN gateways.

4. To configure a new IKE gateway, click **Add**.
5. Fill in the options as described in Table 64 on page 459.
6. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration and return to the main Configuration page, click **Cancel**.

Table 64: IKE Gateway Options

Field	Function	Action
IKE Gateway		
Name	Name to identify the IKE gateway.	Enter a name.
IKE Policy	IKE policy to associate with the IKE gateway. An IKE policy specifies the type of preshared key to use during Phase 1 negotiations as well as which Phase 1 proposal(s) to use.	Select a previously created IKE policy.
External Interface	Outgoing interface to use when establishing security associations (SAs). An interface acts as a doorway through which traffic enters and exits the JUNOS device.	Specify a previously created interface.
NAT Keepalive Interval	The dynamic VPN feature automatically includes support for NAT traversal (NAT-T). The NAT keepalive interval controls how often NAT keepalive packets can be sent so that NAT translation continues.	Specify a maximum interval in seconds at which NAT keepalive packets can be sent. Range: 1 through 300 seconds. Default: 5 seconds.
Local Identity	Local identity of the endpoint computer to send in the IKE exchange. You can identify the local identity in any of the following ways: ■ **IP Address**—Use an IPv4 IP address to identify the endpoint computer. ■ **Hostname**—Use a fully qualified domain name (FQDN) to identify the endpoint computer. ■ **User at Hostname**—Use an e-mail address to identify the endpoint computer. If you do not configure a local identity, the device uses the virtual IP address assigned by the Radius server during the Xauth configuration exchange.	Specify an IP address, hostname, or user-at-hostname.
Dynamic Remote Identifier		
Connections limit	Maximum number of concurrent connections allowed. When the maximum number of connections is reached, no more dynamic VPN endpoint users attempting to access an IPsec VPN are allowed to begin Internet Key Exchange (IKE) negotiations.	Specify the maximum number of concurrent users that can be connected to the gateway (Remote Access Server).

Table 64: IKE Gateway Options *(continued)*

Field	Function	Action
IKE User Hostname	Name or identifier to use when establishing the VPN tunnel. We recommend entering the fully qualified domain name to identify the dynamic peer, but you can enter any name or identifier as long as it is unique.	Specify one primary name or identifier and up to four backups.
Dead Peer Detection		
Enable DPD	Enable dead peer detection (DPD), as outlined in *RFC 3706 Dead Peer Detection*.	Click the check box to disable or enable. (Disabled by default.)
Always Send	Send DPD requests regardless of whether there is outgoing IPsec traffic to the peer.	Click the check box to disable or enable. (Disabled by default.)
Interval	Amount of time that the peer waits for traffic from its destination peer before sending a DPD request packet.	Enter the interval at which to send DPD messages. Range: 1 through 60 seconds. Default: 10.
Threshold	Maximum number of unsuccessful DPD requests that can be sent before the peer is considered unavailable.	Enter the maximum number of unsuccessful DPD requests to be sent. Range: 1 through 5. Default: 5.
XAuth		
Access Profile	Provide extended authentication (XAuth), in addition to IKE authentication, for remote users trying to access a VPN tunnel. **NOTE:** This Access Profile option does not control authentication for users trying to download Access Manager. For client download authentication, use the Access Profile option on the Global Settings Quick Configuration page. For more information, see "Configuring Global Client Download Settings—Quick Configuration (Dynamic VPNs)" on page 473.	Select a previously created access profile.

Configuring an IPsec Phase 2 Proposal—Quick Configuration (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly configure IPsec Phase 2 proposals. (For information about configuring a Phase 2 proposal using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Configuring an IPsec Phase 2 Proposal (Standard and Dynamic VPNs)" on page 387.)

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

Figure 74 on page 461 shows the Quick Configuration page where you can select an existing proposal, or click **Add** to create a new one.

Figure 74: Phase 2 Proposal Quick Configuration Page – Adding a Proposal

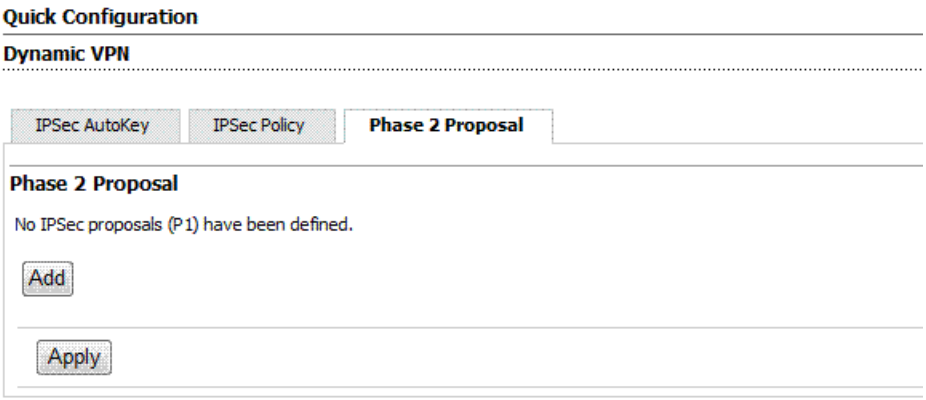

Figure 75 on page 461 shows the Quick Configuration page where you create a new proposal.

Figure 75: Phase 2 Proposal Quick Configuration Page – Configuring a Proposal

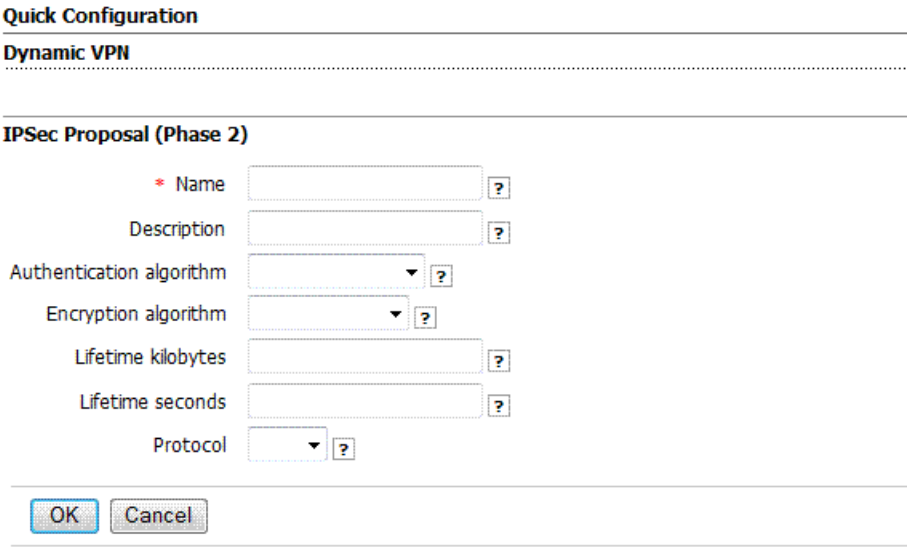

To configure an IPsec Phase 2 proposal with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > IPSec AutoKey**.
2. Select the IPsec **Phase 2 Proposal** tab if it is not selected

3. To modify an existing proposal, click the appropriate link in the **Name** column to go to the proposal's configuration page. Or, select the proposal from among those listed and click one of the following buttons:
   - To apply the configuration, click **Apply**.
   - To delete the configuration, click **Delete**.
4. To configure a new Phase 2 proposal, click **Add**.
5. Fill in the options as described in Table 65 on page 462.
6. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration and return to the main Configuration page, click **Cancel**.

Table 65: IPsec Phase 2 Proposal Options

Field	Function	Action
IPsec Proposal (Phase 2)		
Name	Name to identify the Phase 2 proposal.	Enter a name.
Description	Description of the Phase 2 proposal.	Enter a brief description of the proposal.
Authentication algorithm	Hash algorithm that authenticates packet data. You can choose one of the following: ■ **hmac-md5-96**—Produces a 128-bit digest. ■ **hmac-sha1-96**—Produces a 160-bit digest.	Select a hash algorithm.
Encryption algorithm	IKE algorithm used to encrypt data. You can choose one of the following: ■ **3des-cbc**—Has a block size of 24 bytes; the key size is 192 bits long. ■ **aes-128-cbc**—AES 128-bit encryption algorithm. ■ **aes-192-cbc**—AES 192-bit encryption algorithm. ■ **aes-256-cbc**—AES 256-bit encryption algorithm. ■ **des-cbc**—Has a block size of 8 bytes; the key size is 48 bits long.	Select an encryption algorithm.
Lifetime kilobytes	Lifetime (in kilobytes) of an IPsec security association (SA). The SA is terminated when the specified number of kilobytes of traffic have passed.	Enter a value from 64 through 1,048,576 bytes.
Lifetime seconds	Lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is either replaced by a new SA and security parameter index (SPI) or the SA is terminated.	Enter a value from 180 through 86,400 seconds.

Table 65: IPsec Phase 2 Proposal Options *(continued)*

Field	Function	Action
Protocol	Type of security protocol. Supported options include: ■ **ah**—Authentication Header (AH) protocol verifies the authenticity/integrity of the content and origin of a packet. ■ **esp**—Encapsulating Security Payload (ESP) protocol ensures privacy (encryption) and source authentication and content integrity (authentication).	Select a protocol for the proposal.

Configuring an IPsec Policy—Quick Configuration (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly configure an IPsec policy. (For information about configuring an IPsec policy using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Configuring an IPsec Policy (Standard and Dynamic VPNs)" on page 388.)

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

Figure 76 on page 463 shows the Quick Configuration page where you can select an existing policy, or click **Add** to create a new one.

Figure 76: IPsec Policy Quick Configuration Page – Adding a Policy

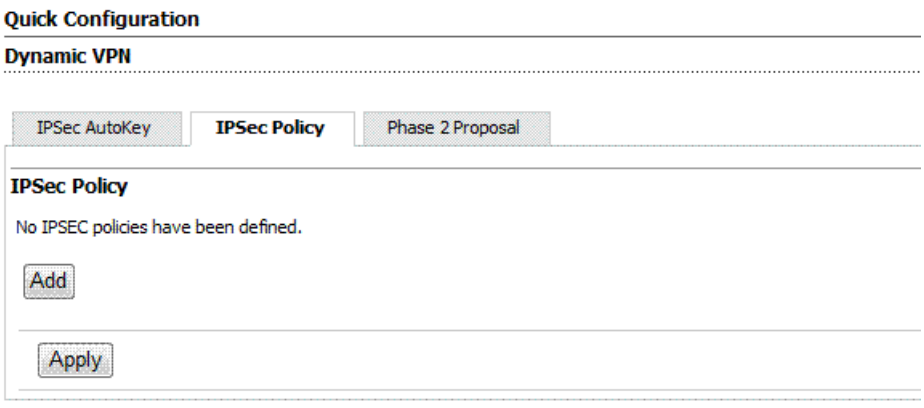

Figure 77 on page 464 shows the Quick Configuration page where you create a new IPsec policy.

Figure 77: IPsec Policy Quick Configuration Page – Configuring a Policy

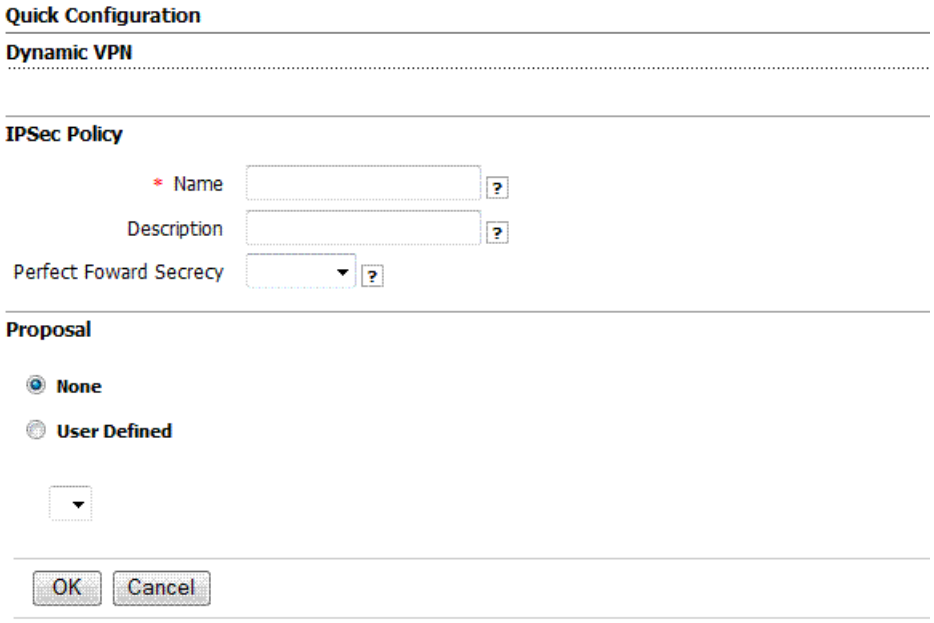

To configure an IPsec policy with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > IPsec Policy**.
2. Select the **IPsec Policy** tab if it is not selected.
3. To modify an existing policy, click the appropriate link in the **Name** column to go to the policy's configuration page. Or, select the policy from among those listed and click one of the following buttons:
   - To apply the configuration, click **Apply**.
   - To delete the configuration, click **Delete**.
4. To configure a new IPsec policy, click **Add**.
5. Fill in the options as described in Table 66 on page 465.
6. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration and return to the main Configuration page, click **Cancel**.

Table 66: IPsec Policy Configuration Options

Field	Function	Action
IPsec Policy		
Name	Name to identify the policy.	Enter a name.
Description	Description of the policy.	Enter a brief description of the policy.
Perfect Forward Secrecy	Method the device uses to generate the encryption key. Perfect Forward Secrecy generates each new encryption key independently from the previous key. ■ **group1**—Diffie-Hellman Group 1. ■ **group2**—Diffie-Hellman Group 2. ■ **group5**—Diffie-Hellman Group 5.	Select a method.
Proposal		
None	Do not use a proposal.	Click **None**.
User Defined	Use up to four Phase 2 proposals that you previously defined. If you include multiple Phase 2 proposals in the IPsec policy, use the same Diffie-Hellman group in all of the proposals.	Click **User Defined**, select a proposal (or proposals) from the pop-up menu, and click **Add**.

Configuring an IPsec Autokey—Quick Configuration (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly configure IPsec AutoKey. (For information about configuring an IPsec AutoKey using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Configuring IPsec AutoKey (Standard and Dynamic VPNs)" on page 390.)

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

Figure 78 on page 466 shows the Quick Configuration page where you can select an existing policy, or click **Add** to create a new one.

Figure 78: IPsec Autokey Quick Configuration Page – Adding an Autokey VPN

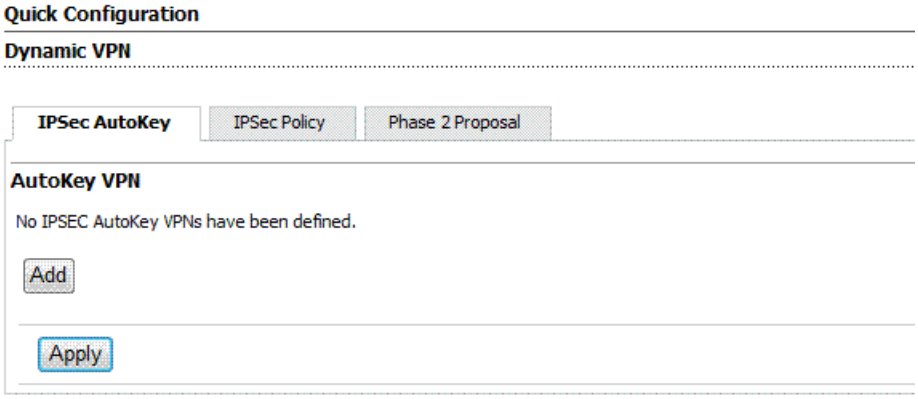

Figure 79 on page 467 shows the Quick Configuration page where you create a new IPsec AutoKey.

Figure 79: IPsec Autokey Quick Configuration Page – Configuring an Autokey VPN

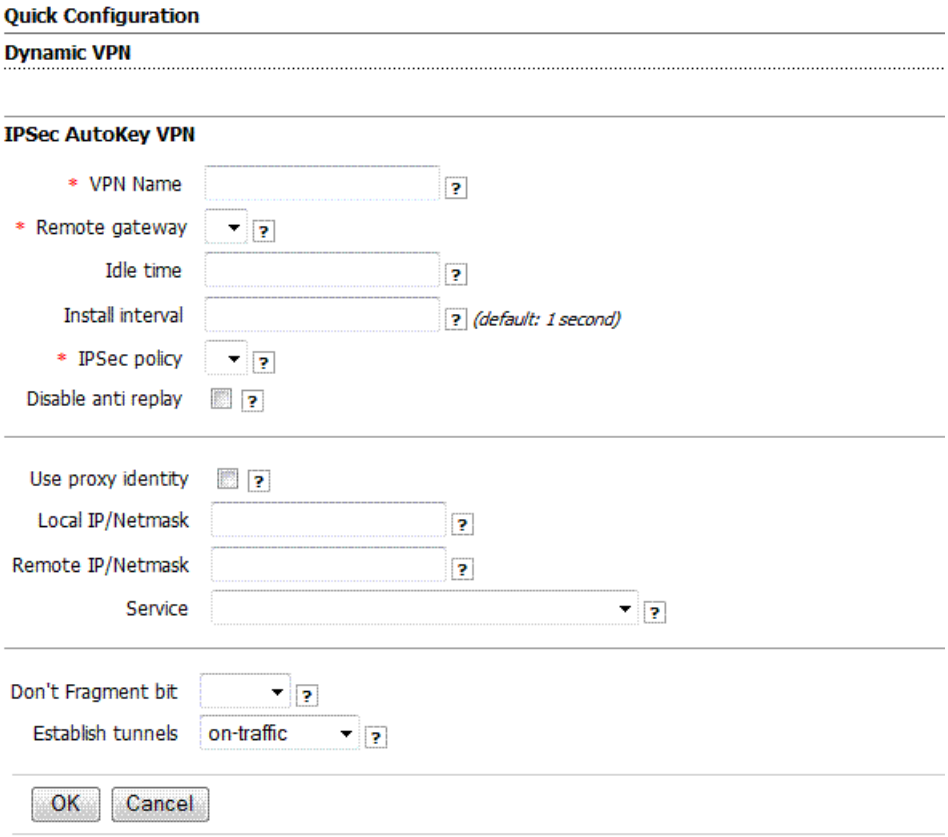

To configure an IPsec AutoKey with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > IPSec Autokey**.
2. Select the **IPSec AutoKey** tab if it is not selected.
3. To modify an existing IPsec AutoKey configuration, click the appropriate link in the **Name** column to go to the configuration page. Or, select the policy from among those listed and click one of the following buttons:
   - To apply the configuration, click **Apply**.
   - To delete the configuration, click **Delete**.
4. To configure a new IPsec AutoKey, click **Add**.
5. Fill in the options as described in Table 67 on page 468.
6. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration and return to the main Configuration page, click **Cancel**.

Table 67: IPsec AutoKey Configuration Options

Field	Function	Action
IPsec Autokey VPN		
VPN Name	Name to identify the IPsec AutoKey.	Enter a name.
Remote gateway	IKE gateway to associate with the IPsec AutoKey. An IKE gateway specifies a variety of IKE configuration options, including which IKE policy to use, how to identify endpoint computers during IKE exchanges, NAT options, dead peer detection options, and Xauth options.	Select a previously created IKE gateway.
Idle time	Maximum amount of time to allow a security association (SA) to remain idle before deleting it.	Specify a value between 60 and 999,999 seconds.
Install interval	Maximum number of seconds to allow the installation of a rekeyed outbound SA on the device.	Specify a value between 0 and 10 seconds.
IPsec policy	IPsec policy to associate with the IPsec AutoKey. An IPsec policy specifies the Diffie-Hellman group to use when generating encryption keys as well as the Phase 2 proposals to use.	Select a previously created IPsec policy.
Disable anti replay	Replay attacks occur when somebody intercepts a series of packets and uses them to flood the system or gain entry into a trusted system. Select this option to enable replay protection.	Click the check box to disable or enable. (Disabled by default.)
Use proxy identity	(Optional) Specify the IPsec proxy identity to use in IKE negotiations. The default behavior is to use the identities taken from the firewall policies.	Click the check box to disable or enable. (Disabled by default.)
Local IP/Netmask	Local IP address and subnet mask for the proxy identity.	Enter an IP address and subnet mask.
Remote IP/Netmask	Remote IP address and subnet mask for the proxy identity.	Enter an IP address and subnet mask.
Service	Service (port and protocol combination) to protect.	Select a service.
Don't fragment bit	Specify how the device should handle the Don't Fragment (DF) bit in the outer header. ■ **clear**—Clear (disable) the DF bit from the outer header. This is the default. ■ **copy**—Copy the DF bit to the outer header. ■ **set**—Set (enable) the DF bit in the outer header.	Choose an option.

Table 67: IPsec AutoKey Configuration Options *(continued)*

Field	Function	Action
Establish tunnels	Specify when to activate IKE: ■ **immediately**—Activate IKE immediately after the VPN is configured and changes are committed. ■ **on-traffic**—Activate IKE only when data traffic flows and must be negotiated.	Choose an option.

Creating a Client Configuration—Quick Configuration (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly create or modify a client configuration that downloads to the user's computer along with the Access Manager client. (For information about creating or modifying a client configuration using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see "Creating a Client Configuration (Dynamic VPNs)" on page 472.)

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

Figure 80: Client Configuration Quick Configuration Page

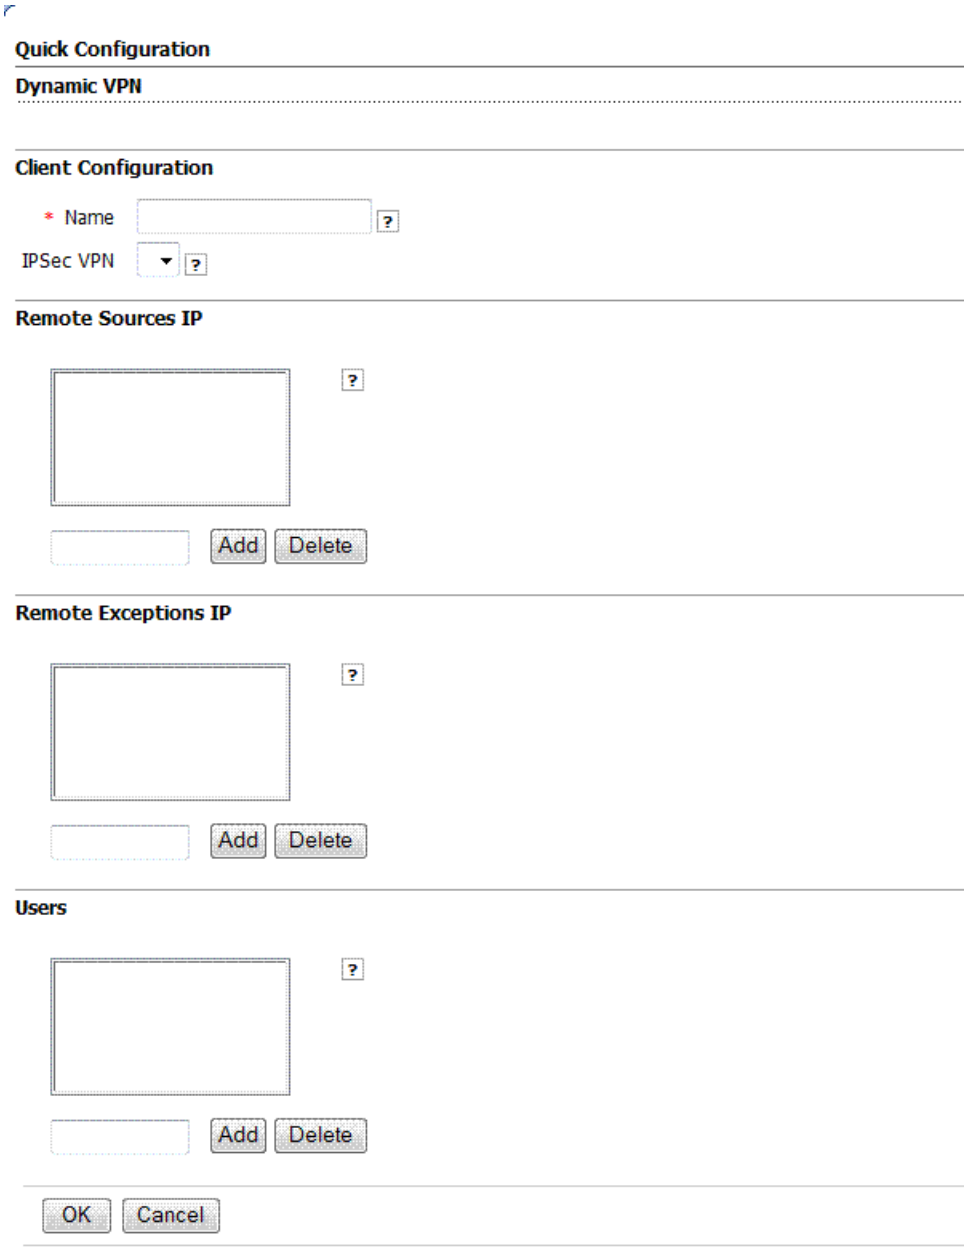

Figure 80 on page 470 shows the Quick Configuration page where you create a new client configuration.

To create or modify a client configuration with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > Global Settings**.
2. To modify an existing configuration, click the appropriate link in the **Client VPN Configuration** column to go to the configuration page. Or, select the client configuration from among those listed and click one of the following buttons:

- To apply the configuration, click **Apply**.
- To delete the configuration, click **Delete**.

3. To create a new client configuration, click **Add**.
4. Fill in the options as described in Table 68 on page 471.
5. Click one of the following buttons:
   - To apply the configuration, click **OK**.
   - To cancel the configuration and return to the main Configuration page, click **Cancel**.

Table 68: Client VPN Configuration Options

Field	Function	Action
Client Configuration		
Name	Name to identify the client configuration.	Enter a name.
IPSec VPN	IKE AutoKey configuration to use when establishing the VPN tunnel.	Select a previously created IKE AutoKey configuration.
Remote Address		
Remote Sources IP	IP address and netmask of a resource behind the firewall. Traffic to the specified resource will go through the VPN tunnel and therefore will be protected by the firewall's security policies. **NOTE:** The device does not validate that the IP/netmask combination that you enter here matches up with your security policies.	Enter an IP address and netmask and click **Add**.
Remote Exceptions IP	IP address and netmask of exceptions to the remote protected resources list.	Enter an IP address and netmask and click **Add**.
Users		
Users	List of users that can use this client configuration. **NOTE:** The server does not validate the names that you enter here, but the names must be the names that the users use to log onto the device when downloading the client.	Enter a user name and click **Add**.

Creating a Client Configuration (Dynamic VPNs)

Client configurations control which resources should be protected by the VPN configuration and specify which users can download the client.

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

In this example, you configure the device to protect the 10.100.100.0/24 resource, but specify that the 0.0.0.0./0, 1.1.1.1/24, and 0.0.0.0/32 resources are exempt. You specify that the dynamic VPN feature should use the ipsec-config VPN configuration to create the VPN tunnel. In addition, you add johndoe and janedoe to the list of users who can download the client.

To configure client configurations, use either the J-Web or the CLI configuration editor. (For information about configuring client configurations using J-Web Quick Configuration pages, see "Creating a Client Configuration—Quick Configuration (Dynamic VPNs)" on page 469.)

This topic covers:

- J-Web Configuration on page 472
- CLI Configuration on page 473

J-Web Configuration

To create a client configuration in J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Dynamic vpn, click **Configure** or **Edit**.
4. Next to Clients, click **Add New Entry** .
5. In the Name field, enter config1.
6. Enter ipsec-config in the **Ipsec vpn** field.
7. Next to Remote protected resources, click **Add New Entry** .
8. In the Resource ip field, enter 10.100.100.0/24 and click **OK**.
9. Next to Remote exceptions, click **Add New Entry** .
10. In the Exception ip field, enter 0.0.0.0/0 and click **OK**.
11. Next to Remote exceptions, click **Add New Entry** .
12. In the Exception ip field, enter 1.1.1.1/24 and click **OK**.

13. Next to Remote exceptions, click **Add New Entry** .
14. In the Exception ip field, enter 0.0.0.0/32 and click **OK**.
15. Next to User, click **Add New Entry** .
16. In the Name field, enter johndoe and click **OK**.
17. Next to User, click **Add New Entry** .
18. In the Name field, enter janedoe and click **OK**.
19. To save and commit the configuration, click **Commit**.

CLI Configuration

To create a client configuration in the CLI editor:

```
user@host# set security dynamic-vpn clients config1 remote-protected-resources
  10.100.100.0/24
user@host# set security dynamic-vpn clients config1 remote-exceptions 0.0.0.0/0,
  1.1.1.1/24, 0.0.0.0/32
user@host# set security dynamic-vpn clients config1 ipsec-vpn ipsec-config
user@host# set security dynamic-vpn clients config1 user johndoe, janedoe
```

Configuring Global Client Download Settings—Quick Configuration (Dynamic VPNs)

You can use J-Web Quick Configuration to quickly configure global dynamic VPN settings. (For information about configuring global dynamic VPN settings using the CLI editor or the standard J-Web configuration pages (not Quick Configuration pages), see “Configuring Global Client Download Settings (Dynamic VPNs)” on page 475.)

Before You Begin

For background information, read:

- “Dynamic VPN Overview” on page 433.
- “Configuring a Dynamic VPN—Overview” on page 449.

Figure 81: Global Settings Quick Configuration Page

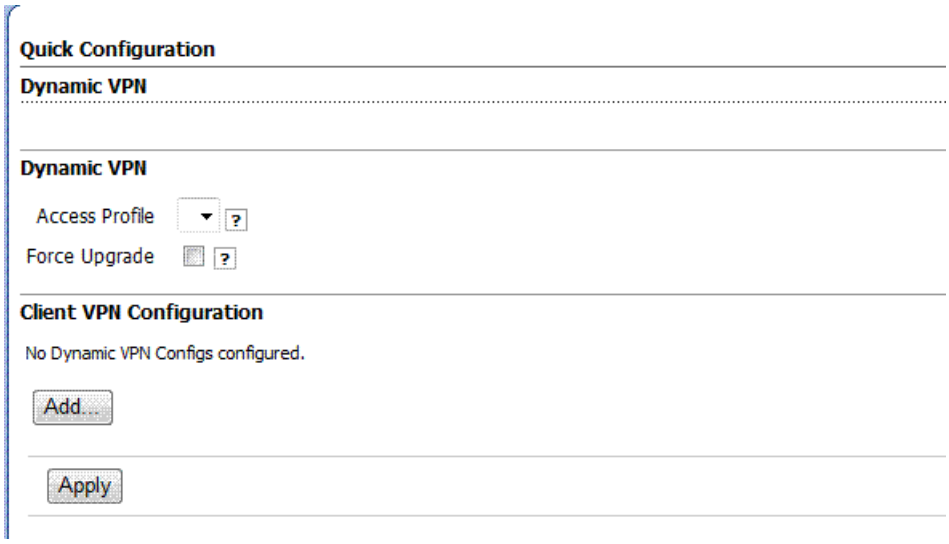

Figure 81 on page 474 shows the quick configuration page where you can modify global settings.

NOTE: For information about creating a client configuration, see "Creating a Client Configuration—Quick Configuration (Dynamic VPNs)" on page 469.

To create or modify global dynamic VPN settings with Quick Configuration:

1. Select **Configure > IPSec VPN > Dynamic VPN > Global Settings**.
2. Fill in the options as described in Table 68 on page 471.
3. To apply the configuration, click **OK**.

Table 69: Client VPN Configuration Options

Field	Function	Action
Dynamic VPN		

Table 69: Client VPN Configuration Options *(continued)*

Field	Function	Action
Access Profile	Access profile that controls the authentication of users who want to download Access Manager. (You will need to select these access profiles when configuring your IKE gateway and dynamic VPN global options. Note that you can use the same access profile to authenticate users in both cases, or you can use separate access profiles to authenticate downloads and VPN sessions.) **NOTE:** This Access Profile option does not control authentication for VPN sessions. For session authentication, use the Access Profile option on the IKE Gateway Quick Configuration page. For more information, see "Configuring an IKE Gateway—Quick Configuration (Dynamic VPNs)" on page 456.	Select a previously created access profile.
Force Upgrade	Enable this option if you want the setup program to automatically download the latest client and install it on the user's computer when the setup program detects a version mismatch between the client and server. Otherwise, the setup program prompts the user the upgrade the client when it detects a version mismatch, but does not force the upgrade. If the user does not choose to upgrade, the setup program will launch the existing client version that resides on the user's computer.	Click the check box to enable or disable. (Enabled by default.)

Configuring Global Client Download Settings (Dynamic VPNs)

Global dynamic VPN settings enable you to control authentication and upgrade options for Access Manager. The following example-based instructions show how to configure global dynamic VPN settings.

Before You Begin

For background information, read:

- "Dynamic VPN Overview" on page 433.
- "Configuring a Dynamic VPN—Overview" on page 449.

In this example, you configure the device to use the "remote-users" access profile to control the authentication of users who try to download Access Manager.You also configure the setup program to automatically upgrade the client on users' machines when it detects that a more current version of the client is available on the server.

NOTE: For information about creating a client configuration, see “Creating a Client Configuration (Dynamic VPNs)” on page 472.

To configure dynamic VPN global settings, use either the J-Web or the CLI configuration editor. (For information about configuring global settings using J-Web Quick Configuration pages, see “Configuring IPSec VPN Global Settings in J-Web (Standard VPNs)” on page 372.)

This topic covers:

- J-Web Configuration on page 476
- CLI Configuration on page 476

J-Web Configuration

To configure global dynamic VPN settings using J-Web:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Dynamic vpn, click **Configure** or **Edit**.
4. From the Access profile list, select **remote-users**.
5. Select the **Force Upgrade** checkbox.
6. To save and commit the configuration, click **Commit**.

CLI Configuration

To configure global dynamic VPN settings using the CLI editor:

```
user@host# set security dynamic-vpn clients remote-users
user@host# set security dynamic-vpn force-upgrade
```

Chapter 14

NetScreen-Remote VPN Client

The Juniper Networks NetScreen-Remote VPN client is a virtual private network (VPN) client that you can install on a PC or laptop computer to send and receive secure communications over the Internet. NetScreen-Remote client is certified by the International Computer Security Association (ICSA) as an IPsec-compliant VPN solution.

When NetScreen-Remote client operates on an unprotected public network, such as the Internet, it can create a VPN tunnel between your PC or laptop computer and a J Series Services Router running JUNOS Software. The tunnel secures IPsec traffic sent across a public or private TCP/IP network. NetScreen-Remote client allows you to specify an internal network IP address to be sent for client-to-gateway communications.

NetScreen-Remote client starts automatically each time the computer starts and runs transparently behind other software applications.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- System Requirements for NetScreen-Remote Client Installation on page 477
- Installing the NetScreen-Remote Client on a PC or Laptop on page 478
- Configuring the Firewall on the Router on page 483
- Configuring the PC or Laptop on page 486
- Logging In to the NetScreen Remote Client on page 496

System Requirements for NetScreen-Remote Client Installation

You can install the NetScreen-Remote client in the environment described in Table 70 on page 477.

Table 70: System Requirements

PC-compatible computer	Pentium processor or its equivalent
Operating system	Microsoft Windows 2000 Professional Microsoft Windows XP Professional or Home Edition.

Table 70: System Requirements *(continued)*

Minimum RAM	64 MB RAM for Windows 2000 or Windows XP.
Available hard disk space	Minimum 5 MB, maximum 35 MB
Software installation	CD-ROM drive, network drive, or website
Communications protocol	IPsec and IKE L2TP with Windows 2000 (*optional*) Native Microsoft TCP/IP
Dial-up connections	Internal or external modem connection—an analog, DSL, or cable modem line connected to your PC or laptop through a serial or USB port Native Microsoft Dial-up Networking PPPoE drivers Compatible with America Online (AOL) 6.0 or later
Network connections	Ethernet Wireless Ethernet (802.11a/b)
Help file viewer	Microsoft Internet Explorer 4.0 or later

Installing the NetScreen-Remote Client on a PC or Laptop

Use one of three installation methods to start installing the NetScreen-Remote client on your PC or laptop computer. Then complete the installation.

Before You Begin
Verify that your PC or laptop meets the necessary system requirements. See " System Requirements."

This topic covers:

- Starting NetScreen-Remote Client Installation on page 478
- Completing NetScreen-Remote Client Installation on page 480

Starting NetScreen-Remote Client Installation

Use one of the following methods to begin installing the NetScreen-Remote Client on your PC or laptop computer:

- Starting Installation from a CD-ROM on page 479
- Starting Installation from a Network Share Drive on page 479
- Starting Installation from a Website on page 479

Starting Installation from a CD-ROM

Figure 82 on page 479 shows the Install menu.

Figure 82: Select Install

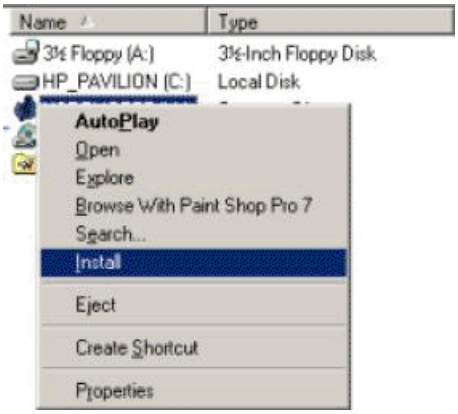

To install the NetScreen-Remote Client from a CD-ROM:

1. Make sure all applications are closed.
2. Insert the NetScreen-Remote client CD into the CD-ROM drive.
3. Right-click the CD-ROM drive on your computer—for example, **D:**. (This example uses D to designate the CD-ROM drive. Your computer might be configured differently.)
4. The menu shown in Figure 82 on page 479 appears.
5. To install the NetScreen-Remote client, select **Install** from the menu.
6. Go on to "Completing NetScreen-Remote Client Installation" on page 480.

Starting Installation from a Network Share Drive

To install the NetScreen-Remote Client from a network share drive:

1. Map to the network drive.
2. Locate the NetScreen-Remote client files.
3. To run the NetScreen-Remote client setup application, double-click **setup.exe**.
4. Go on to "Completing NetScreen-Remote Client Installation" on page 480.

Starting Installation from a Website

To install the NetScreen-Remote Client from a website:

1. Locate the NetScreen-Remote client files on the Juniper Networks website http://juniper.net.
2. Select **setup.exe**, and download the file.
3. After the file is downloaded, unzip the file to C:\temp.
4. To run the NetScreen-Remote client setup application, double-click **setup.exe**.
5. Go on to "Completing NetScreen-Remote Client Installation" on page 480.

Completing NetScreen-Remote Client Installation

Figure 83 on page 480 shows ther Welcome Screen page.

Figure 83: Welcome Screen

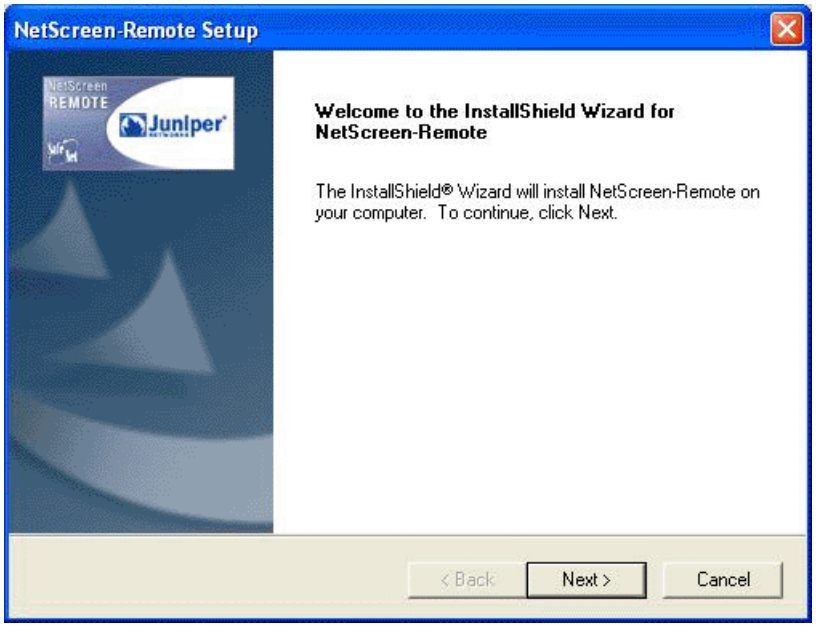

Figure 84 on page 481 shows the License Agreement page.

Figure 84: License Agreement

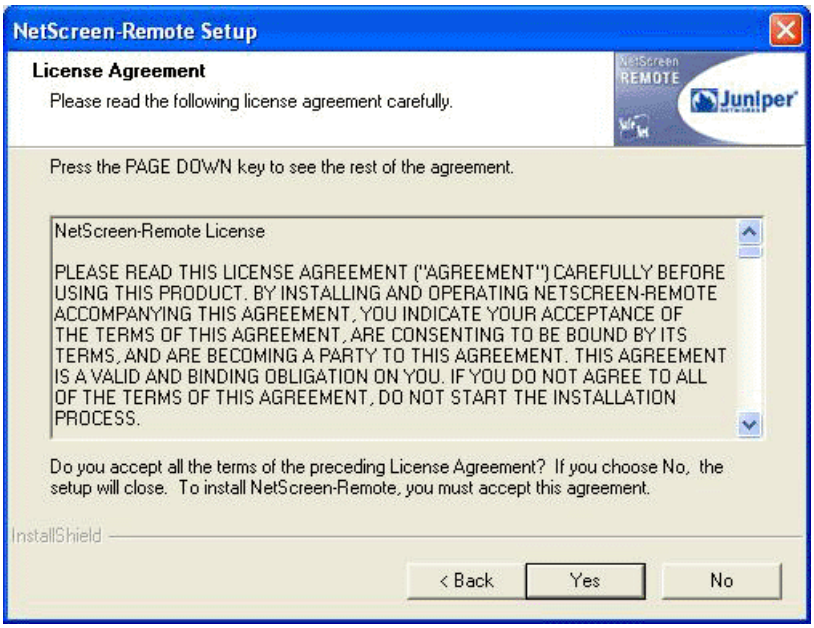

Figure 85 on page 481 shows the Installation Setup Type page.

Figure 85: Installation Setup Type

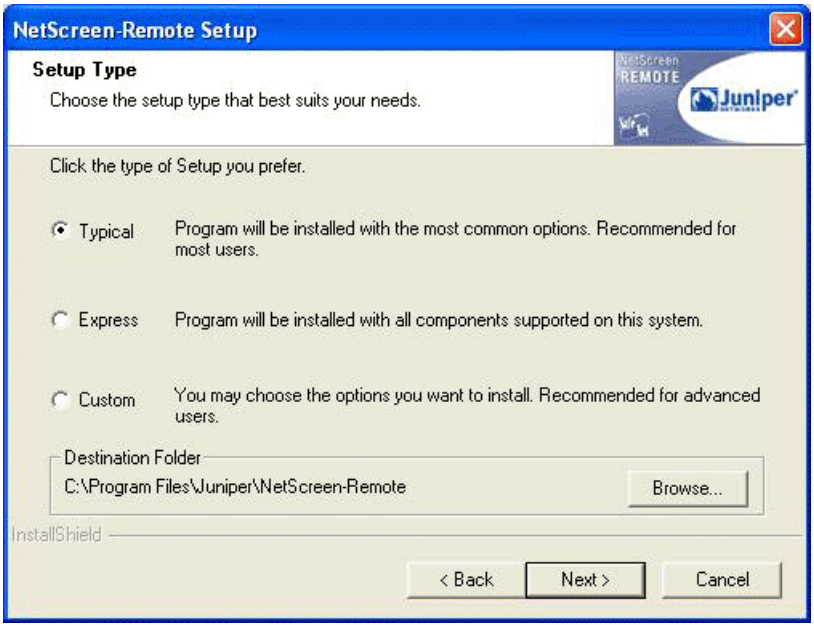

Figure 86 on page 482 shows the Start Copying Files page.

Figure 86: Start Copying Files

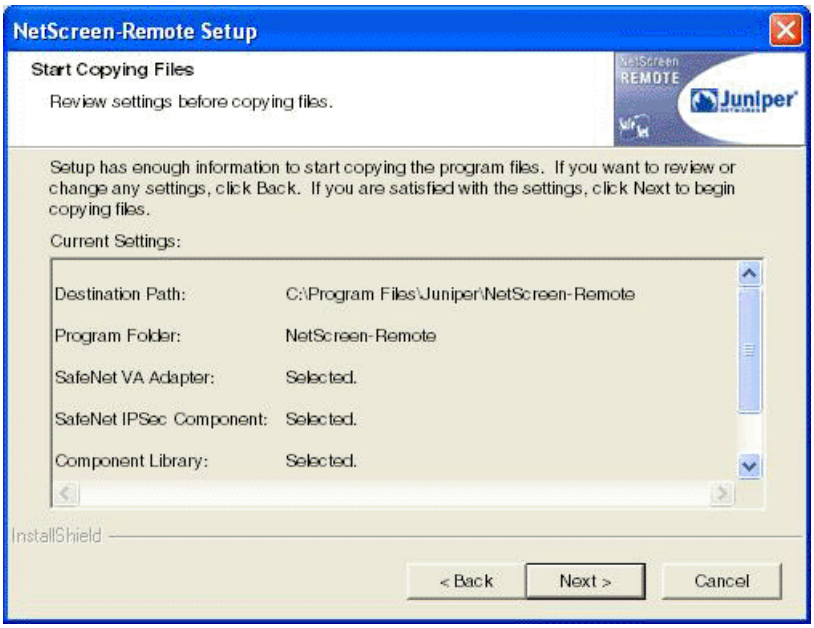

After you have finished initial installation (see “Starting NetScreen-Remote Client Installation” on page 478), complete the NetScreen-Remote client installation:

1. When the InstallShield Wizard starts as shown in Figure 83 on page 480, click **Next**.
2. On the Software License Agreement page (Figure 84 on page 481), read the agreement and click **Yes** to continue.
3. In the Setup Type dialog box (Figure 85 on page 481), select one of the following options:
   - **Typical**—Recommended for most users. This option installs all VPN client components in the default destination folder C:\Program Files\NetScreen\NetScreen-Remote.
   - **Express**—Installs only the components that the system supports.
   - **Custom**—Enables you to individually select the components to install.

   To install the NetScreen-Remote client in the default destination folder C:\Program Files\NetScreen\NetScreen-Remote, click **Next**.
4. Verify your selections in the window that appears (Figure 86 on page 482), and click **Next**.

   The NetScreen-Remote client files are copied to the program folder that you specified. A progress bar shows you the status of the installation.

Your computer automatically reboots after a successful installation. To cancel the reboot process, click **Cancel** in the progress bar before device timeout.

Configuring the Firewall on the Router

For the NetScreen-Remote client to have a VPN tunnel to the router running JUNOS Software, you must configure a firewall on the router.

Before You Begin

Install the NetScreen-Remote client on your PC or laptop. See "Installing the NetScreen-Remote Client on a PC or Laptop" on page 478.

This topic covers:

- Firewall Configuration Overview on page 483
- Configuring a Security Zone on page 483
- Configuring a Tunnel Interface on page 484
- Configuring an Access Profile for XAuth on page 485
- Configuring an IKE Gateway on page 485
- Configuring Policies on page 486

Firewall Configuration Overview

Perform the following tasks to configure the JUNOS Software firewall on the router:

1. Configure a security zone where the VPN client is to terminate. For more information, see "Configuring a Security Zone" on page 483.
2. Configure a tunnel interface to identify the VPN. For more information, see "Configuring a Tunnel Interface" on page 484.
3. Configure the eXtended Authentication (XAuth) settings. For more information, see "Configuring an Access Profile for XAuth" on page 485.
4. Configure an IKE gateway and phase 1 proposal. For more information, see "Configuring an IKE Gateway" on page 485.
5. Configure firewall policies to permit or deny traffic. For more information, see "Configuring Policies" on page 486.

Configuring a Security Zone

You must configure a security zone where the VPN client is to terminate. Terminating the VPN client in a separate security zone allows you to filter traffic coming from the VPN client and apply advanced security functions such as deep-inspection, traffic shaping, and policy actions. For more information on these features, see "Attack Detection and Prevention" on page 761 and "Security Policies Overview" on page 69.

To configure a security zone:

1. Select Configure>Security>Zones.

2. Click Add to create new zones.
3. Create a new zone called vpn. For more information on creating zones, see "Configuring Security Zones—Quick Configuration" on page 51.

Configuring a Tunnel Interface

You must create a tunnel interface and bind it to the security zone (vpn) that you created in "Configuring a Security Zone" on page 483. Then you can bind the VPN to the tunnel interface in the VPN configuration.

Figure 87 on page 484 shows the Interface List page.

Figure 87: Interface List

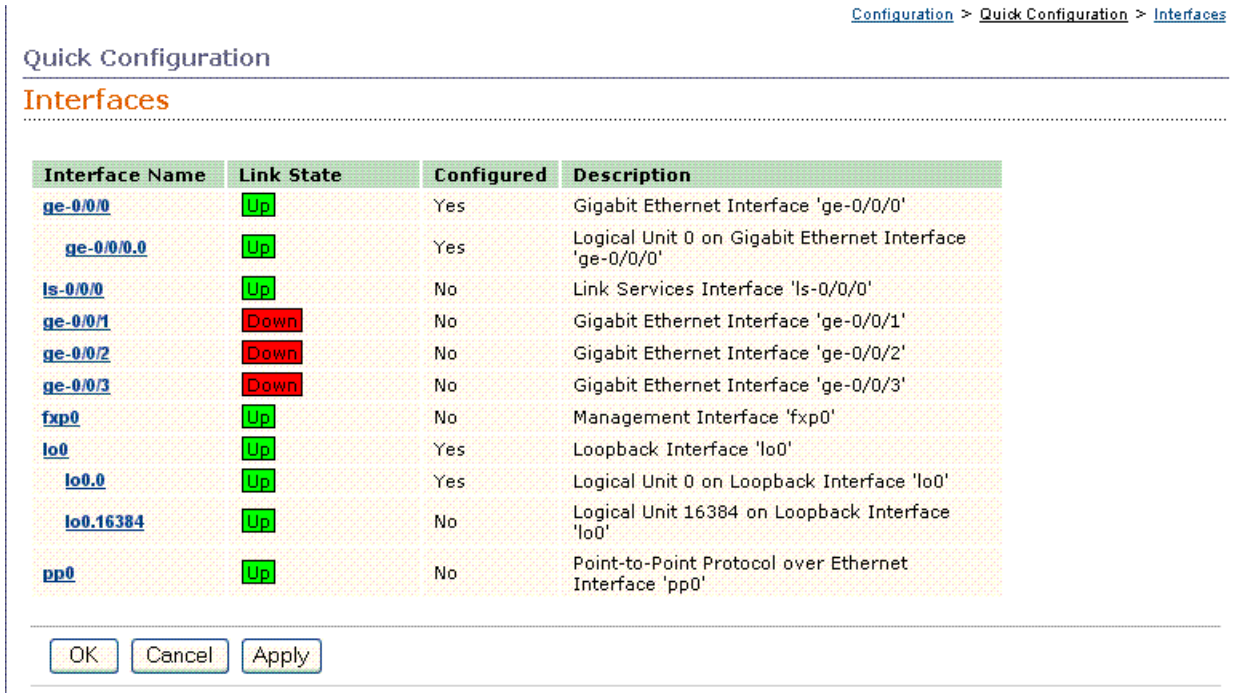

To configure a tunnel interface to be used for the VPN:

1. Select Configure>Interfaces. A list of the network interfaces available on the router appears, as shown in Figure 87 on page 484. The third column indicates whether the interface has been configured.
2. Configure properties for a network interface by selecting the interface name and following the instructions in "Configuring a Gigabit Ethernet Interface—Quick Configuration" on page 63.

   To configure interfaces other than Gigabit Ethernet interfaces, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Configuring an Access Profile for XAuth

The router uses an access profile to verify extended authentication (XAuth) for remote users trying to access a virtual private network (VPN) tunnel. Before you configure XAuth as part of the IKE gateway, you must configure an access profile for remote users. You refer to this access profile from the IKE gateway configuration. When you configure an access profile for XAuth, you must specify RADIUS as the type of authentication server (authentication-order radius).

The following sample commands create an access profile called xauth that sets RADIUS as the authentication method and specifies the IP address of the RADIUS server, the number of the port on which to contact the server (1812 by default), and the secret password shared between the router and the server:

```
user@host# set access profile xauth authentication-order radius
user@host# set access profile xauth radius-server 10.157.90.244 port 1812
user@host# set access profile xauth radius-server 10.157.90.244 secret " $9
```

Configuring an IKE Gateway

IKE configuration is needed to identify the clients using the JUNOS Software firewall. Usually you configure one IKE user per client installed, but because XAuth is used as an extra layer of authentication, the shared IKE ID is used to build Phase 1 of the VPN tunnel.

Perform the following tasks to complete the configuration:

1. Configure an IKE Phase 1 proposal.
2. Configure an IKE policy. In Phase 1 IKE policy configuration, you must set the mode in which the Phase 1 channel is negotiated, specify the type of key exchange to be used, and reference the Phase 1 proposal.
3. Configure the IKE gateway. In Phase 1 gateway configuration, you must configure a gateway and reference the Phase 1 policy. You must configure the router to use a shared IKE ID and limit the number of concurrent connections.

   The following sample commands configure an IKE gateway named jsr_gateway for up to 100 concurrent users. The remote (dynamic) users have an unspecified IP address but use the e-mail address test@juniper.net as their shared IKE ID. The trusted external interface for traffic from these users is fe-3/0/20. The gateway uses extended authentication (XAuth) to verify the authenticity of these users with the previously created access profile xauth. (See "Configuring an Access Profile for XAuth" on page 485)

   ```
   user@host# set security ike gateway jsr_gateway dynamic user-at-hostname "
     test@juniper.net"
   user@host# set security ike gateway jsr_gateway dynamic connections-limit 100
   user@host# set security ike gateway jsr_gateway dynamic ike-user-type
     shared-ike-id
   user@host# set security ike gateway jsr_gateway external-interface fe-3/0/2
   user@host# set security ike gateway jsr_gateway xauth access-profile xauth
   ```

NOTE: Make sure you select the access profile that you set using the CLI statements in "Configuring an Access Profile for XAuth" on page 485.

Configuring Policies

Now that you have the VPN clients terminating in the vpn zone, configure a policy that allows traffic from the VPN zone to any destination zone. Make sure that the policy configuration is not any any, and is an IPsec configuration. For example, the following commands set a security policy for VPN match from zone A to zone B and source and destination address as any and application as any. If these conditions match, then permit traffic through the tunnel ipsec-vpn test.

```
user@host# set security policies from-zone A to-zone B policy VPN match
   source-address any
user@host# set security policies from-zone A to-zone B policy VPN match
   destination-address any
user@host# set security policies from-zone A to-zone B policy VPN match application
   any
user@host# set security policies from-zone A to-zone B policy VPN then permit tunnel
   ipsec-vpn test
```

To configure security policies with J-Web Quick Configuration:

1. Select Configure>Security>Policy>FW Policies.
2. Select the Default Policy Action, Deny All or Permit All.
3. In the From Zone and To Zone boxes, select the zone direction. You must have preconfigured the security zones for which you want to set policies. For more information, see "Configuring Security Zones—Quick Configuration" on page 51.

Configuring the PC or Laptop

Before You Begin
Configure a firewall on the J Series router you are creating a tunnel to. See "Configuring the Firewall on the Router" on page 483.

You can use a preshared key operation when the NetScreen-Remote client has either a fixed or dynamically assigned IP address. There are three steps to setting up a NetScreen-Remote client for a VPN tunnel with a preshared key:

1. "Creating a New Connection" on page 487
2. "Creating the Preshared Key" on page 490
3. "Defining the IPsec Protocols" on page 491

Creating a New Connection

First initiate a new connection. Then name the connection, define it as secure, and determine the identification and location of the other end of the eventual VPN tunnel.

Figure 88 on page 487 shows the NetScreen-Remote Client Icon in the Task Bar page.

Figure 88: NetScreen-Remote Client Icon

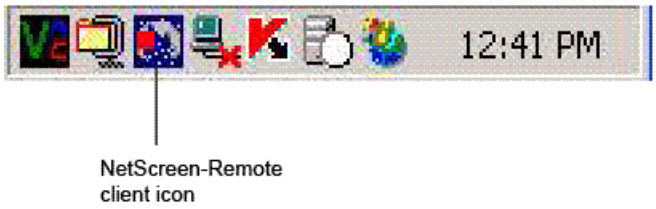

Figure 89 on page 487 shows the Security Policy Editor page.

Figure 89: Security Policy Editor

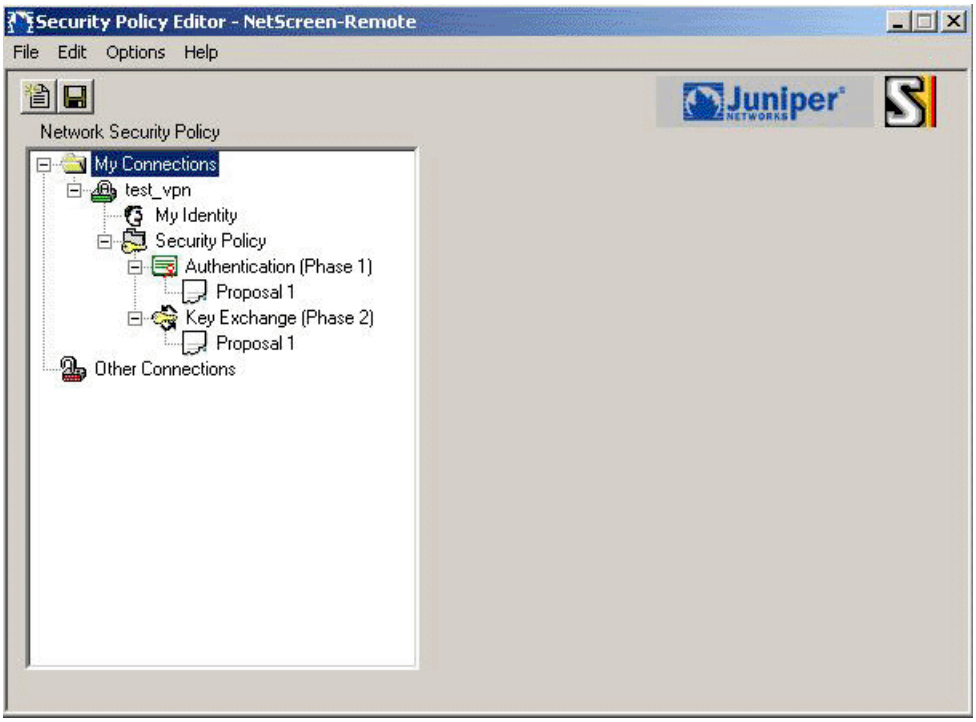

New Connection appears in the Network Security Policy list, as shown in Figure 90 on page 488.

Figure 90: Configure Connection

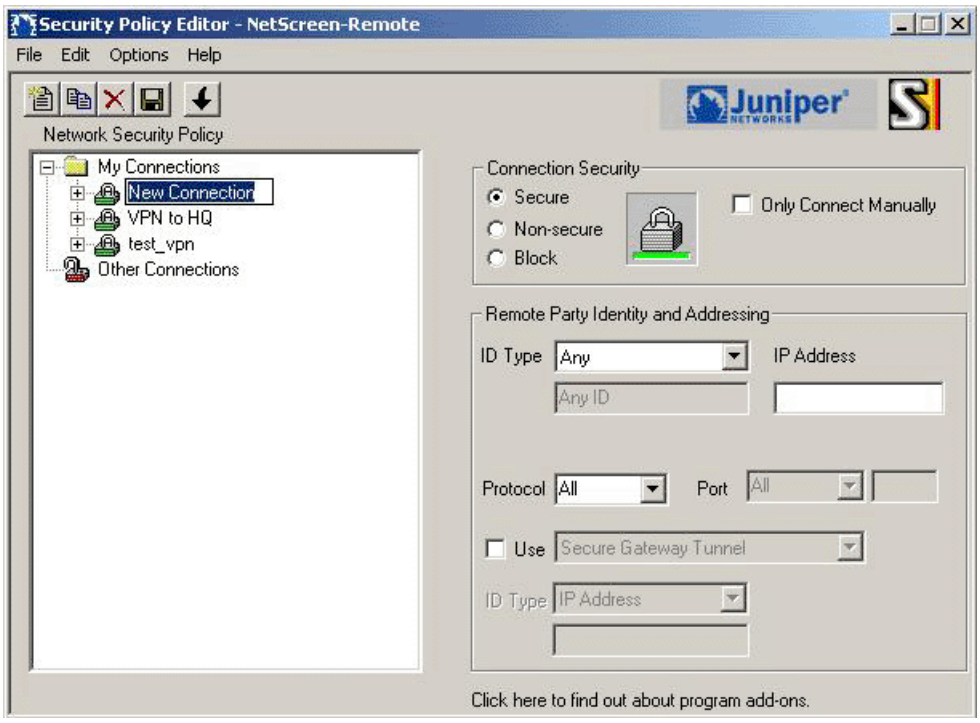

Figure 91 on page 489 shows the Configuring the New Connection page.

Figure 91: Configuring the New Connection

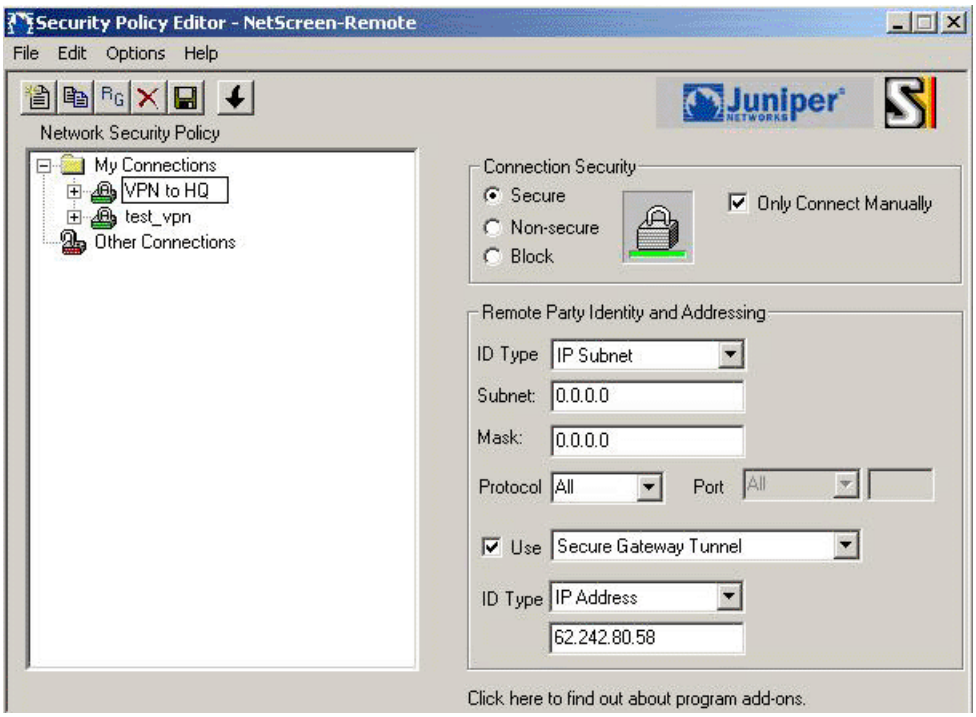

1. Double-click the NetScreen-Remote icon in the Windows taskbar (Figure 88 on page 487). The Security Policy Editor screen appears (Figure 89 on page 487).
2. Click the New Connection icon to create a new connection.
3. Give the new connection a unique name—for example, VPN to HQ.
4. In the Connection Security area (to the right of the Network Security Policy list), select Secure.
5. In the Remote Party Identity and Addressing area, select an identifier for the other party from the ID Type list, and enter the required information.
6. Choose either IP Address or IP Subnet. Other choices will not work.
7. Select the protocol you want to use for the Connection. The default is All.
   - All—This choice allows the connection to use any IP protocol.
   - TCP—Transmission Control Protocol, the protocol that controls data transfer on the Internet
   - UDP—User Datagram Protocol, a protocol within the TCP/IP protocol suite that provides very few error recovery services (for example, a lost packet is simply ignored) and is used primarily for broadcasting

- ICMP—Internet Control Message Protocol, a protocol tightly integrated with the Internet Protocol (IP) that supports packets containing error, control, and informational messages related to network operations
- GRE—Generic Routing Encapsulation, a protocol that encapsulates the packets of one kind of protocol within GRE packets, which can then be contained within the packets of another kind of protocol

8. If you are using tunnel mode to connect to a J Series device running JUNOS Software, select Connect using Secure Gateway Tunnel.

   The Secure Gateway Tunnel ID Type and IP Address fields are enabled.

9. Select IP Address as an identifier for the other party from the ID Type list and enter the IP address. See Figure 91 on page 489.

Creating the Preshared Key

After you have created a new connection called VPN to HQ, create the preshared key to be used in identifying the communicating parties during the Phase 1 negotiations.

Figure 92 on page 490 shows the My Identity and Internet Interface page.

Figure 92: My Identity and Internet Interface

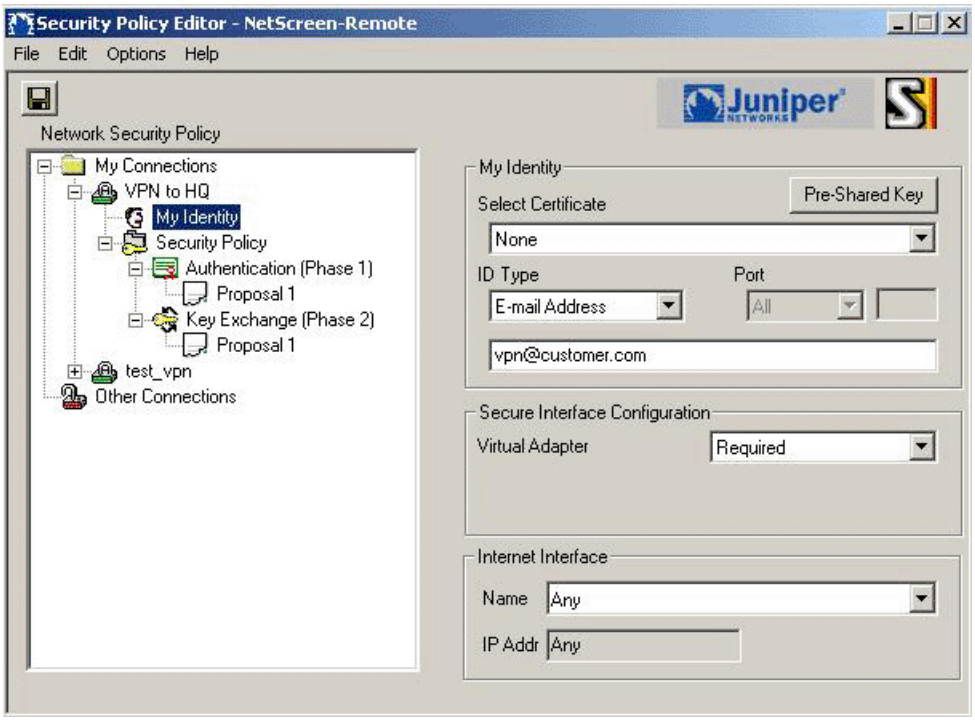

Figure 93 on page 491 shows the Per-Shared Key Dialog Box page.

Figure 93: Pre-Shared Key Dialog Box

1. Double-click the **VPN to HQ** icon from the Security Policy list in the left panel.

   My Identity and Security Policy icons appear in the Network Security Policy list.
2. Click **My Identity**. The My Identity and Internet Interface areas appear in the right panel (Figure 92 on page 490).
3. Select **None** from the Select Certificate drop-down list.
4. From the ID Type drop-down list, select **E-mail Address** and type vpn@customer.com as the ID for the IKE user.
5. Click **Pre-Shared Key**. The Pre-Shared Key dialog box appears (see Figure 93 on page 491).
6. Click the **Enter Key** to enable the Pre-Shared Key field.
7. Type a key with a length between 8 and 58 characters. A longer key length results in stronger encryption.
8. Click **OK** to save the entry.

Defining the IPsec Protocols

The Security Policy area appears on the right, and the Authentication (Phase 1) icon and Key Exchange (Phase 2) icon appear in the Network Security Policy list, as shown in Figure 94 on page 492.

Figure 94: Security Policy

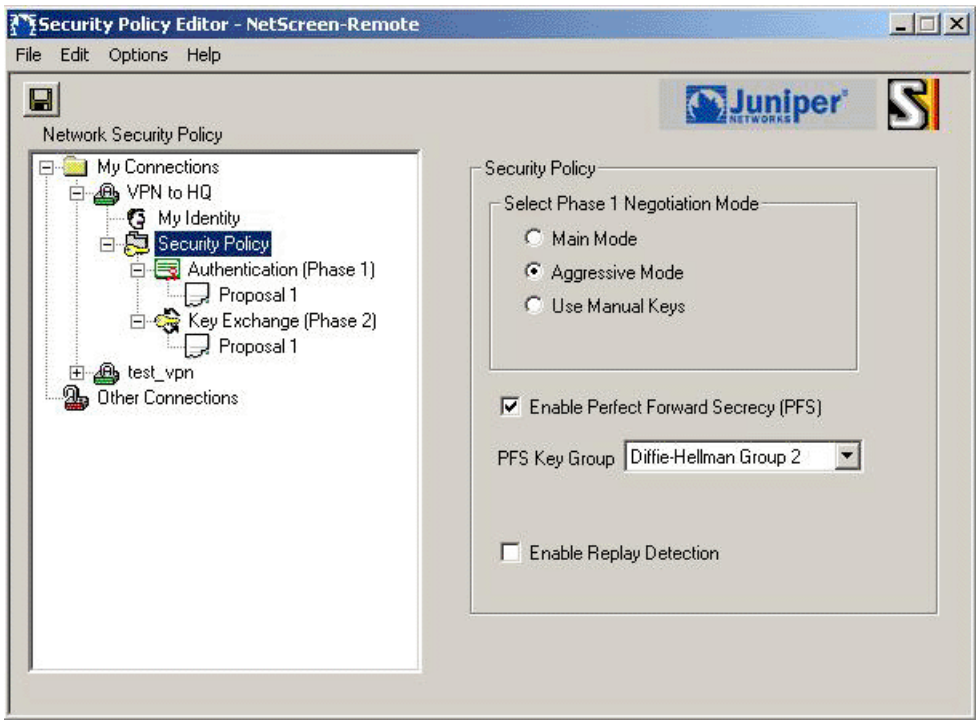

Figure 95 on page 493 shows the Algorithims Area page.

Figure 95: Algorithms Area

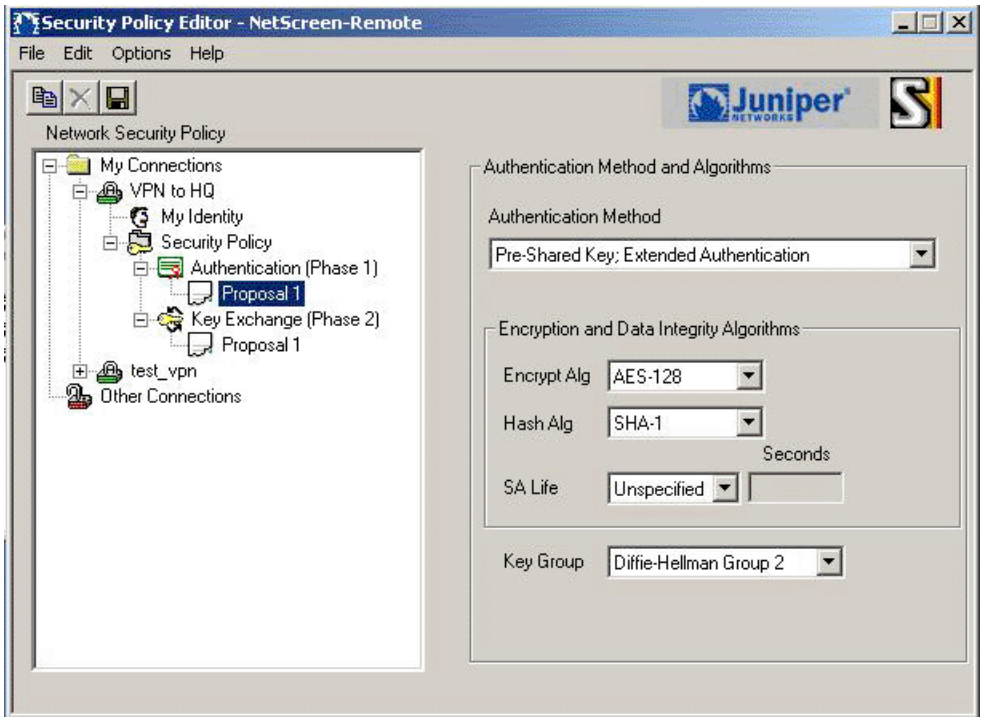

Figure 96: IPsec Protocols Area

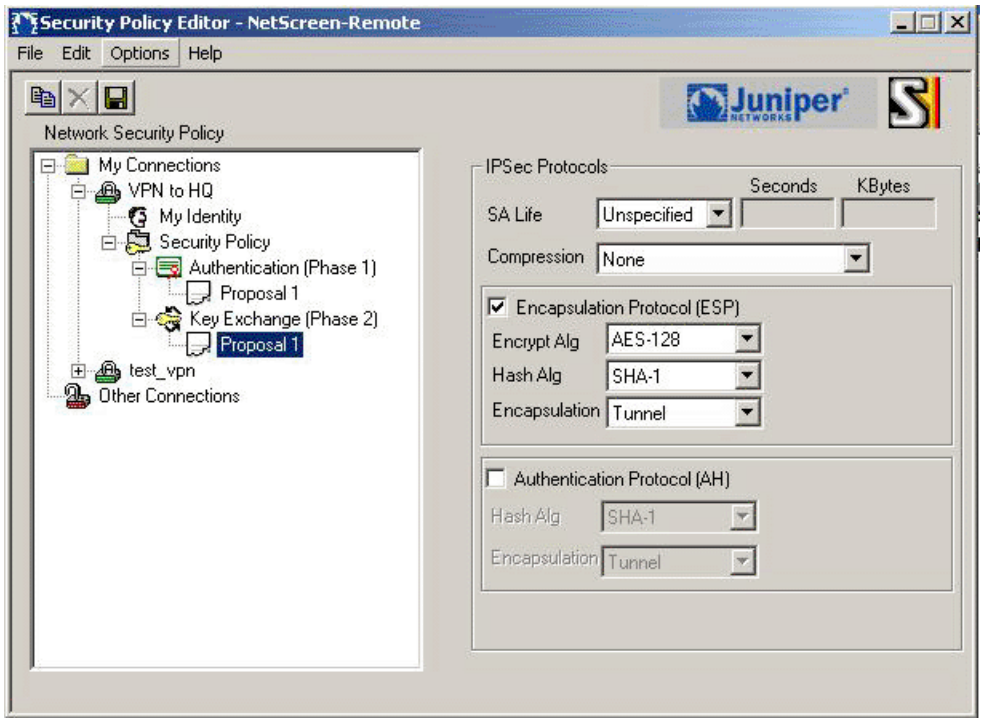

To define the Internet Protocol Security (IPsec) protocols for securing the VPN tunnel:

1. Double-click **Security Policy** in the Network Security Policy list.
2. Select **Aggressive Mode** in the Security Policy area.
3. Select **Enable Perfect Forward Secrecy (PFS)**. PFS allows generation of a new encryption key that is independent from and unrelated to the preceding key.
4. In the PFS Key Group drop-down list, select **Diffie-Hellman Group 2**.
5. In the Security Policy List (left panel), select **Authentication (Phase 1)**. Proposal 1 appears below the Authentication (Phase 1) icon.
6. Select **Proposal 1** to display the Authentication Method and Algorithms area, as shown in Figure 92 on page 490.
7. Select **Pre-Shared Key; Extended Authentication** from the Authentication Method. This allows you to use XAuth.

NOTE: XAuth must also be enabled on the J Series device running JUNOS Software. XAuth allows password-prompt authentication in addition to a preshared key. If enabled, you are prompted for a password when initiating a VPN. See “Configuring an Access Profile for XAuth” on page 485 and “Configuring an IKE Gateway” on page 485 for more information on configuring XAuth.

8. In the Authentication and Algorithms area, define the Encryption Algorithm AES-128 and the Hash Algorithm SHA-1. See Table 71 on page 495 for brief descriptions of these protocols.
9. From the Key Group drop-down list, select **Diffie-Hellman Group 2**.
10. In the left panel, double-click the **Key Exchange Phase (2)** icon. Proposal 1 appears below the icon.
11. Select **Proposal 1** to display the IPsec Protocols area as shown in Figure 96 on page 494.
12. In the IPsec Protocols area, define the SA Life (the lifetime of the security association) in either seconds or bytes, or leave it as **Unspecified**.

NOTE: Unspecified lifetimes (Phase I and II) cause the NetScreen-Remote client to accept the values proposed by the router.

13. Select Encapsulation Protocol (ESP). ESP provides encryption, authentication, and an integrity check for IP datagrams.
14. Select the encryption algorithm AES-128, the hash algorithm SHA-1, and Tunnel.for the encapsulation

NOTE: If you select the Connect using Secure Gateway Tunnel check box when defining Remote Party Identity and Addressing, the encapsulation method must be Tunnel—no other option is available.

15. Click Save in the toolbar, or choose Save Changes from the File menu.

    The configuration for the NetScreen-Remote end of an eventual VPN tunnel using a preshared key is complete.

Table 71: Encryption and Hash Algorithms

DES	Data Encryption Standard. A cryptographic block algorithm with a 56-bit key.
Triple DES	A more powerful version of DES in which the original DES algorithm is applied in three rounds, using a 168-bit key.
AES protocols	Advanced encryption standard. These protocols provide maximum security for the key. The higher the AES value, the more secure the key is. AES values can be AES-128, the least secure, AES-192, medium security, and AES-256, the most secure.
MD5	Message Digest version 5. An algorithm that produces a 128-bit message digest or hash from a message of arbitrary length. The resulting hash is used, like a fingerprint of the input, to verify authenticity.
SHA-1	Secure Hash Algorithm-1. An algorithm that produces a 160-bit hash from a message of arbitrary length. SHA-1 is generally regarded as more secure than MD5 because of the larger hashes it produces.

Logging In to the NetScreen Remote Client

Before You Begin
Configure your computer to connect to the router. See “Configuring the PC or Laptop” on page 486.

Figure 97 on page 496 shows the Login page.

Figure 97: Login

Figure 98 on page 496 shows the Successfully Connected page.

Figure 98: Successfully Connected

Now that you have configured the NetScreen-Remote client connection on your PC or laptop, you can establish connection:

1. From the Windows taskbar, right-click the NetScreen Remote icon and choose Connect. Select your VPN VPN to HQ. If you configured correctly, the login screen appears (Figure 97 on page 496).
2. Type the username and password configured for the XAuth user. See "Configuring an Access Profile for XAuth" on page 485 and "Configuring an IKE Gateway" on page 485 for more information.

NOTE: The username and password must match the RADIUS username and password configured on the RADIUS server.

3. If the connection is successful, the dialog box in Figure 98 on page 496 appears.
4. Click OK.

Part 7
Intrusion Detection and Prevention

- IDP Policies on page 501
- IDP Signature Database on page 571
- IDP Application Identification on page 589
- IDP SSL Inspection on page 601
- Performance and Capacity Tuning on page 607
- IDP Logging on page 609

Chapter 15

IDP Policies

The JUNOS Software Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through an IDP-enabled device. It allows you to define policy rules to match a section of traffic based on a zone, network, and application, and then take active or passive preventive actions on that traffic.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506
- Understanding IDP Rule Match Conditions on page 507
- Understanding IDP Rule Objects on page 508
- Understanding IDP Rule Actions on page 510
- Understanding IDP Rule IP Actions on page 512
- Understanding IDP Rule Notifications on page 514
- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519
- IDP Policies—Quick Configuration on page 522
- Inserting a Rule in the Rulebase on page 530
- Deactivating and Reactivating Rules in a Rulebase on page 532
- Understanding Application Sets on page 533
- Configuring Applications or Services for IDP on page 533
- Configuring Application Sets for IDP on page 535
- Enabling IDP in a Security Policy on page 536
- Understanding IDP Terminal Rules on page 540
- Setting Terminal Rules in Rulebases on page 541
- Understanding Custom Attack Objects on page 544
- Understanding Protocol Decoders on page 560

- Configuring Signature-Based Attacks on page 562
- Configuring Protocol Anomaly-Based Attacks on page 565
- Specifying IDP Test Conditions for a Specific Protocol on page 567
- Configuring DSCP in an IDP Policy on page 568

IDP Policies Overview

An IDP policy defines how your device handles the network traffic. It allows you to enforce various attack detection and prevention techniques on traffic traversing your network.

A policy is made up of *rulebases* and each rulebase contains a set of *rules*. You define rule parameters, such as traffic match conditions, action, and logging requirements, then add the rules to rule bases. After you create an IDP Policy by adding rules in one or more rulebases, you can select that policy to be the active policy on your device.

JUNOS Software allows you to configure multiple IDP policies, but a device can have only one active IDP policy at a time. You can install the same IDP policy on multiple devices, or you can install a unique IDP policy on each device in your network. A single policy can contain only one instance of any type of rulebase.

IDP Policy Terms

Before configuring IDP policies, become familiar with the terms defined in Table 72 on page 502.

Table 72: IDP Terms

Term	Definition
Attacks	Attacks attempt to exploit vulnerabilities in computer hardware and software. Depending on the severity of the attack, it might disable your system completely, allow an attacker to gain confidential information stored on your system, or use your network to attack other networks.
Attack objects	A signature or protocol anomaly that is combined with context information. Attack objects are used in Main rulebase rules to match malicious traffic patterns. Each attack object detects a known attack or protocol anomaly that can be used by an attacker to compromise your network.
False positives	Any situation in which benign traffic causes an intrusion detection service to generate an alert; also known as a false alert.
Protocol anomaly	A deviation from the RFC specifications that dictate how communications between two entities should be implemented. Most legitimate traffic does not deviate from the protocols; when anomalies are detected, they are often a sign of malicious traffic and seen as a threat to the system.
Rule	A user-defined match/action sequence. Rules are represented graphically in the Security Policy Editor, where you can create, modify, delete, and reorder them in a rulebase.

Table 72: IDP Terms *(continued)*

Term	Definition
Rulebase	A set of rules that uses a specific detection mechanism to identify and prevent attacks.
Severity	The designated threat level of an attack (critical, high, medium, low, or informational). Attack objects use the severity setting that matches the threat level of the attack they detect.

Working with IDP Policies

You can perform the following tasks to manage IDP policies:

- Create new IDP policies starting from scratch (see "Defining Rules for an IPS Rulebase" on page 515).
- Create an IDP policy starting with one of the predefined templates provided by Juniper Networks (see "Using Predefined Policy Templates" on page 572).
- Add or delete rules within a rulebase. You can use any of the following IDP objects to create rules:
  - Zone and network objects available in the base system
  - Predefined service objects provided by Juniper Networks
  - Custom application objects
  - Predefined attack objects provided by Juniper Networks
- Create custom attack objects (see "Configuring Signature-Based Attacks" on page 562).
- Update the signature database provided by Juniper Networks. This database contains all predefined objects.
- Maintain multiple IDP policies. Any one of the policies can be applied to the device.

Understanding IDP Policy Rulebases

IDP policies are collections of rules and rulebases. A rulebase is an ordered set of rules that use a specific detection method to identify and prevent attacks.

Before You Begin

For background information, read:

- Introducing JUNOS Software for J Series Services Routers on page 33
- IDP Policies Overview on page 502

Rules are instructions that provide context to detection mechanisms by specifying which part of the network traffic the IDP system should look in to find attacks. When a rule is matched, it means that an attack has been detected in the network traffic, triggering the action for that rule. The IDP system performs the specified action and protects your network from that attack.

Each rulebase can have multiple rules—you determine the sequence in which rules are applied to network traffic by placing them in the desired order. Each rulebase in the IDP system uses specific detection methods to identify and prevent attacks. JUNOS Software supports two types of rulebases—intrusion prevention system (IPS) rulebase and exempt rulebase.

This topic covers:

- Understanding Application-level DDoS Rulebase on page 504
- IPS Rulebase on page 505
- Exempt Rulebase on page 505
- Related Topics on page 506

Understanding Application-level DDoS Rulebase

The application-level DDoS rulebase defines parameters used to protects servers, such as DNS or HTTP, from application-level distributed denial-of-service (AppDDoS) attacks. You can set up custom application metrics based on normal server activity requests to determine when clients should be considered an attack client. The AppDDoS rulebase is then used to defines the source match condition for traffic that should be monitored, then takes the defined action: close server, drop connection, drop packet, or no action. It can also perform an IP action: ip-block, ip-close, ip-notify, or timeout. Table 73 on page 504 summarizes the options that you can configure in the AppDDoS rulebase rules.

Table 73: AppDDoS Rulebase Components

Term	Definition
Match condition	Specify the network traffic you want the device to monitor for attacks.
Action	Specify the actions you want IDP to take when the monitored traffic matches the application-ddos objects specified in the application-level DDoS rule.
IP Action	Enables you to implicitly block a source address to protect the network from future intrusions while permitting legitimate traffic. You can configure one of the following IP action options in the AppDDoS:, none, drop, drop packet, or close server.

Related Topics

- Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview on page 808

- Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview on page 807
- Enabling Protection Against Application-level Distributed Denial-of-Service (AppDDoS) Attack on page 812

IPS Rulebase

The IPS rulebase protects your network from attacks by using attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies. Table 74 on page 505 summarizes the options that you can configure in the IPS-rulebase rules.

Table 74: IPS Rulebase Components

Term	Definition
Match condition	Specify the type of network traffic you want the device to monitor for attacks. For more information about match conditions, see "Understanding IDP Rule Match Conditions" on page 507.
Attack objects/groups	Specify the attacks you want the device to match in the monitored network traffic. Each attack is defined as an attack object, which represents a known pattern of attack. For more information about attack objects, see "Understanding IDP Rule Objects" on page 508.
Terminal flag	Specify a terminal rule. The device stops matching rules for a session when a terminal rule is matched. For more information about terminal rules, see "Understanding IDP Terminal Rules" on page 540.
Action	Specify the action you want the system to take when the monitored traffic matches the attack objects specified in the rules. If an attack triggers multiple rule actions, then the most severe action among those rules is executed. For more information about actions, see "Understanding IDP Rule Actions" on page 510.
IP Action	Enables you to protect the network from future intrusions while permitting legitimate traffic. You can configure one of the following IP action options in the IPS rulebase—notify, drop, or close. For more information about IP actions, see "Understanding IDP Rule IP Actions" on page 512.
Notification	Defines how information is to be logged when action is performed. You can choose to log an attack, create log records with the attack information, and send information to the log server. For more information, see "Understanding IDP Rule Notifications" on page 514.

Exempt Rulebase

The exempt rulebase works in conjunction with the IPS rulebase to prevent unnecessary alarms from being generated. You configure rules in this rulebase to exclude known false positives or to exclude a specific source, destination, or

source/destination pair from matching an IPS rule. If traffic matches a rule in the IPS rulebase, the system attempts to match the traffic against the exempt rulebase before performing the action specified. Carefully written rules in an exempt rulebase can significantly reduce the number of false positives generated by an IPS rulebase.

NOTE: Make sure to configure the IPS rulebase before configuring the exempt rulebase.

Table 75 on page 506 summarizes the options that you can configure in the exempt-rulebase rules.

Table 75: Exempt Rulebase Options

Term	Definition
Match condition	Specify the type of network traffic you want the device to monitor for attacks in the same way as in the IPS rulebase. However, in the exempt rulebase, you cannot configure an application; it is always set to any.
Attack objects/groups	Specify the attack objects that you do *not* want the device to match in the monitored network traffic.

Related Topics

- Understanding IDP Policy Rules on page 506
- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519

Understanding IDP Policy Rules

Each instruction in an IDP policy is called a rule. Rules are created in rulebases.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503

Rulebases are a set of rules that combine to define an IDP policy. Rules provide context to detection mechanisms by specifying which part of the network traffic the IDP system should look in to find attacks. When a rule is matched, it means that an attack has been detected in the network traffic, triggering the action for that rule. The IDP system performs the specified action and protects your network from that attack.

IDP policy rules are made up of the following components:

- Match Conditions (see “Understanding IDP Rule Match Conditions” on page 507)
- Attack Objects and Groups (see “Understanding IDP Rule Objects” on page 508)
- Actions (see “Understanding IDP Rule Actions” on page 510)
- IP Actions (see “Understanding IDP Rule IP Actions” on page 512)
- Notifications (see “Understanding IDP Rule Notifications” on page 514)

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519
- Using Predefined Policy Templates on page 572
- Inserting a Rule in the Rulebase on page 530
- Deactivating and Reactivating Rules in a Rulebase on page 532

Understanding IDP Rule Match Conditions

Match conditions specify the type of network traffic you want IDP to monitor for attacks.

Before You Begin
For background information, read: ■ IDP Policies Overview on page 502 ■ Understanding IDP Policy Rulebases on page 503 ■ Understanding IDP Policy Rules on page 506

Match conditions use the following characteristics to specify the type of network traffic to be monitored:

- From-zone and to-zone—All traffic flows from a source to a destination zone. You can select any zone for the source or destination. You can also use zone exceptions to specify unique to and from zones for each device. Specify any to monitor network traffic originating from and to any zone. The default value is any.
- Source IP Address—Specify the source IP address from which the network traffic originates. You can specify any to monitor network traffic originating from any IP address. You can also specify source-except to specify all sources except the specified addresses. The default value is any.
- Destination IP address—Specify the destination IP address to which the network traffic is sent. You can set this to any to monitor network traffic sent to any IP

address. You can also specify destination-except to specify all destinations except the specified addresses. The default value is any.

- Application—Specify the Application Layer protocols supported by the destination IP address. You can specify any for all applications and default for the application configured in the attack object for the rule.

Related Topics

- Understanding IDP Rule Objects on page 508
- Understanding IDP Rule Actions on page 510
- Understanding IDP Rule IP Actions on page 512
- Understanding IDP Rule Notifications on page 514

Understanding IDP Rule Objects

Objects are reusable logical entities that you can apply to rules. Each object that you create is added to a database for the object type.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

You can configure the following types of objects for IDP rules.

Zone Objects

A zone or security zone is a collection of one or more network interfaces. IDP uses zone objects configured in the base system.

Address or Network Objects

Address objects represent components of your network, such as host machines, servers, and subnets. You use address objects in IDP policy rules to specify the network components that you want to protect.

Application or Service Objects

Service objects represent network services that use Transport Layer protocols such as TCP, UDP, RPC, and ICMP. You use service objects in rules to specify the service an attack uses to access your network. Juniper Networks provides predefined service objects, a database of service objects that are based on industry-standard services.

If you need to add service objects that are not included in the predefined service objects, you can create custom service objects. IDP supports the following types of service objects:

- Any—Allows IDP to match all Transport Layer protocols.
- TCP—Specifies a TCP port or a port range to match network services for specified TCP ports. You can specify junos-tcp-any to match services for all TCP ports.
- UDP—Specifies a UDP port or a port range to match network services for specified UDP ports. You can specify junos-udp-any to match services for all UDP ports.
- RPC—Specifies a remote procedure call (RPC from Sun Microsystems) program number or a program number range. IDP uses this information to identify RPC sessions.
- ICMP—Specifies a type and code that is a part of an ICMP packet. You can specify junos-icmp-all to match all ICMP services.

Attack Objects

IDP attack objects represent known and unknown attacks. IDP includes a predefined attack object database that is periodically updated by Juniper Networks. Attack objects are specified in rules to identify malicious activity. Each attack is defined as an attack object, which represents a known pattern of attack. Whenever this known pattern of attack is encountered in the monitored network traffic, the attack object is matched. There are three main types of attack objects:

Signature Attack Objects

Signature attack objects detect known attacks using stateful attack signatures. An attack signature is a pattern that always exists within an attack; if the attack is present, so is the attack signature. With stateful signatures, IDP can look for the specific protocol or service used to perpetrate the attack, the direction and flow of the attack, and the context in which the attack occurs. Stateful signatures produce few false positives because the context of the attack is defined, eliminating huge sections of network traffic in which the attack would not occur.

Protocol Anomaly Attack Objects

Protocol anomaly attack objects identify unusual activity on the network. They detect abnormal or ambiguous messages within a connection according to the set of rules for the particular protocol being used. Protocol anomaly detection works by finding deviations from protocol standards, most often defined by RFCs and common RFC extensions. Most legitimate traffic adheres to established protocols. Traffic that does not, produces an anomaly, which may be created by attackers for specific purposes, such as evading an IPS.

Compound Attack Objects

A compound attack object combines multiple signatures and/or protocol anomalies into a single object. Traffic must match all of the combined signatures and/or protocol anomalies to match the compound attack object; you can specify the order in which

signatures or anomalies must match. Use compound attack objects to refine your IDP policy rules, reduce false positives, and increase detection accuracy. A compound attack object enables you to be very specific about the events that need to occur before IDP identifies traffic as an attack. You can use And, Or, and Ordered and operations to define the relationship among different attack objects within a compound attack and the order in which events occur.

Attack Object Groups

IDP contains a large number of predefined attack objects. To help keep IDP policies organized and manageable, attack objects can be grouped. An attack object group can contain one or more attack objects of different types. JUNOS Software supports the following two types of attack groups:

- Static groups—Contain a fixed set of attack objects.
- Dynamic groups—Contain attack objects based on a certain matching criteria. For example, a dynamic group can contain all attacks related to an application. During signature update, the dynamic group membership is automatically updated based on the matching criteria for that group.

Related Topics

- Understanding IDP Rule Match Conditions on page 507
- Understanding IDP Rule Actions on page 510
- Understanding IDP Rule IP Actions on page 512
- Understanding IDP Rule Notifications on page 514

Understanding IDP Rule Actions

Actions specify the actions you want IDP to take when the monitored traffic matches the attack objects specified in the rules.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

Table 76 on page 511 shows the actions you can specify for IDP rules:

Table 76: IDP Rule Actions

Term	Definition
No Action	No action is taken. Use this action when you only want to generate logs for some traffic.
Ignore Connection	Stops scanning traffic for the rest of the connection if an attack match is found. IDP disables the rulebase for the specific connection. **NOTE:** This action does not mean ignore an attack.
Diffserv Marking	Assigns the indicated Differentiated Services code point (DSCP) value to the packet in an attack, then passes the packet on normally. Note that DSCP value is not applied to the first packet that is detected as an attack, but is applied to subsequent packets.
Drop Packet	Drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source-IP address. **NOTE:** When an IDP policy is configured using a non-packet context defined in a custom signature for any application and has the action drop packet, when IDP identifies an attack the decoder will promote drop_packet to drop_connection. With a DNS protocol attack, this is not the case. The DNS decoder will not promote drop_packet to drop_connection when an attack is identified. This will ensure that only DNS attack traffic will be dropped and valid DNS requests will continue to be processed. This will also avoid TCP retransmission for the valid TCP DNS requests..
Drop Connection	Drops all packets associated with the connection, preventing traffic for the connection from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
Close Client	Closes the connection and sends an RST packet to the client but not to the server.
Close Server	Closes the connection and sends an RST packet to the server but not to the client.
Close Client and Server	Closes the connection and sends an RST packet to both the client and the server.

Table 76: IDP Rule Actions *(continued)*

Term	Definition
Recommended	All predefined attack objects have a default action associated with them. This is the action that Juniper Networks recommends when that attack is detected. **NOTE:** This action is supported only for IPS rulebases. Recommended —A list of all attack objects that Juniper Networks considers to be serious threats, organized into categories. ■ Attack type groups attack objects by type (anomaly or signature). Within each type, attack objects are grouped by severity. ■ Category groups attack objects by predefined categories. Within each category, attack objects are grouped by severity. ■ Operating system groups attack objects by the operating system to which they apply: BSD, Linux, Solaris, or Windows. Within each operating system, attack objects are grouped by services and severity. ■ Severity groups attack objects by the severity assigned to the attack. IDP has five severity levels: Critical, Major, Minor, Warning, and Info. Within each severity, attack objects are grouped by category.

Related Topics

- Understanding IDP Rule Match Conditions on page 507
- Understanding IDP Rule Objects on page 508
- Understanding IDP Rule IP Actions on page 512
- Understanding IDP Rule Notifications on page 514

Understanding IDP Rule IP Actions

IP actions are actions that apply on future connections that use the same IP action attributes. For example, you can configure an IP action in the rule to block all future HTTP sessions between two hosts if an attack is detected on a session between the hosts. Or you can specify a timeout value that defines that the action should be applied only if new sessions are initiated within that specified timeout value. The default timeout value for IP actions is 0, which means that IP actions are never timed out.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

IP actions are similar to other actions; they direct IDP to drop or close the connection. However, because you now also have the attacker's IP address, you can choose to block the attacker for a specified time. If attackers cannot immediately regain a connection to your network, they might try to attack easier targets. Use IP actions in conjunction with actions and logging to secure your network.

IP action attributes are a combination of the following fields:

- Source IP address
- Destination IP address
- Destination port
- From-zone
- Protocol

Table 77 on page 513 summarizes the types IP actions supported by IDP rules:

Table 77: IDP Rule IP Actions

Term	Definition
Notify	Does not take any action against future traffic, but logs the event. This is the default.
Drop/Block Session	All packets of any session matching the IP action rule are dropped silently.
Close Session	Any new sessions matching this IP action rule are closed by sending RST packets to the client and server.

When traffic matches multiple rules, the most severe IP action of all matched rules is applied. The most severe IP action is the Drop/Block Session action, the next in severity is the Close Session action, and then the Notify action.

Related Topics

- Understanding IDP Rule Match Conditions on page 507
- Understanding IDP Rule Objects on page 508
- Understanding IDP Rule Actions on page 510
- Understanding IDP Rule Notifications on page 514

Understanding IDP Rule Notifications

Notification defines how information is to be logged when an action is performed. When attacks are detected, you can choose to log an attack and create log records with attack information and send that information to the log server.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

By using notifications, you can also configure the following options that instruct the log server to perform specific actions on logs generated for each rule:

- Set Alerts—Specify an alert option for a rule in the IDP policy. When the rule is matched, the corresponding log record displays an alert in the alert column of the Log Viewer. Security administrators use alerts to become aware of and react to important security events.
- Send Emails—Specify this option to send a notification to a specified e-mail address.
- Run Scripts—Run a specified script on a log server.
- Set Severity Level—Set severity levels in logging to support better organization and presentation of log records on the log server. You can use the default severity settings of the selected attack objects or choose a specific severity for your rule. The severity you configure in the rules overrides the inherited attack severity. You can set the severity level to the following levels:
  - Info—2
  - Warning—3
  - Minor—4
  - Major—5
  - Critical—7

Related Topics

- Understanding IDP Rule Match Conditions on page 507
- Understanding IDP Rule Objects on page 508
- Understanding IDP Rule Actions on page 510
- Understanding IDP Rule IP Actions on page 512

Defining Rules for an IPS Rulebase

Each rule is composed of match conditions, objects, actions, and notifications. When you define an IDP rule, you must specify the type of network traffic you want IDP to monitor for attacks by using the following characteristics—source zone, destination zone, source IP address, destination IP address, and the Application Layer protocol supported by the destination IP address. The rules are defined in rulebases, and rulebases are associated with policies.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Policy Rules on page 506
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Create security zones. See "Creating Security Zones" on page 49.
5. Enable IDP in security policies. See "Enabling IDP in a Security Policy" on page 536.

The configuration instructions in this topic describe how to create a policy called base-policy, specify a rulebase for this policy, and then add a rule R1 to this rulebase. In this example, rule R1:

- Specifies the match condition to include any traffic from a previously configured zone called *trust* to another previously configured zone called *untrust*. The match condition also includes a predefined attack group Critical - TELNET. The application setting in the match condition is *default* and matches any application configured in the attack object.
- Specifies an action to drop connection for any traffic that matches the criteria for rule R1,
- Enables attack logging and specifies that an alert flag is added to the attack log.
- Specifies a severity level as *critical*.

After defining the rule, you specify base-policy as the active policy on the device.

You can use either J-Web or the CLI configuration editor to configure an application set.

This topic contains:

- J-Web Configuration on page 516
- CLI Configuration on page 517
- Related Topics on page 519

J-Web Configuration

To define rules for an IPS rulebase:

1. Create a policy by assigning a meaningful name to it. The following tasks specify base-policy as the policy name:
   a. Select **Configure > CLI Tools > Point and Click CLI**.
   b. Next to Security, click **Configure** or **Edit**.
   c. Next to Idp, click Configure.
   d. Next to Idp policy, click Add new entry.
   e. In the Policy name box, type **base-policy**.
2. Associate a rulebase with the policy and add rules to the rulebase. The following tasks associate an IPS rulebase with base-policy and add rule R1 to the rulebase.
   a. Next to Rulebase ips, click Configure.
   b. Next to Rule, click Add new entry.
   c. In the Name box, type **R1**.
3. Define the match criteria for the rule. The following tasks specify that traffic from *trust* zone to *untrust* zone as match criteria for rule R1. The *default* application setting matches any application configured in the attack object.
   a. Next to Match, click Configure.
   b. From the From zone list, select Enter specific value and type **trust** in the Zone box.
   c. From the To zone list, select Enter specific value and type **untrust** in the Zone box.
   d. From the Source list, select Source address.
   e. Next to Source address, click Add new entry.
   f. From the Value list, select Any and click OK.
   g. From the Destination list, select Destination address.
   h. Next to Destination address, click Add new entry.
   i. From the Value list, select Any and click OK.

j. Next to Application, click Add new entry.

k. From the Value list, select default and click OK.

4. Define an attack as match criteria. The following tasks specify predefined attack group Critical - TELNET as match criteria for rule R1.

a. On the Rule R1 page, next to Match, click Configure.

b. Next to Attacks, click Configure.

c. Next to Predefined attack groups, click Add new entry.

d. In the Value box, type **"Critical - TELNET"** and click OK.

5. Specify an action for the rule. The following tasks specify that the connection be dropped for any traffic that matches the criteria defined for rule R1:

a. On the Rule R1 page, next to Then, click Configure.

b. Next to Action, click Configure.

c. From the Action list, select Drop connection and click OK.

6. Specify notification and logging options for the rule. The following tasks enable logging for this attack and specify that an alert flag be added to the attack log:

a. On the Rule R1 page, next to Then, click Configure.

b. Next to Notification, select the check box and click Configure.

c. Next to Log attacks, select the check box and click Configure.

d. Next to Alert, select the check box and click OK.

7. Set the severity level for the rule. The following tasks set a critical severity level for rule R1:

a. On the Rule R1 page, next to Then, select Configure or Edit.

b. From the Severity list, select critical and click OK.

8. Activate the policy. The following tasks specify base-policy as the active policy:

a. On the Idp page, in the Active-policy box, type **base-policy**.

b. Click OK.

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

CLI Configuration

To define rules for an IPS rulebase:

1. Create a policy by assigning a meaningful name to it. The following statement specifies base-policy as the policy name:

```
user@host# set security idp idp-policy base-policy
```

2. Associate a rulebase with the policy. The following statement associates an IPS rulebase with base-policy:

   user@host# **set security idp idp-policy base-policy rulebase-ips**

3. Add rules to the rulebase. The following statement adds a rule R1 to the rulebase:

   user@host# **set security idp idp-policy base-policy rulebase-ips rule R1**

4. Define the match criteria for the rule. The following statement specifies that traffic from *trust* zone to *untrust* zone as match criteria for rule R1. The *default* application setting matches any application configured in the attack object.

   user@host# **set security idp idp-policy base-policy rulebase-ips rule R1 match from-zone trust to-zone untrust source-address any destination-address any application default**

5. Define an attack as match criteria. The following statement specifies predefined attack group Critical - TELNET as match criteria for rule R1:

   user@host# **set security idp idp-policy base-policy rulebase-ips rule R1 match attacks predefined-attack-group "Critical - TELNET"**

6. Specify an action for the rule. The following statement specifies that the connection be dropped for any traffic that matches the criteria defined for rule R1:

   user@host# **set security idp idp-policy base-policy rulebase-ips rule R1 then action drop-connection**

7. Specify notification and logging options for the rule. The following statement enables logging for this attack and specifies that an alert flag be added to the attack log:

   user@host# **set security idp idp-policy base-policy rulebase-ips rule R1 then notification log-attacks alert**

8. Set the severity level for the rule. The following statement sets a critical severity level for rule R1:

   user@host# **set security idp idp-policy base-policy rulebase-ips rule R1 then severity critical**

9. Activate the policy. The following statement specifies base-policy as the active policy:

   user@host# **set security idp active-policy base-policy**

10. If you are finished configuring the router, commit the configuration.

11. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an Exempt Rulebase on page 519
- Using Predefined Policy Templates on page 572
- Inserting a Rule in the Rulebase on page 530
- Deactivating and Reactivating Rules in a Rulebase on page 532

Defining Rules for an Exempt Rulebase

The exempt rulebase works in conjunction with the IPS rulebase. Before you can create exempt rules, you must first create rules in the IPS rulebase. If traffic matches a rule in the IPS rulebase, IDP attempts to match the traffic against the exempt rulebase before performing the specified action or creating a log record for the event. If IDP detects traffic that matches the source/destination pair and the attack objects specified in the exempt rulebase, it automatically exempts that traffic from attack detection.

Configure an exempt rulebase in the following conditions:

- When an IDP rule uses an attack object group that contains one or more attack objects that produce false positives or irrelevant log records.
- When you want to exclude a specific source, destination, or source/destination pair from matching an IDP rule. This prevents IDP from generating unnecessary alarms.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Policy Rules on page 506
2. Create rules in the IPS rulebase. See "Defining Rules for an IPS Rulebase" on page 515.

When you create an exempt rule, you must specify the following:

- Source and destination for traffic you want to exempt. You can set the source or destination to Any to exempt network traffic originating from any source or sent to any destination. You can also set source-except or destination-except to specify all the sources or destinations except the specified source or destination addresses.
- The attacks you want IDP to exempt for the specified source/destination addresses. You must include at least one attack object in an exempt rule.

In this configuration example, you consistently find that your IDP policy generates false positives for the attack FTP:USER:ROOT on your internal network. You configure the rule to exempt attack detection for this attack when the source IP is from your internal network.

You can use either J-Web or the CLI configuration editor to configure an application set.

This topic contains:

- J-Web Configuration on page 520
- CLI Configuration on page 521
- Related Topics on page 521

J-Web Configuration

To define rules for an exempt rulebase:

1. Specify the IPS rulebase for which you want to define and exempt rulebase. The following statement specifies policy P1 as the IPS rulebase:
   a. Select **Configure > CLI Tools > Point and Click CLI**.
   b. Next to Security, click **Configure** or **Edit**.
   c. Next to Idp, click Configure .
   d. Next to Idp policy, click Add new entry.
   e. In the Policy name box, type **P1**.
2. Associate the exempt rulebase with the policy and add a rule to the rulebase. The following tasks associate the exempt rulebase with policy P1 and adds rule R1 to the rulebase:
   a. Next to Rulebase exempt, click Configure.
   b. Next to Rule, click Add new entry.
   c. In the Name box, type **R1**.
3. Specify the attacks that you want to exempt from attack detection. The following configuration statement specifies that any traffic in your company's internal network is exempt from the FTP:USER:ROOT attack:
   a. Next to Match, click Configure.
   b. From the From zone list, select Enter specific value and type **trust** in the Zone box.
   c. From the To zone list, select any.
   d. From the Source list, select Source address.
   e. Next to Source address, select Add new entry.

f. From the Value list, select Enter specific value.

g. In the Address box, type **FTP:USER:ROOT**.

4. Activate the policy. The following tasks specify P1 as the active policy:

   a. On the Idp page, in the Active-policy box, type **P1**.

   b. Click OK.

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

CLI Configuration

To define rules for an exempt rulebase:

1. Specify the IPS rulebase for which you want to define and exempt rulebase. The following statement specifies policy P1 as the IPS rulebase:

   user@host# **set security idp idp-policy P1**

2. Associate the exempt rulebase with the policy and add a rule to the rulebase. The following statement associates the exempt rulebase with policy P1 and adds rule R1 to the rulebase:

   user@host# **set security idp idp-policy P1 rulebase-exempt rule R1**

3. Specify the attacks that you want to exempt from attack detection. The following configuration statement specifies that any traffic in your company's internal network is exempt from the FTP:USER:ROOT attack:

   user@host# **set security idp idp-policy P1 rulebase-exempt R1 match from-zone trust to-zone any source-address internal-devices destination-address any attacks predefined-attacks "FTP:USER:ROOT"**

4. Activate the policy. The following statement specifies policy P1 as the active policy on the device:

   user@host# **set security idp active-policy P1**

5. If you are finished configuring the router, commit the configuration.

6. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Using Predefined Policy Templates on page 572
- Inserting a Rule in the Rulebase on page 530
- Deactivating and Reactivating Rules in a Rulebase on page 532

IDP Policies—Quick Configuration

This topic contains:

- Configuring IDP Policies—Quick Configuration on page 522
- Adding a New IDP Policy—Quick Configuration on page 523
- Adding an IPS Rulebase—Quick Configuration on page 525
- Adding an Exempt Rulebase—Quick Configuration on page 528

Configuring IDP Policies—Quick Configuration

You can use J-Web Quick Configuration to quickly configure an IDP policy.

Figure 99: Quick Configuration page for IDP Policies

To configure an IDP policy with Quick Configuration:

1. Select Configure>Security>Policy>IDP Policies.

   Figure 99 on page 522 shows the Quick Configuration page for IDP policies.
2. Under Configured IDP Policies List, click Add.
3. Fill in the information as described in Table 78 on page 523.
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

Table 78: IDP Policies Quick Configuration Page Summary

Field	Function	Actions
Active IDP Policy	Specifies the name of the active IDP policy enabled on the device.	Displays the active IDP policy.
Security Package Update	Specifies to manually download or install the updated signature database from the specified URL.	Click Security Package Update to download or install the updated signature database. For more information, see "Configuring a Security Package Update—Quick Configuration" on page 580.
Configured IDP Policies List	Specifies the list of all the configured IDP policies on the device.	Displays the list of all the configured IDP Policies on the device.
Configured Rulebase(s)	Specifies to reorder the IPS and Exempt rulebases in the Configured IDP Policies List.	From the Name column, select the policy name. A new page with the configured rulebases appears. From the Move column, select the up or down arrow to reorder the configured rulebase(s) in the IPS Rulebase and Exempt Rulebase tables.
Set as active policy	Specifies if the configured IDP policy is set as the active policy in the Configured IDP Policies List.	Select the check box next to the IDP policy you want to set as an active policy.
Add	Specifies to add a new IDP policy.	Click Add to add a new IDP policy. For more information, see "Adding a New IDP Policy—Quick Configuration" on page 523.
Copy	Specifies to copy an existing IDP policy from the Configured IDP Policies List.	Open a new page, where you can select a policy, and click Copy.
Delete	Specifies to delete an existing IDP policy from the Configured IDP Policies List.	Select the check box corresponding to the policy you want to delete, and click Delete.

Adding a New IDP Policy—Quick Configuration

You can use J-Web Quick Configuration to quickly configure and add a new IDP policy.

Figure 100: Quick Configuration Page for Adding a New IDP Policy

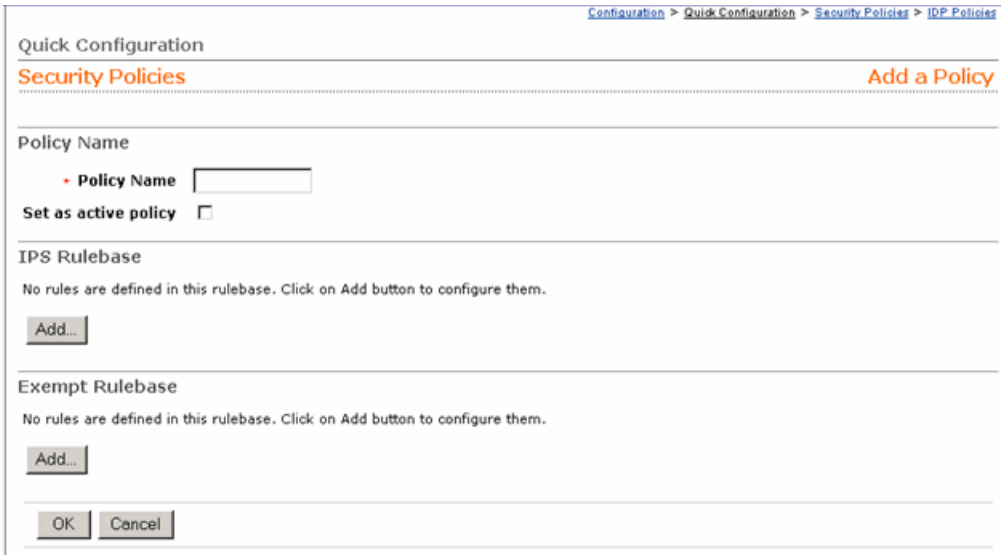

To configure a new IDP policy with Quick Configuration:

1. Select Configure>Security>Policy>IDP Policies.
2. Under Configured IDP Policies List, click Add.

   Figure 100 on page 524 shows the Quick Configuration page for adding a new IDP policy.
3. Fill in the information as described in Table 79 on page 524.
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

Table 79: Adding a New IDP Policy Quick Configuration Page Summary

Field	Function	Actions
Policy Name		
Policy Name	Specifies the name of the IDP policy.	Type a policy name.
Set as active policy	Specifies if the configured IDP policy is set as the active policy.	Select the check box.
IPS Rulebase	Defines the IPS rulebase to protect the network from attacks by using attack objects to detect known and unknown attacks.	Click Add under IPS Rulebase to add a new IPS rulebase to the IDP policy. For more information, see “Adding an IPS Rulebase—Quick Configuration” on page 525.

Table 79: Adding a New IDP Policy Quick Configuration Page Summary *(continued)*

Field	Function	Actions
Exempt Rulebase	Defines the exempt rulebase to exclude known false positives or to exclude a specific source, destination, or source/destination pair from matching an IPS rule.	Click Add under Exempt Rulebase to add an exempt rulebase to the IDP policy. For more information, see "Adding an Exempt Rulebase—Quick Configuration" on page 528.

Adding an IPS Rulebase—Quick Configuration

You can use J-Web Quick Configuration to quickly configure and add an IPS rulebase.

Figure 101: Quick Configuration Page for Adding an IPS Rulebase

To configure an IPS rulebase with Quick Configuration:

1. Select Configure>Security>Policy>IDP Policies.
2. Under Configured IDP Policies List, click Add.
3. In the Policy Name text box, type a policy name.
4. Under IPS Rulebase, click Add to add an IPS rulebase.

   Figure 101 on page 525 shows the Quick Configuration page for IPS rulebase.
5. Fill in the information as described in Table 80 on page 526.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

Table 80: Adding an IPS Rulebase Quick Configuration Page Summary

Field	Function	Actions
Policy Name	Specifies the name of the IDP Policy.	Displays the name of the IDP policy.
Rulebase	Specifies IPS rule to create, modify, delete, and reorder the rules in a rulebase.	Displays the name of the rulebase.
Configure Rule Name and Description		
Rule Name	Specifies the name of the IPS rulebase rule.	Type a rule name.
Description	Specifies the description for the rule.	Type the description for the rule.
Rule Match Criteria		
From-Zone and Source Addresses/Address Sets		
Match	Specifies the match criteria for the source zone for each rule.	Click the option button to enable the match criteria.
Source Address Book	Lists all the from-zone and source addresses/address sets for the policy.	Select the from-zone and source addresses/address sets from the list and do one of the following: ■ To match the from-zone and source address/address sets to the rule, click the left arrow. ■ To make the from-zone exceptions for each rule, click the right arrow.
Except	Specifies the zone exceptions for the from-zone and source address for each rule.	Click the option button to enable the exception criteria.
To-Zone and Destination Addresses/Address Sets		
Match	Specifies the match criteria for the to-zone and source addresses for each rule.	Click the option button to enable the match criteria.
Destination Address Book	Lists all the to-zone and destination addresses/address sets for the policy.	Select the to-zone and destination addresses/address sets from the list and do either one of the following: ■ To match the to-zone and destination addresses/address sets to the rule, click the left arrow. ■ To make the to-zone exceptions for each rule, click the right arrow.
Except	Specifies the except criteria for the to-zone and source address for each rule.	Click the option button to enable exception criteria.
Applications and Application Sets		
Matched	Specifies the type of network traffic you want the device to monitor for attacks.	

Table 80: Adding an IPS Rulebase Quick Configuration Page Summary *(continued)*

Field	Function	Actions
Application/Application Sets	Lists one or multiple configured applications and application sets.	Select the applications and application sets to be matched and do either one of the following: ■ To match the rule to the applications/application sets, click the left arrow. ■ To remove the rule match for the applications/application sets, select the rule match and click the right arrow.
Specify a rule action		
Rule Action	Lists all the rule actions for IDP to take when the monitored traffic matches the attack objects specified in the rules.	Select a rule action from the list.
Attacks and Attack Action		
Predefined Attacks	Specifies predefined attack objects that are used to match the traffic against known attacks.	Select one or multiple custom attacks from the Custom Attacks List and do either one of the following: ■ To match a custom attack to the rule, click the left arrow. ■ To remove the rule match for the custom attack to the rule, select the rule match and click the right arrow.
Predefined Attack Groups	Specifies predefined attack groups that are used to match the traffic against known attack objects.	Select one or multiple custom attacks from the Custom Attacks List and do either one of the following: ■ To match a custom attack to the rule, click the left arrow. ■ To remove the rule match for the custom attack to the rule, select the rule match and click the right arrow.
Custom Attacks	Specifies the custom attack objects to detect new attacks that are unique to your network.	Select one or multiple custom attacks from the Custom Attacks List and do either one of the following: ■ To match a custom attack to the rule, click the left arrow. ■ To remove the rule match for the custom attack to the rule, select the rule match and click the right arrow.
Attack Action		
IP Action	Specifies the action IDP takes against future connections that use the same IP address.	Select an IP action from the list.
IP Target	Specifies the destination IP address.	Select an IP target from the list.

Table 80: Adding an IPS Rulebase Quick Configuration Page Summary *(continued)*

Field	Function	Actions
Timeout	Specifies the number of seconds IP action should remain effective before new sessions are initiated within that specified timeout value.	Type the timeout value, in seconds. Maximum acceptable value is 65535 seconds.
Log IP Action	Specifies if the log attacks are enabled to create a log record that appears in the log viewer.	Select the check box.
Rule Additional Actions		
Severity	Specifies the rule severity levels in logging to support better organization and presentation of log records on the log server.	Select a severity level from the list.
Terminal	Specifies if the terminal rule flag is set or unset.	Select the check box.
Notifications - Attack Logging		
Enable	Specifies if the configuring attack logging alert is enabled.	Select the check box.
Set Alert Flag	Specifies if an alert flag is set.	Select the check box.

Adding an Exempt Rulebase—Quick Configuration

You can use J-Web Quick Configuration to quickly configure and add an exempt rulebase.

Figure 102: Quick Configuration Page for Adding an Exempt Rulebase

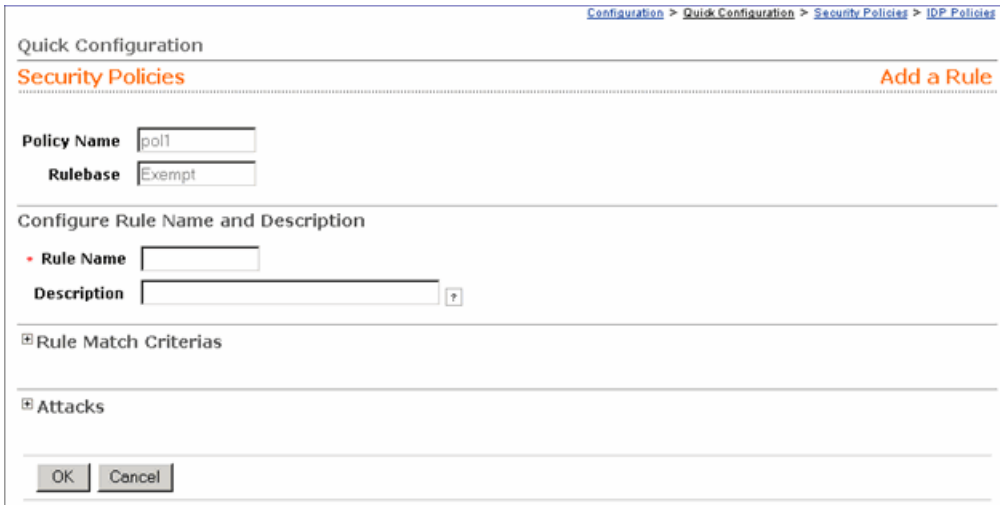

To configure an exempt rulebase with Quick Configuration:

1. Select Configure>Security>Policy>IDP Policies.
2. Under Configured IDP Policies List, click Add.

3. In the Policy Name text box, type a policy name.
4. Under Exempt Rulebase, click Add to add an exempt rulebase.

   Figure 102 on page 528 shows the Quick Configuration page for exempt rulebase.
5. Fill in the information as described in Table 81 on page 529.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

Table 81: Adding an Exempt Rulebase Quick Configuration Page Summary

Field	Function	Actions
Policy Name	Specifies the name of the IDP policy.	Displays the name of the IDP policy.
Rulebase	Specifies IPS rule to create, modify, delete, and reorder the rules in a rulebase.	Displays the name of the rulebase.
Configure Rule Name and Description		
Rule Name	Specifies the name of the IPS rulebase rule.	Type a rule name.
Description	Specifies the description for the rule.	Type the description for the rule.
Rule Match Criteria		
From-Zone and Source Addresses/Address Sets		
Match	Specifies the match criteria for the source zone for each rule.	Click the option button to enable the match criteria.
Source Address Book	Lists all the from-zone and source addresses/address sets for the policy.	Select the from-zone and source addresses/address sets from the list and do one of the following: ■ To match the from-zone and source address/address sets to the rule, click the left arrow. ■ To remove the rule match for the from-zone exceptions for each rule, click the right arrow.
Except	Specifies the zone exceptions for the from-zone and source address for each rule.	Click the option button to enable the exception criteria.
To-Zone and Destination Addresses/Address Sets		
Match	Specifies the match criteria for the to-zone and source addresses for each rule.	Click the option button to enable the match criteria.

Table 81: Adding an Exempt Rulebase Quick Configuration Page Summary *(continued)*

Field	Function	Actions
Destination Address Book	Lists all the to-zone and destination addresses/address sets for the policy.	Select the to-zone and destination addresses/address sets from the list and do either one of the following: ■ To match the to-zone and destination addresses/address sets to the rule, click the left arrow. ■ To remove the rule match for the to-zone exceptions for each rule, click the right arrow.
Except	Specifies the except criteria for the to-zone and source address for each rule.	Click the option button to enable exception criteria.
Attacks		
Predefined Attacks	Specifies predefined attack objects that are used to match the traffic against known attacks.	Select one or multiple custom attacks from the Custom Attacks List and do either one of the following: ■ To match a custom attack to the rule, click the left arrow. ■ To remove the rule match for the custom attack to the rule, select the rule match and click the right arrow.
Predefined Attack Groups	Specifies predefined attack groups that are used to match the traffic against known attack objects.	Select one or multiple custom attacks from the Custom Attacks List and do either one of the following: ■ To match a custom attack to the rule, click the left arrow. ■ To remove the rule match for the custom attack to the rule, select the rule match and click the right arrow.
Custom Attacks	Specifies the custom attack objects to detect new attacks that are unique to your network.	Select one or multiple custom attacks from the Custom Attacks List and do either one of the following: ■ To match a custom attack to the rule, click the left arrow. ■ To remove the rule match for the custom attack to the rule, select the rule match and click the right arrow.

Inserting a Rule in the Rulebase

The IDP rule-matching algorithm starts from the top of the rulebase and checks traffic against all rules in the rulebase that match the specified match conditions. You determine the sequence in which rules are applied to network traffic by placing them in the desired order. When you add a rule to the rulebase, it is placed at the end of

the existing list of rules. To place a rule in any other location than at the end of the rulebase, you *insert* the rule at the desired location in the rulebase.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Policy Rules on page 506
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Define rules in a rulebase. See"Defining Rules for an IPS Rulebase" on page 515.

The configuration instructions in this topic describe how to insert rule R2 before rule R1.

You can use either J-Web or the CLI configuration editor to insert a rule.

This topic contains:

- CLI Configuration on page 531
- Related Topics on page 531

CLI Configuration

To insert a rule in the rulebase:

1. Define the position of the rule in the rulebase based on the order in which you want the rule to be evaluated. The following configuration statement places rule R2 before rule R1 in the IPS rulebase in a policy called base-policy:

   ```
   user@host# insert security idp idp-policy base-policy rulebase-ips rule R2 before
     rule R1
   ```

2. If you are finished configuring the router, commit the configuration.
3. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an Exempt Rulebase on page 519
- Using Predefined Policy Templates on page 572
- Deactivating and Reactivating Rules in a Rulebase on page 532

Deactivating and Reactivating Rules in a Rulebase

In a rulebase, you can disable and enable rules by using the deactivate and activate commands.

The deactivate command comments out the specified statement from the configuration. Rules that have been deactivated do not take effect when you issue the commit command.

The activate command adds the specified statement back to the configuration. Rules that have been activated take effect when you next issue the commit command.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Policy Rules on page 506
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Define rules in a rulebase. See "Defining Rules for an IPS Rulebase" on page 515.

The configuration instructions in this topic describe how to deactivate and reactivate rule R2 in an IPS rulebase that is associated with a policy called base-policy.

You can use either J-Web or the CLI configuration editor to deactivate or activate rules in a rulebase.

This topic contains:

- CLI Configuration on page 532
- Related Topics on page 533

CLI Configuration

To deactivate and activate a rule in a rulebase:

1. Specify the rule that you want to deactivate. The following statement deactivates rule R2 in an IPS rulebase associated with base-policy:

   user@host# **deactivate security idp idp-policy base-policy rulebase-ips rule R2**

2. If you want to reactivate the rule, use the activate command. The following statement reactivates the deactivated rule R2 in the IPS rulebase:

   user@host# **activate security idp idp-policy base-policy rulebase-ips rule R2**

3. If you are finished configuring the router, commit the configuration.
4. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an Exempt Rulebase on page 519
- Using Predefined Policy Templates on page 572
- Inserting a Rule in the Rulebase on page 530

Understanding Application Sets

You specify an application, or service, to indicate that a policy applies to traffic of that type. Sometimes the same applications or a subset of them can be present in multiple policies, making them difficult to manage. JUNOS Software allows you to create groups of applications called *application sets*.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

Application sets simplify the process by allowing you to manage a small number of application sets, rather than a large number of individual application entries.

The application (or application set) is configured as a match criterion for packets. Packets must be of the application type specified in the policy for the policy to apply to the packet. If the packet matches the application type specified by the policy and all other criteria match, then the policy action is applied to the packet. You can use predefined or custom applications and refer to them in a policy.

Related Topics

- Configuring Application Sets for IDP on page 535
- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519

Configuring Applications or Services for IDP

Applications or services represent Application Layer protocols that define how data is structured as it travels across the network. Because the services you support on

your network are the same services that attackers must use to attack your network, you can specify which services are supported by the destination IP to make your rules more efficient. Juniper Networks provides predefined applications and application sets that are based on industry-standard applications. If you need to add applications that are not included in the predefined applications, you can create custom applications or modify predefined applications to suit your needs. To create custom applications, specify a meaningful name for an application and associate parameters with it—for example, inactivity timeout, or application protocol type.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Rule Objects on page 508
   - Configuring Applications or Services for IDP on page 533
2. Establish basic connectivity. (See the Getting Started Guide for your device.)
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Enable IDP in security policies. See "Enabling IDP in a Security Policy" on page 536.

The configuration instructions in this topic describe how to create an application cust-app and specify it as a match condition in the IDP policy ABC. In this example you create a special FTP application running on port 78. You also specify the inactivity timeout value as 6000 seconds:

You can use either J-Web or the CLI configuration editor to configure an application.

This topic contains:

- CLI Configuration on page 534
- Related Topics on page 535

CLI Configuration

To create an application and associate it with an IDP policy:

1. Specify a unique name for the application. The following statement specifies cust-app as the name of the application:

   user@host# **set applications application cust-app**

2. Specify application properties. The following statement specifies an FTP application using the TCP protocol and the port 78. Inactivity timeout for the FTP service is set to 6000 seconds.

   user@host# **set applications application cust-app application-protocol ftp protocol tcp destination-port 78 inactivity-timeout 6000**

3. Specify the application as a match condition in a policy. The following statement adds the cust-app application to the ABC policy:

   user@host# **set security idp idp-policy ABC rulebase-ips rule ABC match application cust-app**

4. If you are finished configuring the router, commit the configuration.
5. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Configuring Application Sets for IDP on page 535
- Defining Rules for an IPS Rulebase on page 515
- Using Predefined Policy Templates on page 572

Configuring Application Sets for IDP

To configure an application set, you add predefined or custom applications separately to an application set and assign a meaningful name to the application set. Once you name the application set you specify the name as part of the policy. For this policy to apply on a packet, the packet must match any one of the applications included in this set.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding Application Sets on page 533
2. Establish basic connectivity. See the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Enable IDP in security policies. See “Enabling IDP in a Security Policy” on page 536.

The configuration instructions in this topic describe how to create an application set SrvAccessAppSet and associate it with an IDP policy ABC. The application set SrvAccessAppSet combines three applications. Instead of specifying three applications in the policy rule, you specify one application set. If all of the other criteria match, any one of the applications in the application set serves as valid matching criteria.

You can use either J-Web or the CLI configuration editor to configure an application set.

This topic contains:

- CLI Configuration on page 536
- Related Topics on page 536

CLI Configuration

To create an application set and associate it with an IDP policy:

1. Create an application set and specify applications to be included in the set. The following statements create the SrvAccessAppSet application set that includes a set of three applications:

   ```
   user@host# set applications application-set SrvAccessAppSet application ssh
   user@host# set applications application-set SrvAccessAppSet application telnet
   user@host# set applications application-set SrvAccessAppSet application
     custApp
   ```

2. Associate the application set with an IDP policy. The following statement associates the application set SrvAccessAppSet to IDP policy ABC:

   ```
   user@host# set security idp idp-policy ABC rulebase-ips rule ABC match
     application SrvAccessAppSet
   ```

3. Specify an action for the policy. The following statement permits traffic from applications specified in the application set:

   ```
   user@host# set security idp idp-policy ABC rulebase-ips rule ABC then action
     no-action
   ```

4. If you are finished configuring the router, commit the configuration.
5. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Configuring Applications or Services for IDP on page 533
- Defining Rules for an IPS Rulebase on page 515
- Using Predefined Policy Templates on page 572

Enabling IDP in a Security Policy

For transit traffic to pass through IDP inspection, you configure a security policy and enable IDP application services on all traffic that you want to inspect. Security policies contain rules defining the types of traffic permitted on the network and the way that the traffic is treated inside the network. Enabling IDP in a security policy directs traffic that matches the specified criteria to be checked against the IDP rulebases.

To allow transit traffic to pass through without IDP inspection, specify a *permit* action for the rule without enabling the IDP application services. Traffic matching the conditions in this rule passes through the device without IDP inspection.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Policy Rules on page 506
2. Establish basic connectivity. See the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Create security zones. See "Creating Security Zones" on page 49.
5. Configure applications. See "Configuring Applications or Services for IDP" on page 533.

In this example, you configure two policies—idp-app-policy-1 and idp-app-policy-2. You configure these policies to enable IDP services on all traffic flowing in both directions on the device. Policy idp-app-policy-1 directs all traffic flowing from previously configured zones Zone1 to Zone2 to be checked against IDP rulebases. The policy idp-app-policy-2 directs all traffic flowing from Zone2 to Zone1 to be checked against IDP rulebases.

NOTE: The action set in the security policy action must be *permit*. You cannot enable IDP for traffic that the device denies or rejects.

You can use either J-Web or the CLI configuration editor to configure the IDP application services.

This topic contains:

- J-Web Configuration on page 537
- CLI Configuration on page 539
- Related Topics on page 540

J-Web Configuration

To enable IDP in a security policy:

1. Create a security policy. The following tasks create a policy idp-app-policy-1 for traffic traversing from Zone1 to Zone2:
   a. Select **Configure > CLI Tools > Point and Click CLI**.
   b. Next to Security, click **Configure** or **Edit**.
   c. Next to Policies, select the check box and click Edit.

d. Next to Policy, click Add new entry.

e. In the From zone name box, type **Zone1**.

f. In the To zone name box, type **Zone2**.

g. Next to Policy, click Add new entry.

h. In the Policy name box, type idp-app-policy-1.

2. Specify the match conditions for the traffic flowing in one direction. The following tasks specify that traffic from any source address, to any destination address and with any application type, matches the criteria for this policy:

a. Next to Match, select the check box and click Configure.

b. From the Source Address choice list, select Source address.

c. Next to Source address, click Add new entry.

d. From the Value keyword list, select Any and click OK.

e. From the Destination Address choice list, select Destination address.

f. Next to Destination address, click Add new entry.

g. From the Value keyword list, select Any and click OK.

h. From the Application choice list, select Application.

i. Next to Application, click Add new entry.

j. From the Value keyword list, select Any and click OK.

3. Specify the action to be taken on traffic that matches the specified conditions. The following tasks permit all traffic matching the specified criteria and direct it to be checked against IDP rulebases:

a. On the Policy name idp-app-policy-1 page, next to Then, select check box and click Configure.

b. From the Action list, select Permit.

c. Next to Permit, click Configure.

d. Next to Application services, click Configure.

e. Next to Idp, select the check box.

4. Create another security policy for traffic in the other direction. The following tasks create another policy idp-app-policy-2 for traffic from Zone2 to Zone1:

a. On the Policy page, next to Policy, click Add new entry.

b. In the From zone name box, type **Zone2**.

c. In the To zone name box, type **Zone1**.

d. Next to Policy, click Add new entry.

e. In the Policy name box, type idp-app-policy-2.

5. Specify the match conditions for the traffic flowing in the other direction. The following tasks specify that traffic from any source, to any destination with any application type, matches the criteria for this policy:

a. Next to Match, select the check box and click Configure.

b. From the Source Address choice list, select Source address.

c. Next to Source address, click Add new entry.

d. From the Value keyword list, select Any and click OK.

e. From the Destination Address choice list, select Destination address.

f. Next to Destination address, click Add new entry.

g. From the Value keyword list, select Any and click OK.

h. From the Application choice list, select Application.

i. Next to Application, click Add new entry.

j. From the Value keyword list, select Any and click OK.

6. Specify the action to be taken on traffic that matches the conditions specified in the policy. The following tasks permit all traffic matching the specified criteria and direct it to be checked against IDP rulebases:

a. On the Policy name idp-app-policy-2 page, next the Then, select check box and click Configure.

b. From the Action list, select Permit.

c. Next to Permit, click Configure.

d. Next to Application services, click Configure.

e. Next to Idp, select the check box.

7. If you are finished configuring the router, commit the configuration.

CLI Configuration

To enable IDP in a security policy:

1. Create a security policy. The following statement creates a policy idp-app-policy-1 for traffic traversing from Zone1 to Zone2:

```
user@host# set security policies from-zone Zone1 to-zone Zone2 policy
  idp-app-policy-1
```

2. Specify the match conditions for the traffic flowing in one direction. The following statement specifies that traffic from any source address, to any destination address and with any application type, matches the criteria for this policy:

   user@host# **set security policies from-zone Zone1 to-zone Zone2 policy idp-app-policy-1 match source-address any destination-address any application any**

3. Specify the action to be taken on traffic that matches the specified conditions. The following statement permits all traffic matching the specified criteria and directs it to be checked against IDP rulebases:

   user@host# **set security policies from-zone Zone1 to-zone Zone2 policy idp-app-policy-1 then permit application-services idp**

5. Create another security policy for traffic in the other direction. The following statement creates another policy idp-app-policy-2 for traffic from Zone2 to Zone1:

   user@host# **set security policies from-zone Zone2 to-zone Zone1 policy idp-app-policy-2**

6. Specify the match conditions for the traffic flowing in the other direction. The following statement specifies that traffic from any source, to any destination with any application type, matches the criteria for this policy:

   user@host# **set security policies from-zone Zone2 to-zone Zone1 policy idp-app-policy-2 match source-address any destination-address any application any**

7. Specify the action to be taken on traffic that matches the conditions specified in the policy. The following statement permits all traffic matching the specified criteria and directs it to be checked against IDP rulebases:

   user@host# **set security policies from-zone Zone2 to-zone Zone1 policy idp-app-policy-2 then permit application-services idp**

8. If you are finished configuring the router, commit the configuration.
9. You can verify the configuration by using the show security policies command. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519

Understanding IDP Terminal Rules

The IDP rule-matching algorithm starts from the top of the rulebase and checks traffic against all rules in the rulebase that match the source, destination, and service. However, you can configure a rule to be *terminal*. A *terminal* rule is an exception to this algorithm. When a match is discovered in a terminal rule for the source, destination, zones, and application, IDP does not continue to check subsequent rules

for the same source, destination, and application. It does not matter whether or not the traffic matches the attack objects in the matching rule.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

You can use a terminal rule for the following purposes:

- To set different actions for different attacks for the same Source and Destination.
- To disregard traffic that originates from a known trusted source. Typically, the action is None for this type of terminal rule.
- To disregard traffic sent to a server that is vulnerable only to a specific set of attacks. Typically, the action is Drop Connection for this type of terminal rule.

Use caution when defining terminal rules. An inappropriate terminal rule can leave your network open to attacks. Remember that traffic matching the source, destination, and application of a terminal rule is not compared to subsequent rules, even if the traffic does not match an attack object in the terminal rule. Use a terminal rule only when you want to examine a certain type of traffic for one specific set of attack objects. Be particularly careful about terminal rules that use any for both the source and destination. Terminal rules should appear near the top of the rulebase before other rules that would match the same traffic.

Related Topics

- Setting Terminal Rules in Rulebases on page 541
- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519

Setting Terminal Rules in Rulebases

By default, rules in the IDP rulebase are not terminal. That means that IDP examines all rules in the rulebase and executes all matches. You can specify that a rule is

terminal; if IDP encounters a match for the source, destination, and service specified in a terminal rule, it does not examine any subsequent rules for that connection.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Terminal Rules on page 540
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Enable IDP application services in a security policy. See "Enabling IDP in a Security Policy" on page 536.
5. Create security zones. See "Creating Security Zones" on page 49.
6. Define rules. See "Defining Rules for an IPS Rulebase" on page 515.

The configuration statements in this topic describe how to define terminal rules. You define a rule R2 to terminate the match algorithm if the source IP of the traffic originates from a known trusted network in your company. If this rule is matched, IDP disregards traffic from the trusted network and does not monitor the session for malicious data.

You can use either J-Web or the CLI configuration editor to configure terminal rules.

This topic contains:

- J-Web Configuration on page 542
- CLI Configuration on page 543
- Related Topics on page 544

J-Web Configuration

To configure terminal rules:

1. Create a policy by assigning a meaningful name to it. The following tasks specify P1 as the policy name:
   a. Select **Configure > CLI Tools > Point and Click CLI**.
   b. Next to Security, click **Configure** or **Edit**.
   c. Next to Idp, click Configure.

d. Next to Idp policy, click Add new entry.

e. In the Policy name box, type **P1**.

2. Associate a rulebase with the policy and add rules to the rulebase. The following tasks associate an IPS rulebase with P1 and add rule R2 to the rulebase.

a. Next to Rulebase ips, click Configure.

b. Next to Rule, click Add new entry.

c. In the Name box, type **R2**.

3. Define the match criteria for the rule. The following tasks specify source address internal and destination address any as the match criteria for rule R2.

a. Next to Match, click Configure.

b. From the Source list, select Source address.

c. Next to Source address, click Add new entry.

d. From the Value list, select Enter Specific value.

e. In the Address box, type **internal**.

f. From the Destination list, select Destination address.

g. Next to Destination address, click Add new entry.

h. From the Value list, select Any and click OK.

4. Set the terminal flag for the rule. The following tasks specify R2 as a terminal rule:

a. On the Rule R2 page, next to Terminal, select the check box.

b. Click OK.

5. If you are finished configuring the router, commit the configuration.

CLI Configuration

To configure terminal rules:

1. Define a rule and add it to a rulebase in a policy. The following statement creates a policy P1, associates an IPS rulebase with the policy, and adds rules R2 to the rulebase:

```
user@host# set security idp idp-policy P1 rulebase-ips rule R2
```

2. Define the match criteria for the rule. The following tasks specify source address internal and destination address any as the match criteria for rule R2:

```
user@host# set security idp idp-policy P1 rulebase-ips rule R2 match
  source-address internal destination-address any
```

3. Set the terminal flag for the rule. The following statement specifies R2 as a terminal rule:

   user@host# **set security idp idp-policy P1 rulebase-ips rule R2 terminal**

4. If you are finished configuring the router, commit the configuration.
5. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an Exempt Rulebase on page 519
- Using Predefined Policy Templates on page 572
- Inserting a Rule in the Rulebase on page 530
- Deactivating and Reactivating Rules in a Rulebase on page 532

Understanding Custom Attack Objects

You can create custom attack objects to detect new attacks or customize predefined attack objects to meet the unique needs of your network.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding Predefined Attack Objects and Groups on page 575

To configure a custom attack object, you specify a unique name for it and then specify additional information, such as a general description and keywords, which can make it easier for you to locate and maintain the attack object.

Certain properties in the attack object definitions are common to all types of attacks, such as attack name, description, severity level, service or application binding, time binding, recommended action, and protocol or port binding. Some fields are specific to an attack type and are available only for that specific attack definition.

This topic covers:

- Attack Name on page 545
- Severity on page 545
- Service or Application Binding on page 545
- Protocol or Port Bindings on page 549
- Time Bindings on page 550

- Attack Properties—Signature Attacks on page 551
- Attack Properties—Protocol Anomaly Attacks on page 557
- Attack Properties—Compound or Chain Attacks on page 558
- Related Topics on page 560

Attack Name

Specify an alphanumeric name for the object. You might want to include the protocol the attack uses in the attack name.

Severity

Specifies the brutality of the attack on your network. Severity categories, in order of increasing brutality, are info, warning, minor, major, critical (see "Understanding IDP Rule Notifications" on page 514). Critical attacks are the most dangerous—typically these attacks attempt to crash your server or gain control of your network. Informational attacks are the least dangerous, and typically are used by network administrators to discover holes in their own security systems.

Service or Application Binding

The service or application binding field specifies the service that the attack uses to enter your network.

NOTE: Specify either the service or the protocol binding in a custom attack. In case you specify both, the service binding takes precedence.

- Any—Specify any if you are unsure of the correct service and want to match the signature in all services. Because some attacks use multiple services to attack your network, you might want to select the Any service binding to detect the attack regardless of which service the attack chooses for a connection.
- Service—Most attacks use a specific service to attack your network. You can select the specific service used to perpetrate the attack as the service binding. Table 82 on page 545 displays supported services and default ports associated with the services.

Table 82: Supported Services for Service Bindings

Service	Description	Default Port
AIM	AOL Instant Messenger. America Online Internet service provider (ISP) provides Internet, chat, and instant messaging applications.	TCP/5190
BGP	Border Gateway Protocol	TCP/179
Chargen	Character Generator Protocol is a UDP- or TCP-based debugging and measurement tool.	TCP/19, UDP/19

Table 82: Supported Services for Service Bindings *(continued)*

Service	Description	Default Port
DHCP	Dynamic Host Configuration Protocol allocates network addresses and delivers configuration parameters from server to hosts.	UDP/67, UDP/68
Discard	Discard protocol is an Application Layer protocol that describes a process for discarding TCP or UDP data sent to port 9.	TCP/9, UDP/9
DNS	Domain Name System translates domain names into IP addresses.	TCP/53, UDP/53
Echo	Echo	TCP/7, UDP/7
Finger	Finger is a UNIX program that provides information about users.	TCP/79, UDP/79
FTP	File Transfer Protocol (FTP) allows the sending and receiving of files between machines.	TCP/21, UDP/21
Gnutella	Gnutella is a public domain file sharing protocol that operates over a distributed network.	TCP/6346
Gopher	Gopher organizes and displays Internet servers' contents as a hierarchically structured list of files.	TCP/70
H225RAS	H.225.0/RAS (Registration, Admission, and Status)	UDP/1718, UDP/1719
HTTP	HyperText Transfer Protocol is the underlying protocol used by the World Wide Web (WWW).	TCP/80, TCP/81, TCP/88, TCP/3128, TCP/7001 (Weblogic), TCP/8000, TCP/8001, TCP/8100 (JRun), TCP/8200 (JRun), TCP/8080, TCP/8888 (Oracle-9i), TCP/9080 (Websphere), UDP/80
ICMP	Internet Control Message Protocol	
IDENT	Identification protocol is a TCP/IP Application Layer protocol used for TCP client authentication.	TCP/113
IKE	Internet Key Exchange protocol (IKE) is a protocol to obtain authenticated keying material for use with ISAKMP.	UDP/500
IMAP	Internet Message Access Protocol is used for retrieving messages.	TCP/143, UDP/143
IRC	Internet Relay Chat (IRC) allows people connected to the Internet to join live discussions.	TCP/6667
LDAP	Lightweight Directory Access Protocol is a set of protocols used to access information directories.	TCP/389

Table 82: Supported Services for Service Bindings *(continued)*

Service	Description	Default Port
lpr	Line Printer Daemon protocol is a TCP-based protocol used for printing applications.	TCP/515
MSN	Microsoft Network Messenger is a utility that allows you to send instant messages and talk online.	TCP/1863
MSRPC	Microsoft Remote Procedure Call	TCP/135, UDP/135
MSSQL	Microsoft SQL is a proprietary database server tool that allows for the creation, access, modification, and protection of data.	TCP/1433, TCP/3306
MYSQL	MySQL is a database management system available for both Linux and Windows.	TCP/3306
NBDS	NetBIOS Datagram Service application, published by IBM, provides connectionless (datagram) applications to PCs connected with a broadcast medium to locate resources, initiate sessions, and terminate sessions. It is unreliable and the packets are not sequenced.	UDP/137 (NBName), UDP/138 (NBDS)
NFS	Network File System uses UDP to allow network users to access shared files stored on computers of different types. SUN RPC is a building block of NFS.	TCP/2049, UDP/2049
nntp	Network News Transport Protocol is a protocol used to post, distribute, and retrieve USENET messages.	TCP/119
NTP	Network Time Protocol provides a way for computers to synchronize to a time reference.	UDP/123
POP3	Post Office Protocol is used for retrieving e-mail.	UDP/110, TCP/110
Portmapper	Service that runs on nodes on the Internet to map an ONC RPC program number to the network address of the server that listens for the program number.	TCP/111, UDP/111
RADIUS	Remote Authentication Dial-In User Service application is a server program used for authentication and accounting purposes.	UDP/1812, UDP/1813
rexec	Rexec	TCP/512
rlogin	RLOGIN starts a terminal session on a remote host.	TCP/513
rsh	RSH executes a shell command on a remote host.	TCP/514
rtsp	Real-Time Streaming Protocol (RTSP) is for streaming media applications	TCP/554
SIP	Session Initiation Protocol (SIP) is an Application-Layer control protocol for creating, modifying, and terminating sessions.	TCP/5060, UDP/5060

Table 82: Supported Services for Service Bindings *(continued)*

Service	Description	Default Port
SMB	Server Message Block (SMB) over IP is a protocol that allows you to read and write files to a server on a network.	TCP/139, TCP/445
SMTP	Simple Mail Transfer Protocol is used to send messages between servers.	TCP/25, UDP/25
SNMP	Simple Network Management Protocol is a set of protocols for managing complex networks.	TCP/161, UDP/161
SNMPTRAP	SNMP trap	TCP/162, UDP/162
SQLMON	SQL monitor (Microsoft)	UDP/1434
SSH	SSH is a program to log into another computer over a network through strong authentication and secure communications on a channel that is not secure.	TCP/22, UDP/22
SSL	Secure Sockets Layer	TCP/443, TCP/80
syslog	Syslog is a UNIX program that sends messages to the system logger.	UDP/514
Telnet	Telnet is a UNIX program that provides a standard method of interfacing terminal routers and terminal-oriented processes to each other.	TCP/23, UDP/23
TNS	Transparent Network Substrate	TCP/1521, TCP/1522, TCP/1523, TCP/1524, TCP/1525, TCP/1526, TCP/1527, TCP/1528, TCP/1529, TCP/1530, TCP/2481, TCP/1810, TCP/7778
TFTP	Trivial File Transfer Protocol	UDP/69
VNC	Virtual Network Computing facilitates viewing and interacting with another computer or mobile router connected to the Internet.	TCP/5800, TCP/5900
Whois	Network Directory Application Protocol is a way to look up domain names.	TCP/43
YMSG	Yahoo! Messenger is a utility that allows you to check when others are online, send instant messages, and talk online.	TCP/5050

Protocol or Port Bindings

Protocol or port bindings allow you to specify the protocol that an attack uses to enter your network. You can specify the name of the network protocol, or the protocol number.

NOTE: Specify either the service or the protocol binding in a custom attack. In case you specify both, the service binding takes precedence.

- IP—You can specify any of the supported network layer protocols using protocol numbers. Table 83 on page 549 lists protocol numbers for different protocols.

Table 83: Supported Protocols and Protocol Numbers

Protocol Name	Protocol Number
IGMP	2
IPIP	4
EGP	8
PUP	12
TP	29
IPV6	41
ROUTING	43
FRAGMENT	44
RSVP	46
GRE	47
ESP	50
AH	51
ICMPV6	58
NONE	59
DSTOPTS	60
MTP	92
ENCAP	98
PIM	103

Table 83: Supported Protocols and Protocol Numbers *(continued)*

Protocol Name	Protocol Number
COMP	108
RAW	255

- ICMP, TCP, and UDP—Attacks that do not use a specific service might use specific ports to attack your network. Some TCP and UDP attacks use standard ports to enter your network and establish a connection.
- RPC—The remote procedure call (RPC) protocol is used by distributed processing applications to handle interaction between processes remotely. When a client makes a remote procedure call to an RPC server, the server replies with a remote program; each remote program uses a different program number. To detect attacks that use RPC, configure the service binding as RPC and specify the RPC program ID.

Table 84 on page 550 displays sample formats for key protocols.

Table 84: Sample Formats for Protocols

Protocol Name	Protocol Number	Description
ICMP	<Port>ICMP</Port>	Specify the protocol name.
IP	<Port>IP/*protocol-number*</Port>	Specify the Network Layer protocol number.
RPC	<Port>RPC/*program-number*</Port>	Specify the RPC program number.
TCP or UDP	■ <Port>TCP </Port> ■ <Port>TCP/*port* </Port> ■ <Port>TCP/*minport-maxport* </Port>	Specifying the port is optional for TCP and UDP protocols. For example, you can specify either of the following: ■ <Port>UDP</Port> ■ <Port>UDP/10</Port> ■ <Port>UDP/10-100</Port>

Time Bindings

Use time bindings to configure the time attributes for the custom attack object. Time attributes control how the attack object identifies attacks that repeat for a certain number of times. By configuring the scope and count of an attack, you can detect a sequence of the same attacks over a period of time (one minute) across sessions.

Scope

Specify the scope within which the count of an attack occurs:

- Source—Specify this option to detect attacks from the source address for the specified number of times, regardless of the destination address. This means that for a given attack, a threshold value is maintained for each attack from the source address. The destination address is ignored. For example, anomalies are detected from two different pairs (ip-a, ip-b) and (ip-a, ip-c) that have the same source address ip-a but different destination addresses ip-b and ip-c. Then the number of matches for ip-a increments to 2. Suppose the threshold value or *count* is also set to 2, then the signature triggers the attack event.
- Destination—Specify this option to detect attacks sent to the destination address for the specified number of times, regardless of the source address. This means that for a given attack, a threshold value is maintained for each attack from the destination address. The source address is ignored. For example, if anomalies are detected from two different pairs (ip-a, ip-b) and (ip-c, ip-b) that have the same destination address ip-b but different source addresses ip-a and ip-c. Then the number of matches for ip-b increments to 2. Suppose the threshold value or *count* is also set to 2, then the signature triggers the attack event.
- Peer—Specify this option to detect attacks between source and destination IP addresses of the sessions for the specified number of times. This means that the threshold value is applicable for a pair of source and destination addresses. Suppose anomalies are detected from two different source and destination pairs (ip-a, ip-b) and (ip-a, ip-c). Then the number of matches for each pair is set to 1, even though both pairs have a common source address.

Count

Count or threshold value specifies the number of times that the attack object must detect an attack within the specified scope before the device considers the attack object to match the attack. If you bind the attack object to multiple ports and the attack object detects that attack on different ports, each attack on each port is counted as a separate occurrence. For example, when the attack object detects an attack on TCP/80 and then on TCP/8080, the count is two.

Once the count match is reached, each attack that matches the criteria causes the attack count to increase by one. This count cycle lasts for a duration of 60 seconds, after which the cycle repeats.

Attack Properties—Signature Attacks

Signature attack objects use a stateful attack signature (a pattern that always exists within a specific section of the attack) to detect known attacks. They also include the protocol or service used to perpetrate the attack and the context in which the attack occurs. The following properties are specific to signature attacks, and you can configure them when configuring signature attack:

NOTE: Attack context, flow type, and direction are mandatory fields for the signature attack definition.

Attack Context

An attack context defines the location of the signature. If you know the service and the specific service context, specify that service and then specify the appropriate service contexts. If you know the service, but are unsure of the specific service context, specify one of the following general contexts:

- first-data-packet—Specify this context to detect the attack in only the first data packet.
- first-packet—Specify this context to detect the attack in only the first packet of a stream. When the flow direction for the attack object is set to any, the device checks the first packet of both the server-to-client and the client-to-server flows. If you know that the attack signature appears in the first packet of a session, choosing first packet instead of packet reduces the amount of traffic the device needs to monitor, which improves performance.
- packet—Specify this context to match the attack pattern within a packet. When you select this option, you must also specify the service binding to define the service header options . Although not required, specifying these additional parameters improves the accuracy of the attack object and thereby improves performance.
- line—Specify this context to detect a pattern match within a specific line within your network traffic.
- normalized-stream—Specify this context to detect the attack in an entire normalized stream. The normalized stream is one of the multiple ways of sending information. In this stream the information in the packet is normalized before a match is performed. Suppose www.yahoo.com/sports is the same as www.yahoo.com/s%70orts. The normalized form to represent both of these URLs might be www.yahoo.com/sports. Choose normalized stream instead of stream, unless you want to detect some pattern in its exact form. For example, if you want to detect the exact pattern www.yahoo.com/s%70orts, then select stream.
- normalized-stream256—Specify this context to detect the attack in only the first 256 bytes of a normalized stream.
- normalized-stream1k—Specify this context to detect the attack in only the first 1024 bytes of a normalized stream.
- normalized-stream-8k—Specify this context to detect the attack in only the first 8192 bytes of a normalized stream.
- stream—Specify this context to reassemble packets and extract the data to search for a pattern match. However, the device cannot recognize packet boundaries for stream contexts, so data for multiple packets is combined. Specify this option only when no other context option contains the attack.
- stream256—Specify this context to reassemble packets and search for a pattern match within the first 256 bytes of a traffic stream. When the flow direction is set to any, the device checks the first 256 bytes of both the server-to-client and

client-to-server flows. If you know that the attack signature will appear in the first 256 bytes of a session, choosing stream256 instead of stream reduces the amount of traffic that the device must monitor and cache, thereby improving performance.

- stream1k—Specify this context to reassemble packets and search for a pattern match within the first 1024 bytes of a traffic stream. When the flow direction is set to any, the device checks the first 1024 bytes of both the server-to-client and client-to-server flows. If you know that the attack signature will appear in the first 1024 bytes of a session, choosing stream1024 instead of stream reduces the amount of traffic that the device must monitor and cache, thereby improving performance.
- stream8k—Specify this context to reassemble packets and search for a pattern match within the first 8192 bytes of a traffic stream. When the flow direction is set to any, the device checks the first 8192 bytes of both the server-to-client and client-to-server flows. If you know that the attack signature will appear in the first 8192 bytes of a session, choosing stream8192 instead of stream reduces the amount of traffic that the device must monitor and cache, thereby improving performance.

Attack Direction

You can specify the connection direction of the attack. Using a single direction (instead of Any) improves performance, reduces false positives, and increases detection accuracy.

- Client to server (detects the attack only in client-to-server traffic)
- Server to client (detects the attack only in server-to-client traffic)
- Any (detects the attack in either direction)

Attack Pattern

Attack patterns are signatures of the attacks you want to detect. A signature is a pattern that always exists within an attack; if the attack is present, so is the signature. To create the attack pattern, you must first analyze the attack to detect a pattern (such as a segment of code, a URL, or a value in a packet header), then create a syntactical expression that represents that pattern. You can also negate a pattern. Negating a pattern means that the attack is considered matched if the pattern defined in the attack does *not* match the specified pattern.

NOTE: Pattern negation is supported for packet, line, and application based contexts only and not for stream and normalized stream contexts.

Protocol-Specific Parameters

Specifies certain values and options existing within packet headers. These parameters are different for different protocols. In a custom attack definition, you can specify

fields for only one of the following protocols—TCP, UDP, or ICMP. Although, you can define IP protocol fields with TCP or UDP in a custom attack definition.

NOTE: Header parameters can be defined only for attack objects that use a packet or first packet context. If you specified a line, stream, stream 256, or a service context you cannot specify header parameters.

If you are unsure of the options or flag settings for the malicious packet, leave all fields blank and IDP attempts to match the signature for all header contents.

Table 85 on page 554 displays fields and flags that you can set for attacks that use the IP protocol.

Table 85: IP Protocol Fields and Flags

Field	Description
Type of Service	Specify a value for the service type. Common service types are: ■ 0000 Default ■ 0001 Minimize Cost ■ 0002 Maximize Reliability ■ 0003 Maximize Throughput ■ 0004 Minimize Delay ■ 0005 Maximize Security
Total Length	Specify a value for the number of bytes in the packet, including all header fields and the data payload.
ID	Specify a value for the unique value used by the destination system to reassemble a fragmented packet.
Time to Live	Specify an integer value in the range of 0–255 for the time-to-live (TTL) value of the packet. This value represents the number of devices the packet can traverse. Each router that processes the packet decrements the TTL by 1; when the TTL reaches 0, the packet is discarded.
Protocol	Specify a value for the protocol used.
Source	Enter the source address of the attacking device.
Destination	Enter the destination address of the attack target.
Reserved Bit	This bit is not used.
More Fragments	When set (1), this option indicates that the packet contains more fragments. When unset (0), it indicates that no more fragments remain.
Don't Fragment	When set (1), this option indicates that the packet cannot be fragmented for transmission.

Table 86 on page 555 displays packet header fields and flags that you can set for attacks that use the TCP protocol.

Table 86: TCP Header Fields and Flags

Field	Description
Source Port	Specify a value for the port number on the attacking device.
Destination Port	Specify a value for the port number of the attack target.
Sequence Number	Specify a value for the sequence number of the packet. This number identifies the location of the data in relation to the entire data sequence.
ACK Number	Specify a value for the ACK number of the packet. This number identifies the next sequence number; the ACK flag must be set to activate this field.
Header Length	Specify a value for the number of bytes in the TCP header.
Data Length	Specify a value for the number of bytes in the data payload. For SYN, ACK, and FIN packets, this field should be empty.
Window Size	Specify a value for the number of bytes in the TCP window size.
Urgent Pointer	Specify a value for the urgent pointer. The value indicates that the data in the packet is urgent; the URG flag must be set to activate this field.
URG	When set, the urgent flag indicates that the packet data is urgent.
ACK	When set, the acknowledgment flag acknowledges receipt of a packet.
PSH	When set, the push flag indicates that the receiver should push all data in the current sequence to the destination application (identified by the port number) without waiting for the remaining packets in the sequence.
RST	When set, the reset flag resets the TCP connection, discarding all packets in an existing sequence.
SYN	When set, the SYN flag indicates a request for a new session.
FIN	When set, the final flag indicates that the packet transfer is complete and the connection can be closed.
R1	This reserved bit (1 of 2) is not used.
R2	This reserved bit (2 of 2) is not used.

Table 87 on page 556 displays packet header fields and flags that you can set for attacks that use the UDP protocol.

Table 87: UDP Header Fields and Flags

Field	Description
Source Port	Specify a value for the port number on the attacking device.
Destination Port	Specify a value for the port number of the attack target.
Data Length	Specify a value for the number of bytes in the data payload.

Table 88 on page 556 displays packet header fields and flags that you can set for attacks that use the ICMP protocol.

Table 88: ICMP Header Fields and Flags

Field	Description
ICMP Type	Specify a value for the primary code that identifies the function of the request or reply packet.
ICMP Code	Specify a value for the secondary code that identifies the function of the request or reply packet within a given type.
Sequence Number	Specify a value for the sequence number of the packet. This number identifies the location of the request or reply packet in relation to the entire sequence.
ICMP ID	Specify a value for the identification number. The identification number is a unique value used by the destination system to associate request and reply packets.
Data Length	Specify a value for the number of bytes in the data payload.

Sample Signature Attack Definition

The following is a sample signature attack definition:

```
<Entry>
<Name>sample-sig</Name>
<Severity>Major</Severity>
<Attacks><Attack>
<TimeBinding><Count>2</Count>
<Scope>dst</Scope></TimeBinding>
<Application>FTP</Application>
<Type>signature</Type>
<Context>packet</Context>
<Negate>true</Negate>
<Flow>Control</Flow>
<Direction>any</Direction>
<Headers><Protocol><Name>ip</Name>
<Field><Name>ttl</Name>
<Match>==</Match><Value>128</Value></Field>
</Protocol><Name>tcp</Name>
```

```
<Field><Name><Match>&lt;</Match>
<value>1500</Value>
</Field></Protocol></Headers>
</Attack></Attacks>
</Entry>
```

Attack Properties—Protocol Anomaly Attacks

A protocol anomaly attack object detects unknown or sophisticated attacks that violate protocol specifications (RFCs and common RFC extensions). You cannot create new protocol anomalies, but you can configure a new attack object that controls how your device handles a predefined protocol anomaly when detected.

NOTE: The service or application binding is a mandatory field for protocol anomaly attacks.

The following properties are specific to protocol anomaly attacks. Both attack direction and test condition are mandatory fields for configuring anomaly attack definitions.

Attack Direction

Attack direction allows you to specify the connection direction of an attack. Using a single direction (instead of Any) improves performance, reduces false positives, and increases detection accuracy:

- Client to server (detects the attack only in client-to-server traffic)
- Server to client (detects the attack only in server-to-client traffic)
- Any (detects the attack in either direction)

Test Condition

Test condition is a condition to be matched for an anomaly attack. Juniper Networks supports certain predefined test conditions. In the following example, the condition is a message that is too long. If the size of the message is longer than the preconfigured value for this test condition, the attack is matched.

```
<Attacks>
<Attack>
<Type>anomaly</Type>
...
<Test>MESSAGE_TOO_LONG</Test>
<Value>yes</Value>
...
</Attack>
</Attacks>
```

Sample Protocol Anomaly Attack Definition

The following is a sample protocol anomaly attack definition:

```
<Entry>
<Name>sample-anomaly</Name>
<Severity>Info</Severity>
<Attacks><Attack>
<TimeBinding><Count>2</Count>
<Scope>peer</Scope></TimeBinding>
<Application>TCP</Application>
<Type>anomaly</Type>
<Test>OPTIONS_UNSUPPORTED</Test>
<Direction>any</Direction>
</Attack></Attacks>
</Entry>
```

Attack Properties—Compound or Chain Attacks

A compound or chain attack object detects attacks that use multiple methods to exploit a vulnerability. This object combines multiple signatures and/or protocol anomalies into a single attack object, forcing traffic to match a pattern of combined signatures and anomalies within the compound attack object before traffic is identified as an attack. By combining and even specifying the order in which signatures or anomalies must match, you can be very specific about the events that need to take place before the device identifies traffic as an attack.

You must specify a minimum of 2 members (attacks) in a compound attack. You can specify up to 32 members in compound attack. Members can be either signature or anomaly attacks.

The following properties are specific to compound attacks:

Scope

Scope allows you to specify if the attack is matched within a session or across transactions in a session. If the specified service supports multiple transactions within a single session, you can also specify whether the match should occur over a single session or can be made across multiple transactions within a session:

- Specify *session* to allow multiple matches for the object within the same session.
- Specify *transaction* to match the object across multiple transactions that occur within the same session.

Order

Use ordered match to create a compound attack object that must match each member signature or protocol anomaly in the order you specify. If you do not specify an ordered match, the compound attack object still must match all members, but the attack pattern or protocol anomalies can appear in the attack in random order.

Reset

Specifies that a new log is generated each time an attack is detected within the same session. If this field is set to no then the attack is logged only once for a session.

Expression (Boolean expression)

Using the boolean expression field disables the ordered match function. The boolean expression field makes use of the member name or member index properties. The following three boolean operators are supported along with parenthesis, which helps determine precedence:

- or—If either of the member name patterns match, the expression matches.
- and—If both of the member name patterns match, the expression matches. It does not matter which order the members appear in.
- oand (ordered and)—If both of the member name patterns match, and if they appear in the same order as specified in the boolean expression, the expression matches.

Suppose you have created six signature members, labelled s1-s5. Suppose you know that the attack always contains the pattern s1, followed by either s2 or s3. You also know that the attack always contains s4 and s5, but their positions in the attack can vary. In this case, you might create the following boolean expression: **((s1 oand s2) or (s1 oand s3)) and (s4 and s5)**

NOTE: You can either define an ordered match or an expression (not both) in a custom attack definition.

Member Index

Member Index is specified in chain attacks to identify a member (attack) uniquely. In the following example, member index is used to identify the members m01 and m02 in the defined expression:

```
<Expression>m02 AND m01</Expression>
<Order>no</Order>
<Reset>no</Reset>
<ScopeOption/>
<Members>
<Attack>
<Member>m01</Member>
<Type>Signature</Type>
...
<Pattern><!CDATA[.*/getlatestversion]]></Pattern>
<Regex/>
</Attack>
<Attack><Member>m02</Member>
<Type>Signature</Type>
...
<Pattern><!CDATA[\[Skype\'.*]]></Pattern>
<Regex/>
</Attack>
<Attack>
```

NOTE: When defining the expression, you must specify the member index for all members.

Sample Compound Attack Definition

The following is a sample compound attack definition:

```
<Entry>
<Name>sample-chain</Name>
<Severity>Critical</Severity>
<Attacks><Attack>
<Application>HTTP</Application>
<Type>Chain</Type>
<Order>yes</Order>
<Reset>yes</Reset>
<Members><Attack>
<Type>Signature</Type>
<Context>packet</Context>
<Pattern><![CDATA[Unknown[]></Pattern>
<Flow>Control</Flow>
<Direction>cts</Direction>
</Attack><Attack>
<Type>anomaly</Type>
<Test>CHUNK_LENGTH_OVERFLOW</Test>
<Direction>any</Direction>
</Attack></Members>
</Attack></Attacks>
</Entry>
```

Related Topics

- Configuring Signature-Based Attacks on page 562
- Defining Rules for an IPS Rulebase on page 515

Understanding Protocol Decoders

Protocol decoders are used by IDP to check protocol integrity and protocol contextual information by looking for anomalies and ensuring that RFC standards are met. An anomaly can be any part of a protocol, such as the header, message body, or other individual fields that deviate from RFC standards for that protocol. For example, in the case of SMTP, if SMTP MAIL TO precedes SMTP HELO, that is an anomaly in the SMTP protocol.

When protocol contextual information is available, protocol decoders check for attacks within those contexts. For example, for SMTP, if an e-mail is sent to user@company.com, user@company.com is the contextual information and SMTP MAIL TO is the context. By using protocol contextual data, rather than the entire packet, for attack detection, protocol decoders improve overall performance and accuracy.

If there is a policy configured with a rule that matches the protocol decoder check for SMTP, the rule triggers and the appropriate action is taken.

This section contains the following topics:

- CLI Configuration Example on page 561
- Multiple IDP Detector Support on page 561

The IDP module ships with a preconfigured set of protocol decoders. These protocol decoders have default settings for various protocol-specific contextual checks they perform. You can use these defaults or you can tune them to meet your site's specific needs. To display the list of available protocol decoders, enter the following command:

```
user@host # show security idp sensor-configuration detector protocol-name ?
```

For a more detailed view of the current set of protocol decoders and their default context values, you can view the *detector-capabilities.xml* file located in the /var/db/idpd/sec-download folder on the device. When you download a new security package, you also receive this file which lists current protocols and default decoder context values.

The configuration instructions in this section provide a tunable context configuration example for one protocol decoder, FTP.

CLI Configuration Example

To configure protocol decoder tunables, refer to the following information:

1. View the list of protocols that have tunable parameters by entering the following command.

   ```
   user@host # set security idp sensor-configuration detector protocol-name
   ```

2. To configure tunable parameters for the protocol in question (in this case, FTP), enter the following:

   ```
   user@host # set security idp sensor-configuration detector protocol-name ftp
     tunable-name ftp_failed_logins tunable-value 4
   user@host # set security idp sensor-configuration detector protocol-name ftp
     tunable-name ftp_failed_flags tunable value 1
   user@host # set security idp sensor-configuration detector protocol-name ftp
     tunable-name ftp_line_length tunable-value 1024
   user@host # set security idp sensor-configuration detector protocol-name ftp
     tunable-name ftp_password_length tunable-value 64
   user@host # set security idp sensor-configuration detector protocol-name ftp
     tunable-name ftp_sitestring_length tunable-value 512
   user@host # set security idp sensor-configuration detector protocol-name ftp
     tunable-name ftp_username_length tunable-value 32
   ```

Multiple IDP Detector Support

When a new security package is received, it contains attack definitions and a detector. In any given version of a security package, the attack definitions correspond to the

capabilities of the included detector. When policy aging is disabled on the device (see the *reset-on-policy* command in the *JUNOS Software CLI Reference* for policy aging commands), only one policy is in effect at any given time. But if policy aging is enabled and there is a policy update, the existing policy is not unloaded when the new policy is loaded. Therefore, both policies can be in effect on the device. In this case, all existing sessions will continue to be inspected by existing policies and new sessions are inspected with new policies. Once all the existing sessions using the older policy have terminated or expired, the older policy is then unloaded.

When a policy is loaded, it is also associated with a detector. If the new policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

Note that a maximum of two detectors can be loaded at any given time. If two detectors are already loaded (by two or more policies), and loading a new policy requires also loading a new detector, then at least one of the loaded detectors must be unloaded before the new detector can be loaded. Before a detector is unloaded, all policies that use the corresponding detector are unloaded as well.

You can view the current policy and corresponding detector version by entering the following command:

```
user@host> show security idp status
```

Configuring Signature-Based Attacks

To configure a custom attack object, you specify a unique name for it and then specify additional information, which can make it easier for you to locate and maintain the attack object.

Certain properties in the attack object definitions are common to all types of attacks, such as attack name, severity level, service or application binding, time binding, and protocol or port binding. Some fields are specific to an attack type and are available only for that specific attack definition.

Signature attack objects use a stateful attack signature (a pattern that always exists within a specific section of the attack) to detect known attacks. They also include the protocol or service used to perpetrate the attack and the context in which the attack occurs. The following properties are specific to signature attacks, and you can configure them when configuring signature attack—attack context, attack direction,

attack pattern, and protocol-specific parameters (TCP, UDP, ICMP, or IP header fields).

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding Custom Attack Objects on page 544
   - Understanding Predefined Attack Objects and Groups on page 575
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.

When configuring signature-based attacks, keep the following in mind:

- Attack context and direction are mandatory fields for the signature attack definition.
- Pattern negation is supported for packet, line, and application-based contexts only and not for stream and normalized stream contexts.
- When configuring the protocol-specific parameters, you can specify fields for only one of the following protocols—IP, TCP, UDP, or ICMP.
- When configuring a protocol binding, you can specify only one of the following—IP, ICMP, TCP, UDP, RPC or applications.
  - IP—Protocol number is a mandatory field.
  - TCP and UDP—You can specify either a single port (minimum-port) or a port range (minimum-port and maximum-port). If you do not specify a port, the default value is taken (0-655325).
  - RPC—Program number is a mandatory field.

The configuration instructions in this topic describe how to create a signature-based attack object. In this example, you create a signature attack named sig1 and assign it the following properties:

- Recommended action (drop packet)—Specify to drop a matching packet before it can reach its destination but does not close the connection.
- Time binding—Specify the scope as source and count as 10. When scope is source, all attacks from the same source are counted, and when the number of attack reaches the count (10) specified, the attack is logged. In this example, every tenth attack from the same source is logged.
- Attack context (packet)—Specify to match the attack pattern within a packet.
- Attack direction (any)—Specify to detect the attack in both directions—client-to-server and server-to-client traffic.
- Protocol (TCP)—Specify time to live (TTL) value of 128.

- Shellcode (Intel)—Set the flag to detect shellcode for Intel platforms.
- Protocol binding—Specify TCP protocol and ports 50 through 100.

Once you have configured a signature-based attack object, you specify the attack as match criteria in an IDP policy rule. For more information, see "Defining Rules for an IPS Rulebase" on page 515.

You can use either J-Web or the CLI configuration editor to create a custom attack object.

This topic contains:

- CLI Configuration on page 564
- Related Topics on page 565

CLI Configuration

To create a signature-based attack object:

1. Specify a name for the attack. The following statement specifies sig1 as the name of the attack.

   user@host# **set security idp custom-attack sig1**

2. Specify common properties for the attack. The following statements specify a recommended action to drop packets and define time binding with scope as source scope and count as 10.

   user@host# **set security idp custom-attack sig1 recommended-action drop-packet**
   user@host#**set security idp custom-attack sig1 time-binding scope source count 10**

3. Specify the attack type and context. The following statement specifies the attack type signature and context packet.

   user@host# **set security idp custom-attack sig1 attack-type signature context packet**

4. Specify the attack direction and the shellcode flag. The following statement specifies the attack direction any and sets the shellcode flag to intel.

   user@host# **set security idp custom-attack sig1 attack-type signature shellcode intel**

5. Set the protocol and its fields. The following statement specifies the IP protocol and the TTL value 128.

   user@host# **set security idp custom-attack sig1 attack-type signature protocol ip ttl value 128 match equal**

6. Specify the protocol binding and ports. The following statement specifies the TCP protocol and the port range from 50 through 100.

```
user@host# set security idp custom-attack sig1 attack-type signature
    protocol-binding tcp minimum-port 50 maximum-port 100
```

7. If you are finished configuring the router, commit the configuration.
8. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- Configuring Protocol Anomaly-Based Attacks on page 565
- Understanding Protocol Decoders on page 560

Configuring Protocol Anomaly-Based Attacks

A protocol anomaly attack object detects unknown or sophisticated attacks that violate protocol specifications (RFCs and common RFC extensions). You cannot create new protocol anomalies, but you can configure a new attack object that controls how your device handles a predefined protocol anomaly when detected.

The following properties are specific to protocol anomaly attacks—attack direction and test condition.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding Custom Attack Objects on page 544
   - Understanding Predefined Attack Objects and Groups on page 575
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.

When configuring protocol anomaly-based attacks, keep the following in mind:

- The service or application binding is a mandatory field for protocol anomaly attacks. Besides the supported applications, services also include IP, TCP, UDP, ICMP, and RPC.
- The attack direction and test condition properties are mandatory fields for configuring anomaly attack definitions.

The configuration instructions in this topic describe how to create a signature-based attack object. In this example, you create a protocol anomaly attack named anomaly1 and assign it the following properties:

- Time binding—Specify the scope as peer and count as 2 to detect anomalies between source and destination IP addresses of the sessions for the specified number of times.
- Severity (info)—Specify to provide information about any attack that matches the conditions.
- Attack direction (any)—Specify to detect the attack in both directions—client-to-server and server-to-client traffic.
- Service (TCP)—Specify to match attacks using the TCP service.
- Test condition (OPTIONS_UNSUPPORTED)—Specify to match certain predefined test conditions. In this example, the condition is to match if the attack includes unsupported options.
- Shellcode (sparc)—Set the flag to detect shellcode for Sparc platforms.

Once you have configured the protocol anomaly-based attack object, you specify the attack as match criteria in an IDP policy rule. For more information, see "Defining Rules for an IPS Rulebase" on page 515.

You can use either J-Web or the CLI configuration editor to create a custom attack object.

This topic contains:

- CLI Configuration on page 566
- Related Topics on page 567

CLI Configuration

To create a protocol anomaly-based attack object:

1. Specify a name for the attack. The following statement specifies anomaly1 as the name of the attack.

   ```
   user@host# set security idp custom-attack anomaly1
   ```

2. Specify common properties for the attack. The following statements specify an info severity level and a time binding with a scope type peer and count 2.

   ```
   user@host# set security idp custom-attack anomaly1 severity info
   user@host#set security idp custom-attack anomaly1 time-binding scope peer
     count 2
   ```

3. Specify the attack type and test condition. The following statement specifies the attack type anomaly and test condition UNSUPPORTED_OPTIONS.

   ```
   user@host# set security idp custom-attack anomaly1 attack-type anomaly test
     UNSUPPORTED_OPTIONS
   ```

4. Specify other properties for the anomaly attack. The following statement specifies the service TCP and attack direction any, and sets the shellcode flag to sparc and specifies .

```
user@host# set security idp custom-attack sa attack-type anomaly service TCP
user@host# set security idp custom-attack sa attack-type anomaly direction
  any
user@host# set security idp custom-attack sa attack-type anomaly shellcode
  sparc
```

5. If you are finished configuring the router, commit the configuration.
6. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- Configuring Signature-Based Attacks on page 562

Specifying IDP Test Conditions for a Specific Protocol

When configuring IDP custom attacks, you can specify test conditions for a specific protocol. For example, to configure test conditions for ICMP:

1. List supported test conditions for ICMP and choose the one you want to configure. The supported test conditions are available in the CLI at the [edit security idp custom-attack test1 attack-type anomaly] hierarchy level.

```
user@router#set test icmp?

Possible completions:
 <test>               Protocol anomaly condition to be checked

  ADDRESSMASK_REQUEST
  DIFF_CHECKSUM_IN_RESEND
  DIFF_CHECKSUM_IN_RESPONSE
  DIFF_LENGTH_IN_RESEND
```

2. Configure the service for which you want to configure the test condition.

```
user@router# set service ICMP
```

3. Configure the test condition (specifying the protocol name is not required):

```
user@router# set test ADDRESSMASK_REQUEST
```

4. If you are finished configuring the router, commit the configuration.

Related Topics

- Configuring Protocol Anomaly-Based Attacks on page 565
- Attack Properties—Protocol Anomaly Attacks on page 557

Configuring DSCP in an IDP Policy

Configuring Differentiated Services code point (DSCP) values in IDP policies provides a method of associating class-of-service (CoS) values—thus different levels of reliability—for different types of traffic on the network.

Before You Begin

1. For background information, read:
   - Class of Service chapters in the *JUNOS Software Interfaces and Routing Configuration Guide*.
   - IDP Policies Overview on page 502
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Enable IDP application services in a security policy. See "Enabling IDP in a Security Policy" on page 536.
5. Create security zones. See "Creating Security Zones" on page 49.
6. Define rules. See "Defining Rules for an IPS Rulebase" on page 515.

DSCP is an integer value encoded in the 6-bit field defined in IP packet headers. It is used to enforce CoS distinctions. CoS allows you to override the default packet forwarding behavior and assign service levels to specific traffic flows.

You can configure DSCP value as an action in an IDP policy rule. You first define the traffic by defining match conditions in the IDP policy and then associate a DiffServ marking action with it. Based on the DSCP value, behavior aggregate classifiers set the forwarding class and loss priority for the traffic deciding the forwarding treatment the traffic receives.

All packets that match the IDP policy rule have the CoS field in their IP header rewritten with the DSCP value specified in the matching policy. If the traffic matches multiple rules with differing DSCP values, the first IDP rule that matches takes effect and this IDP rule then applies to all traffic for that session.

The configuration instructions in this topic describe how to create a policy called policy1, specify a rulebase for this policy, and then add a rule R1 to this rulebase. In this example, rule R1:

- Specifies the match condition to include any traffic from a previously configured zone called zone1 to another previously configured zone called zone2. The match condition also includes a predefined attack group called Critical - HTTP. The application setting in the match condition is specified as default and matches any application configured in the attack object.
- Specifies an action to rewrite the CoS field in the IP header with the DSCP value 50 for any traffic that matches the criteria for rule R1,

You can use either J-Web or the CLI configuration editor to configure the DSCP value in an IDP policy.

This topic contains:

- CLI Configuration on page 569
- Related Topics on page 570

CLI Configuration

To configure DSCP values in an IDP policy:

1. Create a policy by assigning a meaningful name to it. The following statement specifies policy1 as the policy name:

   user@host# **set security idp idp-policy policy1**

2. Associate a rulebase with the policy. The following statement associates an IPS rulebase with policy1:

   user@host# **set security idp idp-policy policy1 rulebase-ips**

3. Add rules to the rulebase The following statement adds a rule R1 to the rulebase:

   user@host# **set security idp idp-policy policy1 rulebase-ips rule R1**

4. Define the match criteria for the rule. The following statements specify that any traffic from zone1 to zone2 that includes a predefined attack group Critical - HTTP matches the criteria for rule R1. The default application setting matches any application configured in the attack object.

   user@host# **set security idp idp-policy policy1 rulebase-ips R1 match from-zone zone1 to-zone zone2 source-address any destination-address any application default**
   user@host# **set security idp idp-policy policy1 rulebase-ips R1 match attacks predefined-attack-group “Critical - HTTP”**

5. Specify an action for the rule. The following statement specifies that for all traffic matching the criteria defined for rule R1, the CoS field in the IP header is rewritten with the DSCP value 50:

   user@host# **set security idp idp-policy policy1 rulebase-ips R1 then action mark-diffserv 50**

6. Continue to specify any notification or logging options for the rule, if required.

7. Activate the policy. The following specifies policy1 as the active policy:

   user@host# **set security idp active-policy policy1**

8. If you are finished configuring the router, commit the configuration.
9. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an Exempt Rulebase on page 519
- Using Predefined Policy Templates on page 572

Chapter 16

IDP Signature Database

The signature database is one of the major components of Intrusion Detection and Prevention (IDP). It contains definitions of different objects—such as attack objects, application signatures objects, and service objects—that are used in defining IDP policy rules. As a response to new vulnerabilities, Juniper Networks periodically provides a file containing attack database updates on the Juniper website. You can download this file to protect your network from new threats.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding the IDP Signature Database on page 571
- Using Predefined Policy Templates on page 572
- Understanding Predefined Attack Objects and Groups on page 575
- Updating the Signature Database Overview on page 577
- Updating the Signature Database Manually on page 578
- Configuring a Security Package Update—Quick Configuration on page 580
- Updating the Signature Database Automatically on page 582
- Understanding the Signature Database Version on page 584
- Verifying the Signature Database on page 584

Understanding the IDP Signature Database

The IDP signature database is stored on the IDP enabled device and contains definitions of predefined attack objects and groups. These attack objects and groups are designed to detect known attack patterns and protocol anomalies within the network traffic. You can configure attack objects and groups as match conditions in IDP policy rules.

NOTE: You must install the IDP signature-database-update license key on your device for downloading and installing daily signature database updates provided by Juniper Networks. For license details, see the *JUNOS Software Administration Guide*.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rulebases on page 503
- Understanding IDP Policy Rules on page 506

You can perform the following tasks to manage the IDP signature database:

- Update the signature database—Download the attack database updates available on the Juniper Networks website. New attacks are discovered daily, so it is important to keep your signature database up to date.
- Verify the signature database version—Each signature database has a different version number with the latest database having the highest number. You can use the CLI to display the signature database version number.
- Update the protocol detector engine—You can download the protocol detector engine updates along with downloading the signature database. The IDP protocol detector contains Application Layer protocol decoders. The detector is coupled with the IDP policy and is updated together. It is always needed at policy update time, even if there is no change in the detector.
- Schedule signature database updates—You can configure the IDP-enabled device to automatically update the signature database after a set interval.

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- Understanding the Signature Database Version on page 584
- Updating the Signature Database Manually on page 578
- Updating the Signature Database Automatically on page 582

Using Predefined Policy Templates

Juniper Networks provides predefined policy templates that you can use as a starting point for creating your own policies. Each template is set of rules of a specific rulebase type that you can copy and then update according to your requirements. These templates are available in the templates.xml file on a secured Juniper Networks

website. To start using a template, you run a command from the CLI to download and copy this file to a /var/db/scripts/commit directory.

Before You Begin

1. For background information, read:
   - IDP Policies Overview on page 502
   - Understanding IDP Policy Rulebases on page 503
   - Understanding IDP Policy Rules on page 506
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.

Each policy template contains rules that use the default actions associated with the attack objects. You should customize these templates to work on your network by selecting your own source and destination addresses and choosing IDP actions that reflect your security needs.

Table 89 on page 573 summarizes the predefined IDP policy templates provided by Juniper Networks.

Table 89: Predefined IDP Policy Templates

Template Name	Description
All With Logging	Includes all Attack Objects and enables packet logging for all rules.
All Without Logging	Includes all Attack Objects but does not enable packet logging.
DMZ Services	Protects a typical demilitarized zone (DMZ) environment.
DNS Server	Protects Domain Name System (DNS) services.
File Server	Protects file sharing services, such as Network File System (NFS), FTP, and others.
Getting Started	Contains very open rules. Useful in controlled lab environments, but should not be deployed on heavy traffic live networks.
IDP Default	Contains a good blend of security and performance.
Recommended	Contains only the attack objects tagged as *recommended* by Juniper Networks. All rules have their Actions column set to take the recommended action for each attack object.
Web Server	Protects HTTP servers from remote attacks.

To use predefined policy templates:

- Download the policy templates from the Juniper Networks website.
- Install the policy templates.
- Enable the templates.xml script file. Commit scripts in the /var/db/scripts/commit directory are ignored if they are not enabled.
- Choose a policy template that is appropriate for you and customize it if you need to.
- Activate the policy that you want to run on the system. Activating the policy might take a few minutes. Even after a commit complete message is displayed in the CLI, the system might continue to compile and push the policy to the dataplane.

NOTE: Occasionally, the compilation process might fail for a policy. In this case, the active policy showing in your configuration might not match the actual policy running on your device. Run the show security idp status command to verify the running policy. Additionally, you can view the IDP log files to verify the policy load and compilation status (see "Verifying the Signature Database" on page 584).

- Delete or deactivate the commit script file. By deleting the commit script file, you avoid the risk of overwriting modifications to the template when you commit the configuration. Deactivating the statement adds an inactive tag to the statement, effectively commenting out the statement from the configuration. Statements marked inactive do not take effect when you issue the commit command.

You can use either J-Web or the CLI configuration editor to configure an application set.

This topic contains:

- CLI Configuration on page 574
- Related Topics on page 575

CLI Configuration

To download and use a predefined policy template:

1. Download the script file templates.xml to the/var/db/idpd/sec-download/sub-download directory. This script file contains predefined IDP policy templates.

   user@host> **request security idp security-package download policy-templates**

2. Copy the templates.xml file to the /var/db/scripts/commit directory and rename it to templates.xsl.

   user@host> **request security idp security-package install policy-templates**

3. Enable the templates.xsl scripts file. At commit time, the JUNOS management process (mgd) looks in the /var/db/scripts/commit directory for scripts and runs the script against the candidate configuration database to ensure the configuration conforms to the rules dictated by the scripts.

   user@host# **set system scripts commit file templates.xsl**

4. Commit the configuration. Committing the configuration saves the downloaded templates to the JUNOS configuration database and makes them available in the CLI at the [edit security idp idp-policy] hierarchy level.

5. Display the list of downloaded templates.

```
user@host#set security idp active-policy ?


Possible completions:
 <active policy> Set active policy
All_With_Logging
  All_Without_Logging
  DMZ_Services
  DNS_Service
  File_Server
  Getting_Started
  IDP_Default
  Recommended
  Web_Server
```

6. Activate the predefined policy. The following statement specifies the *Recommended* predefined IDP policy as the active policy:

   user@host# **set security idp active-policy Recommended**

7. Delete or deactivate the commit script file. By deleting the commit script file, you avoid the risk of overwriting modifications to the template when you commit the configuration. Run one of the following commands:

   user@host# **delete system scripts commit file templates.xsl**
   user@host# **deactivate system scripts commit file templates.xsl**

8. If you are finished configuring the router, commit the configuration.

9. You can verify the configuration by using the show security idp status command. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519

Understanding Predefined Attack Objects and Groups

The security package for IDP contains a database of predefined IDP attack objects and IDP attack object groups that you can use in IDP policies to match traffic against

known and unknown attacks. Juniper Networks updates the predefined attack objects and groups on a regular basis with newly discovered attack patterns.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding IDP Policy Rules on page 506
- Understanding IDP Rule Objects on page 508

Updates to the attack object database can include:

- New descriptions or severities for existing attack objects
- New attack objects
- Deletion of obsolete attack objects

Predefined Attack Objects

Predefined attack objects are listed in an alphabetical order. These attack objects have unique names that help you identify the attack. The first part of the name indicates the group to which the attack object belongs. For example:

- FTP:USER:ROOT—Belongs to the FTP:USER group. It detects attempts to log in to an FTP server using the root account.
- HTTP:HOTMAIL:FILE-UPLOAD—Belongs to the HTTP:HOTMAIL group. It detects files attached to e-mails sent via the Web-based e-mail service Hotmail.

Predefined Attack Object Groups

The predefined attack groups list displays the attack objects in the categories described below. A set of recommended attack objects that Juniper Networks considers to be serious threats are also available in this list. The recommended attack objects are organized into the following categories:

- Attack Type—Groups attack objects by type (anomaly or signature). Within each type, attack objects are grouped by severity.
- Category— Groups attack objects by predefined categories. Within each category, attack objects are grouped by severity.
- Operating System—Groups attack objects by the operating system to which they apply: BSD, Linux, Solaris, or Windows. Within each operating system, attack objects are grouped by services and severity.
- Severity—Groups attack objects by the severity assigned to the attack. IDP has five severity levels: Critical, Major, Minor, Warning, Info. Within each severity, attack objects are grouped by category.
- Web Services—Groups attack objects by common Web services. These services are grouped by severity levels—Warning, Critical, Major, Minor, Info.

- Miscellaneous—Groups attack objects by performance level. Attack objects affecting IDP performance over a certain level are grouped under this category.
- Response—Groups attack objects in traffic flowing in the server to client direction.

Related Topics

- Updating the Signature Database Manually on page 578
- Defining Rules for an IPS Rulebase on page 515
- Defining Rules for an Exempt Rulebase on page 519

Updating the Signature Database Overview

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks website. This database includes attack object groups that you can use in IDP policies to match traffic against known attacks. Although you cannot create, edit, or delete predefined attack objects, you can use the CLI to update the list of attack objects that you can use in IDP policies.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding the IDP Signature Database on page 571
- Understanding IDP Policy Rules on page 506
- Understanding Predefined Attack Objects and Groups on page 575

To update the signature database, you download a security package from the Juniper Networks website. The security package consists of the following IDP components:

- Attack objects
- Attack object groups
- Application objects
- Updates to the IDP Detector Engine
- IDP Policy templates (Policy templates are downloaded independently. See “Using Predefined Policy Templates” on page 572.)

By default, when you download the security package, you download the following components into a Staging folder in your device: the latest version of the complete attack object groups table, application objects table, and the updates to the IDP Detector Engine. Because the attack objects table is typically of a large size, by default the system downloads only updates to the attack objects table. However, you can download the complete attack objects table by using the full-update configuration option.

After downloading the security package, you must install the package to update the security database with the newly downloaded updates from the Staging folder in your device.

After installing a security package, when you commit the configuration, all policies are checked for their syntax (not only the active policy). This checking is the same as a commit check. If an attack configured in any of the existing policies is removed from the new signature database that you download, the commit check fails. When you update the IDP signature database, attacks configured in policies are not updated automatically. For example, suppose you configure a policy to include an attack FTP:USER:ROOT that is available in the signature database version 1200 on your system. Then, you download signature database version 1201, which no longer includes the attack FTP:USER:ROOT. Because an attack configured in your policy is missing from the newly downloaded database, the commit check in the CLI fails. To successfully commit your configuration, you must remove the attack (FTP:USER:ROOT) from your policy configuration.

CAUTION: IDP signature updates might fail if a new IDP policy load fails for any reason. When a new IDP policy load fails, the last known good IDP policy is loaded. Once the issue with the new policy load is resolved, and the new valid policy is active, signature updates will work properly.

Related Topics

- Understanding the Signature Database Version on page 584
- Updating the Signature Database Automatically on page 582
- Updating the Signature Database Manually on page 578
- Configuring a Security Package Update—Quick Configuration on page 580

Updating the Signature Database Manually

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks website. This database includes attack object groups that you can use in IDP policies to match traffic against known attacks. Although you cannot create, edit, or delete predefined attack objects, you can use the CLI to update the list of attack objects that you can use in IDP policies. After

downloading the security package, you must install the package to update the security database with the newly downloaded updates from the Staging folder in your device.

Before You Begin

1. For background information, read:
   - Updating the Signature Database Overview on page 577
   - IDP Policies Overview on page 502
   - Understanding the IDP Signature Database on page 571
   - Understanding Predefined Attack Objects and Groups on page 575
2. Establish basic connectivity. See the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.

The configuration instructions in this topic describe how to download the security package with the complete table of attack objects and attack object groups, create a policy, and specify the new policy as the active policy. This example then describes how to download only the updates that Juniper Networks has recently uploaded and then update the attack database, running policy, and detector with these new updates.

You can use either J-Web or the CLI configuration editor to manually download and update the signature database.

This topic contains:

- CLI Configuration on page 579
- Related Topics on page 580

CLI Configuration

To manually download and update the signature database:

1. Download the security package. The security package includes the detector and the latest attack objects and groups.

   user@host> **request security idp security-package download full-update**

2. Update the attack database, the active policy, and the detector with the new package.

   user@host> **request security idp security-package install**

3. Check the attack database update status with the following command. The command output displays information about the downloaded and installed versions of attack database versions.

   user@host> **request security idp security-package install status**

4. Commit the configuration.

5. After committing the configuration, the attack objects and groups are available in the CLI under the predefined-attack-groups and predefined-attacks configuration statements at the [edit security idp idp-policy] hierarchy level.
6. Associate attack objects or attack object groups with the policy. The following statement associates the recommended attack object group Response_Critical-TELNET with policy1:

   user@host# **set security idp idp-policy policy1 rulebase-ips rule rule1 match attacks predefined-attack-groups "Response_Critical - TELNET"**

7. Activate the policy. The following statement makes policy1 the active policy on the device:

   user@host# **set security idp active-policy policy1**

8. Commit the configuration.
9. After a week, if you want to download only the updates that Juniper Networks has recently uploaded, use the following command:

   user@host> **request security idp security-package download**

10. Update the attack database, active policy, detector with the new changes:

    user@host> **request security idp security-package install**

11. If you are finished configuring the router, commit the configuration.
12. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding the Signature Database Version on page 584
- Updating the Signature Database Automatically on page 582
- Updating the Signature Database Manually on page 578
- Configuring a Security Package Update—Quick Configuration on page 580

Configuring a Security Package Update—Quick Configuration

You can use J-Web Quick Configuration to quickly configure the Security Package Update.

Figure 103: Quick Configuration Page for Configuring Security Package Update

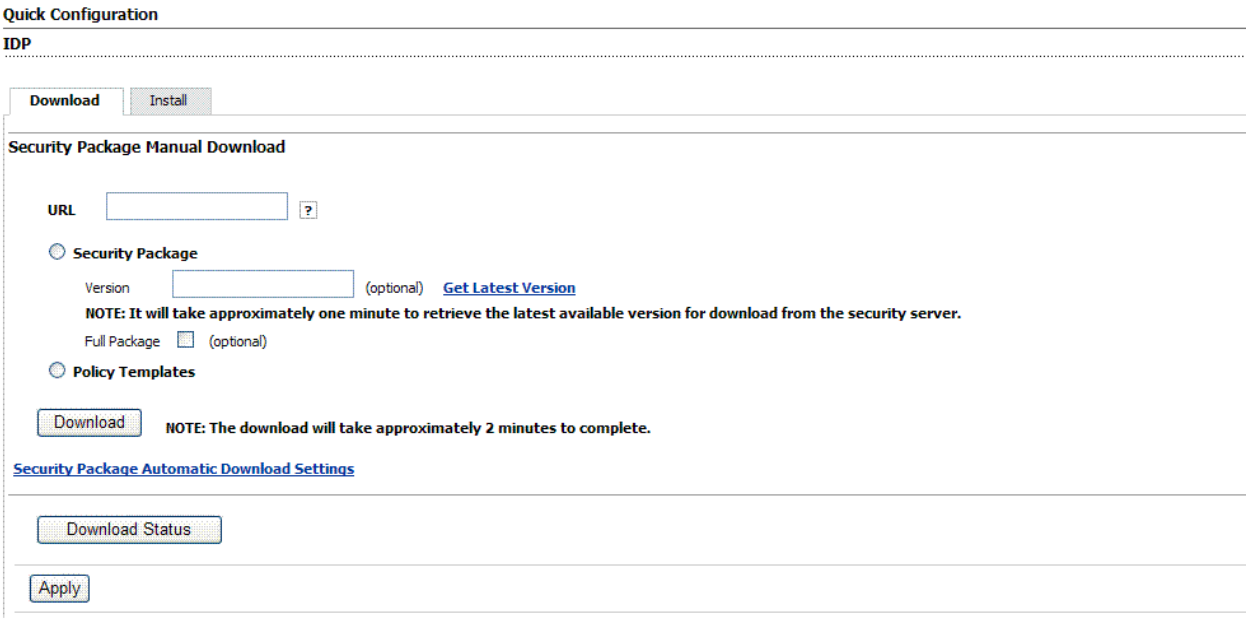

To configure security package update with Quick Configuration:

1. Select Configure>IDP>Signature Update.

   Figure 103 on page 581 shows the Security Package Manual Download page.
2. Fill in the information as described in Table 90 on page 581.
3. Click Apply.

Table 90: Security Package Update Quick Configuration Summary

Field	Function	Actions
Security Package Manual Download		
URL	Specifies the predefined default URL used by the device to download the signature database.	Type a URL with the following format: http://xmlexport.secteam.juniper.net.
Security Package	Specifies to manually download the updated signature database from the specified URL.	Click the option button.
Version (Optional)	Specifies the version number of the security package from the portal.	Type an integer.
Get Latest Version	Specifies to get the latest version of the signature database running on the device.	Click Get Latest Version to get the latest version of the signature database.
Full Package (Optional)	Specifies to enable the device to download the latest security package with the full set of attack signature tables from the portal.	Select the check box.

Table 90: Security Package Update Quick Configuration Summary *(continued)*

Field	Function	Actions
Policy Templates	Specifies to enable the device to download the latest policy templates from the portal.	Click the option button.
Download	Specifies to download the updated signature database.	Click Download.
Security Package Automatic Download Settings	Specifies to automatically download the updated signature database.	Click Security Package Automatic Download Settings.
URL	Specifies the predefined default URL used by the device to download the signature database.	Type a URL with the following format: http://xmlexport.secteam.juniper.net.
Interval	Specifies the amount of time that the device waits before updating the signature database. You should insert a default value.	Type a time interval between 24 and 336 hours.
Enable Schedule Update	Specifies to enable the device to automatically download the updated signature database from the specified URL.	Click the option button. By default, Schedule update is disabled.
Download Status	Specifies the download status of the security package or policy template.	Click Download Status.
Security Package Manual Installation		
Currently Installed Security Package		
Security Package Version	Specifies the information of the currently installed security package version.	Displays the currently installed security package version.
Detector	Specifies the version number of the IDP protocol detector currently running on the device.	Displays the version number of the IDP protocol detector.
Install Options		
Signature Update	Specifies to install the updated signature database.	Click the option button.
Do not set to active after installed	Specifies to activate the installed security package.	Select the check box.
Policy Templates	Specifies to install the latest policy templates from the portal.	Click the option button.
Install	Specifies to install the updated security package.	Click Install.
Install Status	Specifies the status of the installation of the security package.	Click Install Status.

Updating the Signature Database Automatically

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks website. This database includes attack object

groups that you can use in IDP policies to match traffic against known attacks. You can configure your device to download the signature database updates automatically at a specified interval.

Before You Begin

For background information, read:

- Updating the Signature Database Overview on page 577
- Understanding the IDP Signature Database on page 571

The configuration instructions in this topic describe how to download the security package with the complete table of attack objects and attack object groups every 48 hours starting at 11:59 pm on December 10.

You can use either J-Web or the CLI configuration editor to update the signature database automatically.

This topic contains:

- CLI Configuration on page 583
- Related Topics on page 584

CLI Configuration

To download and update predefined attack objects:

1. Specify the URL for the security package. The security package includes the detector and the latest attack objects and groups. The following statement specifies http://sec-pack.juniper.net as the URL for downloading signature database updates:

   user@host# **set security idp security-package url http://sec-pack.juniper.net**

2. Specify the time and interval for download. The following statement sets the interval as 48 hours and the start time as 11:59 pm on December 10:

   user@host# **set security idp security-package automatic interval 48 start-time 12-10.23:59**

3. Enable an automatic download and update of the security package.

   user@host# **set security idp security-package automatic enable**

4. If you are finished configuring the router, commit the configuration.
5. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding the Signature Database Version on page 584
- Updating the Signature Database Manually on page 578
- Configuring a Security Package Update—Quick Configuration on page 580

Understanding the Signature Database Version

New attack objects are added to the signature database server frequently; downloading these updates and installing them on your managed devices regularly ensures that your network is effectively protected against the latest threats. As new attack objects are added to the signature database server, the version number of the database is updated with the latest database version number. Each signature database has a different version number with the latest database having the highest number.

Before You Begin

For background information, read:

- Understanding the IDP Signature Database on page 571

When updating the signature database, the signature database update client connects to the Juniper Networks website and obtains the update using an HTTPS connection. This update—difference between the existing signature database and latest signature database—is calculated based on the version number that is assigned to each signature database. After you download the updates, the updated information is merged with the existing signature database and the version number is set to that of the latest signature database.

Related Topics

- Verifying the Signature Database on page 584
- Updating the Signature Database Manually on page 578
- Updating the Signature Database Automatically on page 582

Verifying the Signature Database

This topic contains:

- Verifying the Policy Compilation and Load Status on page 585
- Verifying the Signature Database Version on page 586

Verifying the Policy Compilation and Load Status

Purpose Display the IDP log files to verify the IDP policy load and compilation status. When activating an IDP policy, you can view the IDP logs and verify if the policy is loaded and compiled successfully.

Action To track the load and compilation progress of an IDP policy, configure either one or both of the following in the CLI:

- You can configure a log file, which will be located in /var/log/, and set trace option flags to record these operations:

```
user@host# set security idp traceoptions file idpd
user@host# set security idp traceoptions flag all
```

- You can configure your device to log system log messages to a file in the /var/log directory:

```
user@host# set system syslog file messages any any
```

After committing the configuration in the CLI, enter either of the following commands from the shell prompt in the UNIX-level shell:

Sample Output

```
user@host> start shell
user@host% tail -f /var/log/idpd
Jun  9 18:15:40 logmsg <valid license found for feature 20>
Jun  9 18:15:40 IDP feature license status: Valid license installed.
Jun  9 18:15:40 idpd commit start...
Jun  9 18:15:40 Entering enable processing.
Jun  9 18:15:40 Enable value (default)
Jun  9 18:15:40 IDP processing default.
...
Jun  9 18:15:40 Apply policy configuration, policy ops bitmask = 45
Jun  9 18:15:40 Starting policy (idpengine) compile...
Jun  9 18:16:10 policy compilation memory estimate: 57126048
Jun  9 18:16:10 ...Passed       (Shows that the policy compilation is
successful)Jun  9 18:16:10 Starting policy package...
Jun  9 18:16:12 ...Policy Packaging Passed
Jun  9 18:16:12 Starting policy load...
Jun  9 18:16:12 Loading policy(/var/db/idpd/bins/idpengine.bin.gz.v +
/var/db/idpd/sec-repository/libidp-detector.so.gz.v +
/var/db/idpd/bins/compiled_ai.bin)...
Jun  9 18:16:12 idpd_dev_add_ipc_connection called..
...
Jun  9 18:16:20 Reading sensor config...
Jun  9 18:16:20 sensor/idp node does not exist, apply defaults
Jun  9 18:16:20 idpd_dev_add_ipc_connection called...
Jun  9 18:16:20 idpd_dev_add_ipc_connection: done.
...
Jun  9 18:16:20 sensor conf successful
Jun  9 18:16:20
...idpd commit end

Jun  9 18:16:20 Returning from commit mode, status = 0. (Shows the policy load
is successful)
```

Sample Output

```
user@host> start shell
```

```
user@host% tail -f /var/log/messages

Jun 24 17:34:38  turtlebert mgd[4786]: UI_COMMIT_PROGRESS: Commit operation in
progress: activating '/var/run/db/juniper.data'
Jun 24 17:34:38  turtlebert mgd[4786]: UI_COMMIT_PROGRESS: Commit operation in
progress: notifying daemons of new configuration
Jun 24 17:34:38  turtlebert mgd[4786]: UI_COMMIT_PROGRESS: Commit operation in
progress:  notifying idpd(62)
Jun 24 17:34:38  turtlebert mgd[4786]: UI_COMMIT_PROGRESS: Commit operation in
progress: signaling 'IDP policy daemon', pid 4699, signal 1, status 0 with
notification errors enabled
...
Jun 24 17:34:45  turtlebert idpd[4699]: IDP_POLICY_LOAD_SUCCEEDED: IDP
policy[/var/db/idpd/bins/test.bin.gz.v] and
detector[/var/db/idpd/sec-repository/libidp-detector.so.gz.v] loaded successfully.
IDPD Trace file:
...
Jun 24 12:10:27 idpd_policy_load: idp policy pre-install succeeded
Jun 24 12:10:27 idpd_comm_server_get_event:478: evGetNext got event.
Jun 24 12:10:27 idpd_comm_server_get_event:486: evDispatch OK
...
Jun 24 12:10:27 idpd_policy_load: idp policy install succeeded
Jun 24 12:10:27 idpd_comm_server_get_event:486: evDispatch OK
...
Jun 24 12:10:27 idpd_policy_load: idp policy post-install succeeded
Jun 24 12:10:28 Reading sensor config...
Jun 24 12:10:28 sensor/idp node does not exist, apply defaults


Jun 24 12:10:28 sensor conf successful
Jun 24 12:10:28

...idpd commit end
Jun 24 12:10:28 Returning from commit mode, status = 0.
```

Meaning Displays log messages showing the procedures that run in the background after you commit the set security idp active-policy command. This sample output shows that the policy compilation, sensor configuration, and policy load are successful.

Related Topics To display the policy that is currently active, run the show security idp status command. For a complete description of this command, see the *JUNOS Software CLI Reference*.

Verifying the Signature Database Version

Purpose Display the signature database version.

Action From the operational mode in the CLI, enter show security idp security-package-version.

Sample Output

```
user@host> show security idp security-package-version
Attack database version:31(Wed Apr 16 15:53:46 2008)
  Detector version :9.1.140080400
  Policy template version :N/A
```

Meaning The output displays the version numbers for the signature database, protocol detector, and the policy template on the IDP-enabled device. Verify the following information:

- Attack database version—On April 16, 2008, the version of the signature database active on the device is 31.
- Detector version—Displays the version number of the IDP protocol detector currently running on the device.
- Policy template version—Displays the version of the policy template that is installed in the /var/db/scripts/commit directory when you run the request security idp security-package install policy-templates configuration statement in the CLI.

Related Topics For a complete description of show security idp security-package-version output, see the *JUNOS Software CLI Reference*.

Chapter 17

IDP Application Identification

Juniper Networks provides predefined application signatures that detect TCP and UDP applications running on non-standard ports. Identifying these applications allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on non-standard ports. It also improves performance by narrowing the scope of attack signatures for applications without decoders.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- IDP Application Identification Support on Different Device Types on page 589
- Understanding Application Identification on page 590
- Understanding Service and Application Bindings on page 591
- Understanding Application System Cache on page 593
- Configuring IDP Policies for Application Identification on page 593
- Disabling Application Identification on page 595
- Setting Memory and Session Limits on page 596
- Verifying Application Identification on page 598

IDP Application Identification Support on Different Device Types

The following table lists key IDP application identification features, specifies whether the features are supported on various device types, and indicates where you can find more information about each feature.

Table 91: Support Information: IDP Application Identification

Feature	J Series Services Routers	SRX Series Services Gateways	More Information
Application identification	No	Yes	"Understanding Application Identification" on page 590
Service and application bindings	No	Yes	"Understanding Service and Application Bindings" on page 591

Table 91: Support Information: IDP Application Identification *(continued)*

Feature	J Series Services Routers	SRX Series Services Gateways	More Information
Application system cache	No	Yes	"Understanding Application System Cache" on page 593

Understanding Application Identification

The IDP sensor monitors the network and detects suspicious and anomalous network traffic based on specific rules defined in IDP rulebases. It applies attack objects to traffic based on protocols or applications. Application signatures enable the sensor to identify known and unknown applications running on non-standard ports and to apply the correct attack objects.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding the IDP Signature Database on page 571

Application signatures are available as part of the security package provided by Juniper Networks. You download predefined application signatures along with the security package updates. You cannot create application signatures. For information on downloading the security package, see "Updating the Signature Database Manually" on page 578.

The application signatures identify an application by matching patterns in the first packet of a session. The IDP sensor matches patterns for both client-to-server and server-to-client sessions.

Application identification is enabled by default and is automatically turned on when you configure the default application in the IDP policy. However, when you specify an application in the policy rule, application identification is disabled and attack objects are applied based on the specified application. This specific application configuration overwrites the automatic identification process. For instructions on specifying applications in policy rules, see "Configuring Applications or Services for IDP" on page 533.

Related Topics

- Understanding Service and Application Bindings on page 591
- Configuring IDP Policies for Application Identification on page 593
- Disabling Application Identification on page 595
- Setting Memory and Session Limits on page 596

Understanding Service and Application Bindings

Attack objects can bind to applications and services in different ways:

- Attack objects can bind to an application implicitly and not have a service definition. They bind to an application based on the name of a context or anomaly.
- Attack objects can bind to a service using a service name.
- Attack objects can bind to a service using TCP or UDP ports, ICMP types or codes or RPC program numbers.

Before You Begin

For background information, read:

- IDP Policies Overview on page 502
- Understanding the IDP Signature Database on page 571

Whether the specified application or service binding applies or not depends on the complete attack object definition as well as the IDP policy configuration:

- If you specify an application in an attack object definition, the service field is ignored. The attack object binds to the application instead of the specified service. However, if you specify a service and no application in the attack object definition, the attack object binds to the service. Table 92 on page 591 summarizes the behavior of application and service bindings with application identification.

Table 92: Applications and Services with Application Identification

Attack Object Fields	Binding Behavior	Application Identification
:application (http) **:service (smtp)**	■ Binds to the application HTTP. ■ The service field is ignored.	Enabled
:service (http)	Binds to the application HTTP.	Enabled
:service (tcp/80)	Binds to TCP port 80.	Disabled

For example in the following attack object definition, the attack object binds to the application HTTP, the application identification is enabled, and the service field SMTP is ignored.

```
: ("http-test"
   :application ("http")
   :service ("smtp")
   :rectype (signature)
```

```
        :signature (
          :pattern (".*TERM=xterm; export TERM=xterm; exec bash – i\x0a\x.*")
          :type (stream)
        )
        :type (attack-ip)
      )
```

- If an attack object is based on service specific contexts (for example http-url) and anomalies (for example tftp_file_name_too_long), both application and service fields are ignored. Service contexts and anomalies imply application, thus when you specify these in the attack object, application identification is applied.
- If you configure a specific application in a policy, you overwrite the application binding specified in an attack object. Table 93 on page 592 summarizes the binding with the application configuration in the IDP policy.

Table 93: Application Configuration in an IDP Policy

Application Type in the Policy	Binding Behavior	Application Identification
Default	Binds to the application or service configured in the attack object definition.	■ Enabled for application-based attack objects. ■ Disabled for service-based attack objects.
Specific application	Binds to the application specified in the attack object definition.	Disabled
Any	Binds to all applications.	Disabled

- If you specify an application in an IDP policy, the application type configured in the attack object definition and in the IDP policy must match. The policy rule cannot specify two different applications (one in the attack object and the other in the policy).

Related Topics

- Configuring IDP Policies for Application Identification on page 593
- Disabling Application Identification on page 595

Understanding Application System Cache

Application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service.

Before You Begin

For background information, read:

- Understanding Application Identification on page 590
- IDP Policies Overview on page 502
- Understanding the IDP Signature Database on page 571

Once an application is identified, its information is saved in the cache so that only one pattern matching is required for an application running on a particular system, thereby expediting the identification process.

A mapping is saved in the cache only if the matched signature contains both client-to-server and server-to-client patterns. This process protects the system from hackers who might send malicious packets through a legitimate server port so that it is interpreted as a different application.

By default, the application system cache saves the mapping information for 3600 seconds. However, you can configure the cache timeout value by using the CLI.

To minimize the impact on performance, application system cache is refreshed only when TCP or UDP traffic triggers a cache lookup. Without a cache lookup, the entries in the ASC remain unchanged even after cache timeout.

Related Topics

- Understanding Service and Application Bindings on page 591
- Configuring IDP Policies for Application Identification on page 593

Configuring IDP Policies for Application Identification

For application identification to work, you must choose the default configuration option as the application type in an IDP policy rule. If you specify an application

instead, the application identification feature is disabled and IDP matches the traffic with the specified application.

Before You Begin

1. For background information, read:
   - Understanding Application Identification on page 590
   - IDP Policies Overview on page 502
   - Understanding the IDP Signature Database on page 571
2. Establish basic connectivity. See the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Download the signature database. See "Updating the Signature Database Manually" on page 578

The configuration instructions in this topic describe how to configure IDP policy for application identification.

You can use either J-Web or the CLI configuration editor to configure IDP policy for application identification.

This topic contains:

- CLI Configuration on page 594
- Related Topics on page 595

CLI Configuration

To configure IDP policy for application identification:

1. Create an IDP policy, associate a rulebase with the policy, and define rules in the rulebase. The following statement creates an IDP policy ABC and defines rule 123 in the IPS rulebase:

   ```
   user@host# set security idp idp-policy ABC rulebase-ips rule 123
   ```

2. Specify the application type as a match condition in the policy. The following statement specifies default as the application type:

   ```
   user@host# set security idp idp-policy ABC rulebase-ips rule 123 match
     application default
   ```

3. Continue to configure other match conditions and actions for the policy (see "Defining Rules for an IPS Rulebase" on page 515).
4. If you are finished configuring the router, commit the configuration.
5. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Disabling Application Identification on page 595
- Setting Memory and Session Limits on page 596
- Verifying Application Identification on page 598

Disabling Application Identification

Application identification is enabled by default. You can disable application identification with the CLI.

Before You Begin

For background information, read:

- Understanding Application Identification on page 590
- IDP Policies Overview on page 502
- Understanding the IDP Signature Database on page 571

The configuration instructions in this topic describe how to disable application identification.

You can use either J-Web or the CLI configuration editor to disable application identification.

This topic contains:

- CLI Configuration on page 595
- Related Topics on page 596

CLI Configuration

To disable and application identification:

1. Specify the disable configuration option.

   user@host# **set security idp sensor-configuration application-identification disable**

2. If you want to reenable application identification, delete the configuration statement that specifies disabling of application identification.

   user@host# **delete security idp sensor-configuration application-identification disable**

3. If you are finished configuring the router, commit the configuration.
4. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding Service and Application Bindings on page 591
- Configuring IDP Policies for Application Identification on page 593
- Setting Memory and Session Limits on page 596
- Verifying Application Identification on page 598

Setting Memory and Session Limits

Although you cannot create application signatures, you can configure sensor settings to limit the number of sessions running application identification and also limit memory usage for application identification.

Before You Begin

1. For background information, read:
   - Understanding Application Identification on page 590
   - IDP Policies Overview on page 502
   - Understanding the IDP Signature Database on page 571
2. Establish basic connectivity. See the Getting Started Guide for your device.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Download the signature database. See "Updating the Signature Database Manually" on page 578

- Memory limit for a session—You can configure the maximum amount of memory bytes that can be used to save packets for application identification for one TCP or UDP session. You can also configure a limit for global memory usage for application identification. Application identification is disabled for a session after the system reaches the specified memory limit for the session. However, IDP continues to match patterns. The matched application is saved to cache so that the next session can use it. This protects the system from attackers trying to bypass application identification by purposefully sending large client-to-server packets.
- Number of sessions—You can configure the maximum number of sessions that can run application identification at the same time. Application identification is disabled after the system reaches the specified number of sessions. You limit the number of sessions so that you can prevent a denial-of-service (DOS) attack,

which occurs when too many connection requests overwhelm and exhaust all the allocated resources on the system.

Table 94 on page 597 provides the capacity of a central point (CP) session numbers for SRX3400, SRX3600, SRX5600, and SRX5800 devices.

Table 94: Maximum CP Session Numbers

SRX Series Devices	Maximum Sessions	Central Point (CP)
SRX3400	2.25 million	Combo-mode CP
SRX3600	2.25 million	Combo-mode CP
SRX5600	10 million	Full CP
	2.25 million	Combo-mode CP
SRX5800	10 million	Full CP
	2.25 million	Combo-mode CP

In the configuration instructions for this example, you configure the limit so that only 600 sessions can run application identification at the same time. You also configure 5000 memory bytes as the maximum amount of memory that can be used for saving packets for application identification for one TCP session.

You can use the CLI configuration editor to configure memory and session limits for application identification.

This topic contains:

- CLI Configuration on page 597
- Related Topics on page 598

CLI Configuration

To configure memory and session limits for application identification:

1. Specify the session limit for application identification. In the following statement you set the maximum number of sessions that can run application identification at the same time as 600:

   user@host# **set security idp sensor-configuration application-identification max-sessions 600**

2. Specify the memory limit for application identification. In the following statement you configure a maximum of 5000 memory bytes to save packets for application identification:

   user@host# **set security idp sensor-configuration application-identification max-tcp-session-packet-memory 5000**

3. If you are finished configuring the router, commit the configuration.
4. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Understanding Application Identification on page 590
- Understanding Service and Application Bindings on page 591
- Configuring IDP Policies for Application Identification on page 593
- Verifying Application Identification on page 598

Verifying Application Identification

This topic covers:

- Verifying the Application System Cache on page 598
- Verifying Application Identification Counters on page 599

Verifying the Application System Cache

Purpose Verify the IDP application system cache (ASC) statistics.

Action From the CLI, enter the show security idp application-identification application-system-cache command.

Sample Output

```
user@host> show security idp application-identification application-system-cache
IDP Application System Cache statistics:

   Vsys-ID     IP address        Port     Protocol    Service

  0           20.0.0.4          23       tcp         TELNET
  0           20.0.0.6          23       tcp         TELNET
  0           20.0.0.2          23       tcp         TELNET
  0           20.0.0.2          25       tcp         SMTP
  0           20.0.0.6          25       tcp         SMTP
  0           20.0.0.4          25       tcp         SMTP
  0           20.0.0.3          135      tcp         MSRPC
  0           20.0.0.5          139      tcp         SMB
  0           20.0.0.7          139      tcp         SMB
  0           20.0.0.3          143      tcp         IMAP
  0           20.0.0.5          143      tcp         IMAP
  0           20.0.0.3          139      tcp         SMB
  0           20.0.0.7          143      tcp         IMAP
  0           20.0.0.3          80       tcp         HTTP
  0           20.0.0.5          80       tcp         HTTP
  0           20.0.0.7          80       tcp         HTTP
```

Meaning The output shows a summary of the ASC statistics information. Verify the following information:

- Vsys-ID—Displays the virtual system identification number.
- IP address—Displays the destination address.
- Port—Displays the destination port on the server.
- Service—Displays the name of the service or application identified on the destination port.

Related Topics For a complete description of show security idp application-identification application-system-cache output, see the *JUNOS Software CLI Reference*.

Verifying Application Identification Counters

Purpose Verify the IDP counters for the application identification processes.

Action From the CLI, enter the show security idp counters application-identification command.

Sample Output

```
user@host> show security idp counters application-identification
IDP counters:

 IDP counter type                                          Value
 AI cache hits                                             2682
 AI cache misses                                           3804
 AI matches                                                74
 AI no-matches                                             27
 AI-enabled sessions                                       3804
 AI-disabled sessions                                      2834
 AI-disabled sessions due to cache hit                     2682
 AI-disabled sessions due to configuration                 0
 AI-disabled sessions due to protocol remapping            0
 AI-disabled sessions due to non-TCP/UDP flows             118
 AI-disabled sessions due to no AI signatures              0
 AI-disabled sessions due to session limit                 0
 AI-disabled sessions due to session packet memory limit   34
 AI-disabled sessions due to global packet memory limit    0
```

Meaning The output shows a summary of the application identification counters. Verify the following information:

- AI cache hits—Displays the number of hits on the application identification cache
- AI cache misses—Displays the number of times the application matches but the application identification cache entry is not added.
- AI matches—Displays the number of times the application matches, and an application identification cache entry is added.
- AI no-matches—Displays the number of times when application does not match.
- AI-enabled sessions—Displays the number of sessions on which application identification is enabled.
- AI-disabled sessions—Displays the number of sessions on which application identification is enabled.

- AI-disabled sessions due to cache hit—Displays the number of sessions on which application identification is disabled after a cache entry is matched. Application identification process is discontinued for this session.
- AI-disabled sessions due to configuration—Displays the number of sessions on which application identification is disabled because of the sensor configuration.
- AI-disabled sessions due to protocol remapping—Displays the number of sessions for which application identification is disabled because you have configured a specific service in the IDP policy rule definition.
- AI-disabled sessions due to non-TCP/UDP flows—Displays the number of sessions for which application identification is disabled because the session is not a TCP or UDP session.
- AI-disabled sessions due to no AI signatures—Displays the number of sessions for which application identification is disabled because no match is found on the application identification signatures.
- AI-disabled due to session limit—Displays the number of sessions for which application identification is disabled because sessions have reached the maximum limit configured. Application identification is disabled for future sessions too.
- AI-disabled due to session packet memory limit—Displays the sessions for which application identification is disabled because sessions have reached the maximum memory limit on TCP or UDP flows. Application identification is disabled for future sessions too.
- AI-disabled due to global packet memory limit—Displays the sessions for which application identification is disabled because the maximum memory limit is reached. Application identification is disabled for future sessions too.

Related Topics For a complete description of show security idp counters output, see the *JUNOS Software CLI Reference*.

Chapter 18
IDP SSL Inspection

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- IDP SSL Overview on page 601
- Supported Ciphers on page 602
- Key Exchange on page 603
- Server Key Management and Policy Configuration on page 603
- Displaying Keys and Servers on page 604
- Adding Keys and Servers on page 604
- Deleting Keys and Servers on page 605
- Configuring SSL Inspection on page 605

IDP SSL Overview

Secure Sockets Layer (SSL), also called Transport Layer Security (TLS), is a protocol suite for Web security that provides authentication, confidentiality and message integrity. Authentication guards against fraudulent transmissions by enabling a Web browser to validate the identity of a Web server. Confidentiality mechanisms ensure that communications are private. SSL enforces confidentiality by encrypting data to prevent unauthorized users from eavesdropping on electronic communications. Finally, message integrity ensures that the contents of a communication have not been tampered with.

Each SSL session begins with a handshake during which the client and server agree on the specific security key and the encryption algorithms to use for that session. At this time, the client also authenticates the server. Optionally, the server can authenticate the client. Once the handshake is complete, transfer of encrypted data can begin.

Juniper Networks provides Intrusion Detection and Prevention (IDP) SSL inspection that uses the SSL protocol suite consisting of different SSL versions, ciphers, and key exchange methods. Combined with the Application Identification feature, the SSL Inspection feature enables SRX Series devices to inspect HTTP traffic encrypted in SSL on any port. The following SSL protocols are supported:

- SSLv2

- SSLv3
- TLS

Supported Ciphers

An SSL cipher comprises encryption cipher, authentication method, and compression. JUNOS Software supports all OPENSSL supported ciphers that do not involve the use of temporary private keys. For authentication, NULL, MD5, and SHA-1 authentication methods are supported.

NOTE: Compression and SSLv2 ciphers are not supported. Currently, most SSL servers automatically upgrade to a TLS cipher when an SSLv2 cipher is received in a client "hello" message. Check your browser to see how strong the ciphers can be and which ones your browser supports. (If the cipher is not in the list of supported ciphers, the session is ignored for deep packet inspection.)

Table 95 on page 602 shows the encryption algorithms supported by the SRX Series devices.

Table 95: Supported Encryption Algorithms

Cipher	Exportable	Type	Key Material	Expanded Key Material	Effective Key Bits	IV Size
NULL	No	Stream	0	0	0	N/A
DES_CBC	No	Block	8	8	56	8
3DES_EDE_CBC	No	Block	24	24	168	8
AES_128_CBC	No	Block	16	16	128	16
AES_256_CBC	No	Block	32	32	256	16

For more information on encryption algorithms, see "Encapsulating Security Payload (ESP) Protocol" on page 359. Table 96 on page 602 shows the supported SSL ciphers.

Table 96: Supported SSL Ciphers

Version	Cipher Suites	Value

Table 96: Supported SSL Ciphers *(continued)*

	TLS_RSA_WITH_NULL_MD5	0x0001
	TLS_RSA_WITH_NULL_SHA	0x0002
	TLS_RSA_WITH_DES_CBC_SHA	0x0009
	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0x000A
	TLS_RSA_WITH_AES_128_CBC_SHA	0x002F
	TLS_RSA_WITH_AES_256_CBC_SHA	0x0035

NOTE: RC4 and IDEA ciphers are not supported because of license and OPENSSL library availability.

Key Exchange

Internet Key Exchange establishes a premaster secret that is used to generate symmetric keys for bulk data encryption and authentication. Section F.1.1 of RFC 2246 defines TLS authentication and key exchange methods. The two key exchange methods are:

- RSA—A key exchange algorithm that governs the way participants create symmetric keys or a secret that is used during an SSL session. RSA key exchange algorithm is the most commonly used method.
- Diffie-Hellman—A Diffie-Hellman (DH) key exchange method allows the participants to produce a shared secret value. The strength of the technique is that it allows the participants to create the secret value over an unsecured medium without passing the secret value through the wire.

Both RSA and Diffie-Hellman key exchange methods can use either a fixed or a temporary server key. IDP can successfully retrieve the premaster secret only if a fixed server key is used. JUNOS Software supports only the RSA key exchange method. For more information on Internet Key Exchange, see “Understanding Public Key Cryptography” on page 402.

NOTE: Juniper IDP does not decrypt SSL sessions that use Diffie-Hellman key exchange.

Server Key Management and Policy Configuration

The device can support up to 1000 server private keys. Each key can have up to 100 servers that use it. This capacity is the same for the SRX5600 and SRX5800 devices, independent of the number of SPUs available because essentially each SPU needs to be able to access all the keys.

Multiple servers can share the same private key; however, one server can have only one private key. SSL decryption is disabled by default and can be enabled by using CLI configuration. Server keys and policy are configured exclusively through the CLI. Both plain and encrypted keys are supported.

NOTE: JUNOS Software does not encrypt SSL keys file for Release 9.3.

Displaying Keys and Servers

You can display all installed server keys and associated servers by using the following CLI command:

```
router> show security idp ssl-inspection key
```

Displays all server keys and IP addresses bound to those keys. The following example shows CLI output when the show security idp ssl-inspection key command is used:

```
Total SSL keys : 2
SSL server key and ip address :
   Key : key1, server : 1.1.1.1
   Key : key2, server : 2.2.2.2
   Key : key2, server : 2.2.2.3
```

To display IP addresses bound to a specific key, use the following CLI command:

```
router> show security idp ssl-inspection key <key-name>
```

The following is an example of the CLI output received when the show security idp ssl-inspection key < key-name > command is used:

```
Key : key1, server : 1.1.1.1
```

Adding Keys and Servers

To install a Privacy-Enhanced Mail (PEM) key use the following CLI command:

```
router> request security idp ssl-inspection key add <key-name> [file <file-path>] server
   <server-ip> [password <password-string>]
```

NOTE: When you are installing a key, you can password protect the key and also associate it to a server. You can also associate the key with a server at a later time by using the add server CLI command.

A server can be associated with only one key. To associate a server to the installed key, use the following CLI command:

```
router> request security idp ssl-inspection key add <key-name> server  <server-ip>
```

NOTE: Maximum key name length is 32 bytes, including the ending “\0”.

Maximum password string length is 32 bytes, including the ending “\0”.

Deleting Keys and Servers

To delete all keys and servers, use the following CLI command:

```
router> request security idp ssl-inspection key delete
```

All installed keys are deleted along with any associated servers.

To delete a specific key and all associated servers with that key, use the following CLI command:

```
router> request security idp ssl-inspection key delete  <key-name>
```

Deletes the specified key and all servers associated with that key.

To delete a single server, use the following CLI command:

```
router> request security idp ssl-inspection key delete  <key-name> server  <server-ip>
```

Deletes the specified server that is bound to the specified key.

Configuring SSL Inspection

SSL inspection is disabled by default. It is enabled if any configurations are found.

```
[edit security]
   idp {
   sensor-configuration {
      ssl-inspection {
         sessions <number>;
      }
   }
```

The sensor now inspects traffic for which it has a key/server pair.

NOTE: Maximum supported sessions per SPU: default value is 10,000 and range is 1 to 100,000. The session limit is per SPU and it is the same on the SRX5600 and SRX5800 devices.

Chapter 19

Performance and Capacity Tuning

If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Performance and Capacity Tuning for IDP Overview on page 607
- Configuring Session Capacity for IDP on page 608

Performance and Capacity Tuning for IDP Overview

Using the maximize-idp-sessions command, you can tune the device to assign a weight value to functions such as firewall and IDP. By default, the device assigns a greater weight value to firewall functions. By tuning the device accordingly, you can shift this weight value to increase IDP session capacity.

NOTE: You should not tune the device to increase IDP session capacity if you are not using an IDP policy.

The device ships with an implicit default session capacity setting. This default value gives added weight to firewall sessions. You can manually override the default by adding the maximize-idp-sessions setting to the system. When you do this, you can choose between several weight values: equal, firewall, and IDP. The following table displays the available session capacity weight values and the throughput changes that result when you configure the device to use these values.

Table 97: Session Capacity and Resulting Throughput

Weight Value	Firewall Capacity	IDP Capacity	Firewall Throughput	IDP Throughput
Default	1,000,000	256,000	10 Gpbs	2.4 Gpbs
Equal	1,000,000	1,000,000	8.5 Gpbs	2 Gpbs
Firewall	1,000,000	1,000,000	10 Gpbs	2.4 Gpbs

Table 97: Session Capacity and Resulting Throughput *(continued)*

Weight Value	Firewall Capacity	IDP Capacity	Firewall Throughput	IDP Throughput
IDP	1,000,000	1,000,000	5.5 Gpbs	1.4 Gpbs

NOTE: Throughput numbers are intended as a guide, actual performance may vary depending on security policy and other factors. Refer to datasheet for more information.

Configuring Session Capacity for IDP

Use the CLI to tune session capacity for IDP policies. You do this by adding the maximize-idp-sessions command and then adding the weight command to specify IDP sessions.

NOTE: The weight option is dependent upon the maximize-idp-sessions command being set.

■ CLI Configuration on page 608

CLI Configuration

1. If you have an active IDP policy, you can tune the device to increase IDP session capacity by entering following command:

   ```
   user@host# set security forwarding-process application-services
     maximize-idp-sessions weight idp
   ```

2. Commit your changes. Reboot the device for session capacity settings to take effect.

NOTE: If the device has maximize-idp-sessions weight enabled for IDP, and you do not have an IDP policy configured, a warning message appears when you commit your configuration. If you see this warning, you should remove your configured settings.

To turn the maximize-idp-sessions settings off, remove the maximize-idp-sessions configuration.

NOTE: You must reboot the device for any maximize-idp-sessions setting changes to take effect.

Chapter 20

IDP Logging

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding IDP Logging on page 609
- Configuring Log Suppression Attributes on page 610
- IDP Attack Log Information Usage on the Infranet Controller on page 611

Understanding IDP Logging

The basic JUNOS system logging continues to function after IDP is enabled. An IDP-enabled device continues to record events that occur because of routine operations, such as a user login into the configuration database. It records failure and error conditions, such as failure to access a configuration file. You can configure files to log system messages and also assign attributes, such as severity levels, to messages. In addition to the regular system log messages, IDP generates event logs for attacks. For information about monitoring events and managing system log files, see the *JUNOS Software Administration Guide*.

Before You Begin

For background information, read:

- Chapter on monitoring events and managing system log files in the *JUNOS Software Administration Guide*.
- Chapter on configuring packet capture in the *JUNOS Software Administration Guide*

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled. When you configure a rule for logging, the device creates a log entry for each event that matches that rule. You can use the CLI or J-Web to configure the policy rules to generate event logs. For more information about configuring IDP policies, see “Defining Rules for an IPS Rulebase” on page 515.

Because IDP event logs are generated during an attack, log generation happens in bursts, generating a much larger volume of messages during an attack. In comparison to other event messages, the message size is also much larger for attack generated messages. The log volume and message size are important concerns for log

management. To better manage the volume of log messages, IDP supports log suppression.

By configuring log suppression you can suppress multiple instances of the same log occurring from the same or similar sessions over the same period of time. Enabling log suppression ensures that minimal numbers of logs are generated for the same event or attack that occurs multiple times.

Related Topics

- Understanding IDP Policy Rules on page 506
- IDP Policies Overview on page 502

Configuring Log Suppression Attributes

Log suppression ensures that minimal numbers of logs are generated for the same event or attack that occurs multiple times. Log suppression is enabled by default. You can configure certain log suppression attributes to suppress logs according to your needs. When configuring log suppression, keep in mind that log suppression can negatively impact sensor performance if you set the reporting interval too high.

Before You Begin

1. For background information, read "Understanding IDP Logging" on page 609.
2. Establish basic connectivity. See the Getting Started Guide for your secure router.
3. Configure network interfaces. See the *JUNOS Software Interfaces and Routing Configuration Guide*.
4. Download the signature database. See "Updating the Signature Database Manually" on page 578

You can configure the following log suppression attributes:

- Include destination addresses while performing log suppression—You can choose to combine log records for events with a matching source address. By default, the IDP sensor does not consider destination when matching events for log suppression.
- Number of log occurrences after which log suppression begins—You can specify the number of instances that a specific event must occur before log suppression begins. By default, log suppression begins after the first occurrence.
- Maximum number of logs that log suppression can operate on—When log suppression is enabled, IDP must cache log records so that it can identify when multiple occurrences of the same event occur. You can specify how many log records are tracked simultaneously by IDP. By default, the maximum number of log records that IDP can operate on is 16384.
- Time after which suppressed logs are reported—When log suppression is enabled, IDP maintains a count of occurrences of the same event. After the specified

number of seconds have passed, IDP writes a single log entry containing the count of occurrences. By default, IDP reports suppressed logs after 5 seconds.

In the configuration instructions for this example, you configure log suppression to begin after the second occurrence of an event. You also specify that logs are reported after 20 seconds.

You can use either J-Web or the CLI configuration editor to configure log suppression.

This topic contains:

- CLI Configuration on page 611
- Related Topics on page 611

CLI Configuration

To configure log suppression attributes:

1. Specify the log number after which you want to start log suppression. In the following statement you specify that log suppression starts after the second instance of an event.

   ```
   user@host# set security idp sensor-configuration log suppression start-log 2
   ```

2. Specify the maximum time after which suppressed logs are reported. In the following statement you specify that IDP reports suppressed logs after 20 seconds.

   ```
   user@host#  set security idp sensor-configuration log suppression
     max-time-report 20
   ```

3. If you are finished configuring the router, commit the configuration.
4. From configuration mode in the CLI, enter the show security idp command to verify the configuration. For more information, see the *JUNOS Software CLI Reference*.

Related Topics

- Defining Rules for an IPS Rulebase on page 515
- IDP Policies Overview on page 502

IDP Attack Log Information Usage on the Infranet Controller

The infranet controller for the United Access Control (UAC) appliance can use IDP attack log information sent from the Juniper Networks device to apply access policies for traffic in which IDP logs indicate an attack has been detected. Using a secure channel of communication, these IDP logs are sent to the infranet controller directly and securely. IDP attack logs are sent to the infranet controller through the JUEP communication channel.

This section contains the following topics:

- IDP Log Message Filtering to the Infranet Controller on page 612
- Configuration Information on page 612

IDP Log Message Filtering to the Infranet Controller

When you configure the infranet controller to receive IDP log messages, you set certain filtering parameters on the infranet controller. Without this filtering, the infranet controller could potentially receive too many log messages. The filtering parameters could include the following:

- The infranet controller should only receive communications from IDP for sessions it has authenticated. See the *Unified Access Control Administration Guide* for details.
- You can create infranet controller filters for receiving IDP logs files based on the their severity. For example, if on the infranet controller the severity is set to high, then IDP only sends logs which have a severity greater than or equal to high. See the *Unified Access Control Administration Guide* for details.
- From the infranet controller, you can disable the receiving of all IDP logs. See the *Unified Access Control Administration Guide* for details.

Configuration Information

All the configuration for receiving and filtering IDP logs is done on the infranet controller. You should refer to the *Unified Access Control Administration Guide* for configuration information for receiving IDP logs and details on the JUEP communication channel.

Part 8
Unified Threat Management

- Unified Threat Management Overview on page 615
- Antispam Filtering on page 619
- Full File-based Antivirus Protection on page 639
- Express Antivirus Protection on page 685
- Content Filtering on page 707
- Web Filtering on page 723

Chapter 21

Unified Threat Management Overview

Unified Threat Management (UTM) is a term used to describe the consolidation of several security features into one device, protecting against multiple threat types. The advantage of UTM is streamlined installation and management of these multiple security capabilities.

The security features provided as part of the UTM solution are:

- Antispam — E-mail spam consists of unwanted e-mail messages, usually sent by commercial, malicious, or fraudulent entities. The antispam feature examines transmitted e-mail messages to identify e-mail spam. When the device detects an e-mail message deemed to be spam, it either drops the message or tags the message header or subject field with a preprogrammed string. The antispam feature uses a constantly updated spam block list (SBL). Sophos updates and maintains the IP-based SBL. The antispam feature is a separately licensed subscription service.
- Full File-Based Antivirus — A virus is executable code that infects or attaches itself to other executable code to reproduce itself. Some malicious viruses erase files or lock up systems. Other viruses merely infect files and overwhelm the target host or network with bogus data. The full file-based antivirus feature provides file-based scanning on specific Application Layer traffic checking for viruses against a virus signature database. It collects the received data packets until it has reconstructed the original application content, such as an e-mail file attachment, and then scans this content. Kaspersky Lab provides the internal scan engine. The full file-based antivirus scanning feature is a separately licensed subscription service.
- Express Antivirus — Express antivirus scanning is offered as a less CPU intensive alternative to the full file-based antivirus feature. The express antivirus feature, like the full antivirus feature, scans specific Application Layer traffic for viruses against a virus signature database. However, unlike full antivirus, express antivirus does not reconstruct the original application content. Rather, it just sends (streams) the received data packets, as is, to the scan engine. With express antivirus, the virus scanning is executed by a hardware pattern matching engine. This improves performance while scanning is occurring, but the level of security provided is lessened. Juniper Networks provides the scan engine. The express antivirus scanning feature is a separately licensed subscription service.
- Content Filtering — Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, protocol command, and embedded object type. Content filtering does not require a separate license.
- Web Filtering — Web filtering lets you manage Internet usage by preventing access to inappropriate Web content. There are three types of Web filtering

solutions. In the case of the integrated Web filtering solution, the decision-making for blocking or permitting Web access is done on the device after it identifies the category for a URL either from user-defined categories or from a category server (Websense provides the CPA Server). The integrated Web filtering feature is a separately licensed subscription service. The redirect Web filtering solution intercepts HTTP requests and forwards the server URL to an external URL filtering server provided by Websense to determine whether to block or permit the requested Web access. Redirect Web filtering does not require a separate license. With Juniper Local Web Filtering, the decision-making for blocking or permitting Web access is done on the device after it identifies the category for a URL from user-defined categories stored on the device. With Local filtering, there is no additional Juniper license or remote category server required.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Custom Objects on page 616
- Platform Support and Licensing on page 616

Custom Objects

Before you can configure most UTM features, you must first configure the custom objects for the feature in question. Custom objects are global parameters for UTM features. This means that configured custom objects can be applied to all UTM policies where applicable, rather than only to individual policies. Certain UTM features make use of certain custom objects. Refer to the feature documentation for information on the custom objects that must be configured for that feature.

- Anti-Virus“J-Web Configuration” on page 658
- Web Filtering “J-Web Configuration” on page 728
- Anti-Spam “J-Web Configuration” on page 628
- Content Filtering “J-Web Configuration” on page 711

Platform Support and Licensing

The majority of UTM features function as a subscription service requiring a license. You can redeem this license once you have purchased your subscription license SKUs. You redeem your license by entering your authorization code and chassis serial number into the Customer Service LMS interface. Once your entitlement is generated, you can use the CLI from your device to send a license update request to the LMS server. The LMS server then sends your subscription license directly to the device.

The command for applying your subscription license to the device is as follows:

```
user@host> request system license update
```

Table 98: UTM Feature Subscription Service License Requirements

UTM Feature	Requires License
Antispam	Yes
Antivirus: full	Yes
Antivirus: express	Yes
Content Filtering	No
Web Filtering: integrated	Yes
Web Filtering: redirect	No
Web Filtering: local	No

Chapter 22

Antispam Filtering

Spam consists of unwanted e-mail messages, usually sent by commercial, malicious, or fraudulent entities. The antispam feature examines transmitted e-mail messages to identify spam. When the device detects a message deemed to be spam, it blocks the e-mail message or tags the e-mail message header or subject with a preprogrammed string.

Antispam filtering allows you to use both a third-party server-based spam block list (SBL) and to optionally create your own local white lists (benign) and black lists (malicious) for filtering against e-mail messages. The antispam feature is not meant to replace your antispam server, but to complement it.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Using Server-Based Spam Filtering on page 619
- Configuring Server-Based Spam Filtering on page 620
- Using Local List Spam Filtering on page 626
- Configuring Local List Spam Filtering on page 627
- Understanding Spam Message Handling on page 636
- Blocking Detected Spam on page 636
- Tagging Detected Spam on page 636
- Verifying Antispam Configurations on page 637

Using Server-Based Spam Filtering

NOTE: Server-based spam filtering supports only IP-based spam block list black list lookup. Sophos updates and maintains the IP-based spam block list. Server-based antispam filtering is a separately licensed subscription service. When your antispam license key expires, you can continue to use locally defined black and white lists.

The server-based antispam feature requires Internet connectivity with the spam block list (SBL) server. Domain Name Service (DNS) is required to access the SBL server. The firewall performs SBL lookups through the DNS protocol. The lookups are against the IP address of the sender (or relaying agent) of the e-mail, adding the name of

the SBL server as the authoritative domain. The DNS server then forwards each request to the SBL server, which returns a DNS response to the device. The device then interprets the DNS response to determine if the e-mail sender is a spammer.

IP addresses that are included in the block lists are generally considered to be invalid addresses for mail servers or easily compromised addresses. Criteria for listing an IP address as a spammer on the SBL can include:

- Running an SMTP open relay service
- Running open proxy servers (of various kinds)
- Being a zombie host possibly compromised by a virus, worm, Trojan, or spyware
- Using a dynamic IP range
- Being a confirmed spam source with a known IP address

By default, the device first checks incoming e-mail against local white and black lists. If there are no local lists, or if the sender is not found on local lists, the device proceeds to query the SBL server over the Internet. When both server-based spam filtering and local list spam filtering are enabled, checks are done in the following order:

1. The local white list is checked. If there is a match, no further checking is done. If there is no match...
2. The local black list is checked. If there is a match, no further checking is done. If there is no match...
3. The SBL server list is checked.

NOTE: SBL server matching stops when the antispam license key is expired.

Configuring Server-Based Spam Filtering

You can use J-Web or the CLI to configure the device to use server-based spam filtering.

This topic contains:

- Configuration Overview on page 620
- J-Web Configuration on page 621
- J-Web Point and Click CLI Configuration on page 623
- CLI Configuration on page 625

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any) for the feature in question. Custom objects are global parameters for UTM features. This means that configured

custom objects can be applied to all UTM policies where applicable, rather than only to individual policies.

The CLI command for setting custom objects is :

```
user@host# set security utm custom-objects
```

2. Configure main feature parameters, called *feature profiles*.

   The CLI command for setting antispam feature profiles is:

   ```
   user@host# set security utm feature-profile anti-spam
   ```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for SMTP and attaching that policy to a profile are:

NOTE: At this time, the antispam feature is only supported for the SMTP protocol.

```
user@host# set security utm utm-policy <name>
user@host# set security utm utm-policy utmp1 anti-spam smtp-profile smtp1
```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

   ```
   user@host# set security policies
   user@host# set security policies from-zone trust to-zone untrust policy p1 then
     permit application-services utm-policy utmp1
   ```

J-Web Configuration

To configure server-based antispam using the J-Web Configuration editor:

1. Select **Configure > Security > UTM > Anti-Spam**.
2. In the Anti-Spam profiles configuration window, click **Add** to configure a profile for the SBL server.

   The profile configuration pop-up window appears. (To edit an existing item, select it and click the **Edit** button.)

3. In the Profile name box, enter a unique name for the antispam profile you are creating.
4. If you are using the default server, select **Yes** next to Default SBL server. If you are not using the default server, select **No**.

   The SBL server is predefined on the device. It ships knowing the name and address of the SBL server. If you do not select Yes, you are disabling server-based spam filtering. You would disable it if you are using only local lists or if you do not have a license for server-based spam filtering.

5. In the Custom tag string box, enter a custom string for identifying a message as spam.

   By default, the devices uses ***SPAM***.

6. In the antispam action list, select one of the following: **Tag subject**, **Block email**, or **Tag header**.

   Here you are selecting the action to be taken by the device when spam is detected.

7. Click **OK** to save your profile.

8. If the profile is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you configure a UTM policy for SMTP to which you attach the antispam profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.

2. In the UTM policy configuration window, click **Add** to configure a UTM policy.

3. In the policy configuration window, select the **Main** tab.

4. In the Policy name box, enter a unique name for the UTM policy you are creating.

5. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.

6. In the Session per client over limit list, select one of the following: **Log and permit** or **Block**.

   This is the action the device takes when the session per client limit for this UTM policy is exceeded.

7. Select the **Anti-Spam profiles** tab in the pop-up window.

8. From the SMTP profile list, select the antispam profile you are attaching to this UTM policy.

9. Click **OK** to save your policy.

10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.

2. In the Security Policy window, click **Add** to configure a security policy with UTM.

   This takes you to the policy configuration pop-up window. (To edit an existing item, select it and click the **Edit** button.)

3. In the Policy tab, enter a name in the Policy Name box.

4. Next to From Zone, select a zone from the list.

5. Next to To Zone, select a zone from the list.
6. Choose a **Source Address**.
7. Choose a **Destination Address**.
8. Choose an **Application**.

   Do this by selecting **junos-smtp** (for antispam) in the Application Sets box and clicking the —> button to move it to the Matched box.
9. Next to Policy Action, select one of the following: **Permit**, **Deny**, or **Reject**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

10. Select the **Application Services** tab in the pop-up window.
11. Next to UTM Policy, select the appropriate policy from the list. This attaches your UTM policy to the security policy.

NOTE: There are several fields in this pop-up window that are not described in this section. See the section on Security Policies for detailed information on configuring security policies and all the available fields.

12. Click **OK** to save your policy.
13. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

NOTE: You must activate your new policy to apply it.

J-Web Point and Click CLI Configuration

To configure server-based antispam using the J-Web Point and Click CLI:

1. Select the **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature Profile, click **Configure**.
5. Next to Anti spam, click **Configure**.
6. Next to sbl, select the **Yes** check box and click **Configure**. (Because you are configuring server-based spam, you do not select a local Address whitelist or an Address blacklist.)
7. Next to Profile, click **Add new entry**.

8. In the Name box, enter a unique name for the antispam profile you are creating.
9. In the Custom tag string box, enter a custom string for identifying a message as spam. By default, the devices uses ***SPAM***.
10. Select the **Yes** check box beside the Default SBL default server if you are using the default server. Otherwise, select the No check box.
11. In the Spam action list, select one of the following: **tag subject** (of e-mail), **block** (e-mail), or **tag header** (of e-mail). Here you are selecting the action to be taken by the device when spam is detected.
12. Click **OK**. See Figure 104 on page 624.

Figure 104: Antispam Server-Based Profile Configuration, Point and Click CLI Configuration

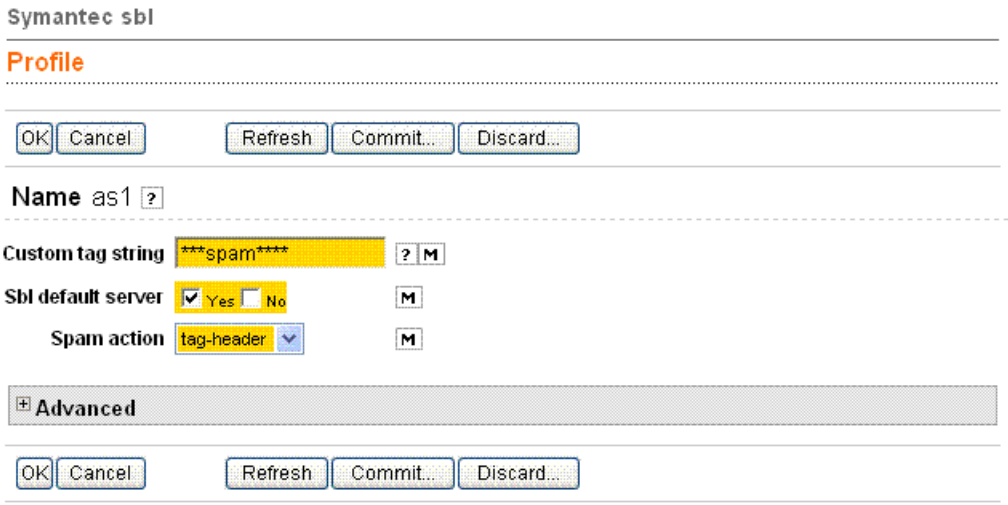

Next, you configure a UTM policy for SMTP to which you attach the antispam profile you have configured.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.
5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Anti spam, click **Configure**.
7. In the Smtp profile box, enter the name of the antispam profile you created earlier.
8. Click **OK**.
9. Click **OK** again to return to main UTM configuration page. Your UTM antispam policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Security, click **Configure**.
4. Next to Policy, select the **Yes** check box click **Edit**.
5. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

6. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
7. Click **OK**.

CLI Configuration

1. To configure server-based antispam using the CLI, first create a profile name.

```
user@host# set security utm feature-profile anti-spam sbl profile sblprofile1
```

2. Configure the default SBL server lookup as enabled or disabled. If you are using server-based spam filtering, you should enter sbl-default-server to enable the default SBL server. (The SBL server is predefined on the device. It ships knowing the name and address of the SBL server.) Entering no-sbl-default-server disables server-based spam filtering. You would disable it if you are using only local lists or if you do not have a license for server-based spam filtering.

```
user@host# set security utm feature-profile anti-spam sbl profile sblprofile1
  sbl-default-server
```

3. Configure the action to be taken by the device when spam is detected (block, tag-header, tag-subject).

```
user@host# set security utm feature-profile anti-spam sbl profile sblprofile1
  sbl-default-server spam-action block
```

4. Configure a custom string for identifying a message as spam.

```
user@host# set security utm feature-profile anti-spam sbl profile sblprofile1
  sbl-default-server custom-tag-string ***spam***
```

5. Configure a UTM policy for SMTP to which you attach the antispam feature profile you have configured.

```
user@host# set security utm utm-policy spampolicy1
```

6. Attach the spam feature profile to the UTM policy you have created.

```
user@host# set security utm utm-policy spampolicy1 anti-spam smtp-profile
  sblprofile1
```

7. Configure a security policy for UTM and attach the UTM policy you created to the security policy.

```
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 match source-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 match application junos-smtp
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 then permit application-services utm-policy spampolicy1
```

NOTE: A default antispam policy ships with the device. It is called junos-as-defaults. It contains the following configuration parameters:

```
anti-spam {
   sbl {
   profile junos-as-defaults {
      sbl-default-server;
      spam-action block;
   custom-tag-string "***SPAM***";
      }
   }
}
```

Using Local List Spam Filtering

When creating your own local white list and black list for spam filtering, you can filter against domain names, e-mail addresses, and/or IP addresses. Pattern matching works a bit differently depending upon the type of matching in question. For example, pattern matching for domain names uses a longest suffix match algorithm. If the sender e-mail address has a domain name of aaa.bbb.ccc, the device tries to match "aaa.bbb.ccc" in the list. If no match is found, it tries to match "bbb.ccc", and then "ccc". IP address matching, however, does not allow for partial matches.

Spam filtering uses local lists for matching in the following manner:

1. Sender IP: The sender IP is checked against the local white list, then the local black list, and then the SBL IP-based server (if enabled).
2. Sender Domain: The domain name is checked against the local white list and then against the local black list.
3. Sender E-mail Address: The sender e-mail address is checked against the local white list and then against the local black list.

By default, the device first checks incoming e-mail against local white and black lists. If the sender is not found on either list, the device proceeds to query the SBL server over the Internet. When both server-based spam filtering and local list spam filtering are enabled, checks are done in the following order:

1. The local white list is checked. If there is a match, no further checking is done. If there is no match...
2. The local black list is checked. If there is a match, no further checking is done. If there is no match...
3. The SBL server list is checked.

NOTE: Local black and white list matching continues after the antispam license key is expired.

Configuring Local List Spam Filtering

You can use J-Web or the CLI to configure the device to use local list spam filtering.

This topic contains:

- Configuration Overview on page 627
- J-Web Configuration on page 628
- J-Web Point and Click CLI Configuration on page 631
- CLI Configuration on page 634

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any), for the feature in question. Custom objects are global parameters for UTM features. This means that configured custom objects apply to all UTM policies where applicable, rather than only to individual policies.

   The CLI command for setting antispam url custom objects is :

   ```
   user@host# set security utm custom-objects url-pattern
   ```

2. Configure main feature parameters, called feature profiles.

   The CLI command for setting antispam feature profiles is :

   ```
   user@host# set security utm feature-profile anti-spam
   ```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for SMTP and attaching that policy to a profile are:

```
user@host# set security utm utm-policy <name>
user@host# set security utm utm-policy utmp1 anti-spam smtp-profile smtp1
```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

```
user@host# set security policies
user@host# set security policies from-zone trust to-zone untrust policy p1 then
    permit application-services utm-policy utmp1
```

J-Web Configuration

To configure local list antispam using the J-Web Configuration editor, you must first create your local white and black list custom objects.

Configure a URL Pattern List Custom Object as follows:

1. Select **Configure > Security > UTM > Custom Objects**.
2. In the UTM custom objects configuration window, select the **URL Pattern List** tab.
3. Click **Add** to create URL pattern lists.
4. Next to URL Pattern Name, enter a unique name for the list you are creating.

   If you are creating a white list, it is helpful to indicate this in the list name. The same applies to a black list. The name you enter here becomes available in the Address Whitelist and Address Blacklist fields when you are configuring your antispam profiles.
5. Next to URL Pattern Value, enter the URL pattern for white list or black list antispam filtering.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

6. Click **Add** to add your URL pattern to the Values list box.

   The list can contain up to 8192 items.
7. Click **OK**.
8. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

NOTE: Because you use URL Pattern Lists to create Custom URL Category Lists, you must configure URL Pattern List Custom Objects before you configure a Custom URL Category List. URL Pattern List Custom Objects are described in the previous section.

Configure a Custom URL Category List Custom Object as follows:

1. Select **Configure > Security > UTM > Custom Objects**.
2. In the UTM custom objects configuration window, select the **URL Category List** tab.
3. Click **Add** to create URL category lists.
4. Next to URL Category Name, enter a unique name for the list you are creating. This name appears in the Address Whitelist list when you configure antispam global options.
5. In the Available Values box, select a URL Pattern List name from the list for bypassing scanning and click the right arrow button – **>** to move it to the Selected Values box.

   Continue to add URLs or IP addresses in this manner.
6. Click **OK** to save the selected values as part of the custom URL list you have created.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Now that you have created your white list and black list custom objects, you can configure antispam filtering to use those lists.

1. Select **Configure > Security > UTM > Global options**.
2. Select the **Anti-Spam** tab in the right pane.
3. Under Anti-Spam , select an **Address Whitelist** and/or an **Address Blacklist** from the list for local lists for spam filtering.

   These lists are configured as Custom Objects and are made available for selection here.
4. Click **OK**.
5. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.
6. Select the **Anti-Spam** tab in the left pane under Security.
7. Click **Add** to configure an anti-spam profile.

   The profile configuration pop-up window appears.
8. In the Profile name box, enter a unique name for the antispam profile you are creating.

9. If you are using the default server, select **Yes** beside Default SBL server. If you are not using the default server, select **No**.

   If you select No, you are disabling server-based spam filtering. You would disable it if you are using only local lists or if you do not have a license for server-based spam filtering. If you are using the default SBL server, refer to "Configuring Server-Based Spam Filtering" on page 620.

10. In the Custom tag string box, enter a custom string for identifying a message as spam.

    By default, the devices uses ***SPAM***.

11. In the Actions list, select one of the following: **Tag subject**, **Block email**, or **Tag header**.

    Here you are selecting the action to be taken by the device when spam is detected.

12. Click **OK** to save your profile.

13. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you configure a UTM policy for SMTP to which you attach the antispam profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.

2. In the UTM policy configuration window, click **Add** to configure a UTM policy.

   The policy configuration pop-up window appears.

3. Select the **Main** tab in pop-up window.

4. In the Policy name box, enter a unique name for the UTM policy you are creating.

5. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.

6. In the Session per client over limit list, select one of the following: **Log and permit**, **Block**.

   This is the action the device takes when the session per client limit for this UTM policy is exceeded.

7. Select the **Anti-Spam profiles** tab in the pop-up window.

8. From the SMTP profile list, select the antispam profile you are attaching to this UTM policy.

9. Click **OK** to save your policy.

10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.
2. In the Security Policy window, click **Add** to configure a security policy with UTM.

   The policy configuration pop-up window appears.
3. In the Policy tab, enter a name in the Policy Name box.
4. Next to From Zone, select a zone from the list.
5. Next to To Zone, select a zone from the list.
6. Choose a **Source Address**.
7. Choose a **Destination Address**.
8. Choose an **Application**.

   Do this by selecting **junos-smtp** (for antispam) in the Application Sets box and clicking the **—>** button to move it to the Matched box.
9. Next to Policy Action, select one of the following: **Permit**, **Deny**, or **Reject**.

NOTE: When you select Permit for policy action, several additional fields become available in the Applications Services tab, including UTM Policy.

10. Select the **Application Services** tab in the pop-up window.
11. Next to UTM Policy, select the appropriate policy from the list.

    This attaches your UTM policy to the security policy.

NOTE: There are several fields in this pop-up window that are not described in this section. See the section on Security Policies for detailed information on configuring security policies and all the available fields.

12. Click **OK** to save your policy.
13. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

NOTE: You must activate your new policy to apply it.

J-Web Point and Click CLI Configuration

To configure local list antispam using the J-Web Point and Click CLI, you must first create your local white and black list custom objects.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to URL pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the URL pattern for white list or black list antispam filtering.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

9. Click **OK**. See Figure 105 on page 632.

Figure 105: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

NOTE: Continue to create as many Whitelist and Blacklist entries as necessary in your lists. You cannot use URL patterns directly. You must configure Custom URL category lists to which you apply your URL patterns.

Configure a Custom URL Category List Custom Object as follows: (See "URL White List" on page 650 for overview information.)

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.

3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Custom url category, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the name of the Url pattern list you created for bypassing scanning.
9. Click **OK**. See Figure 106 on page 633.

Figure 106: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

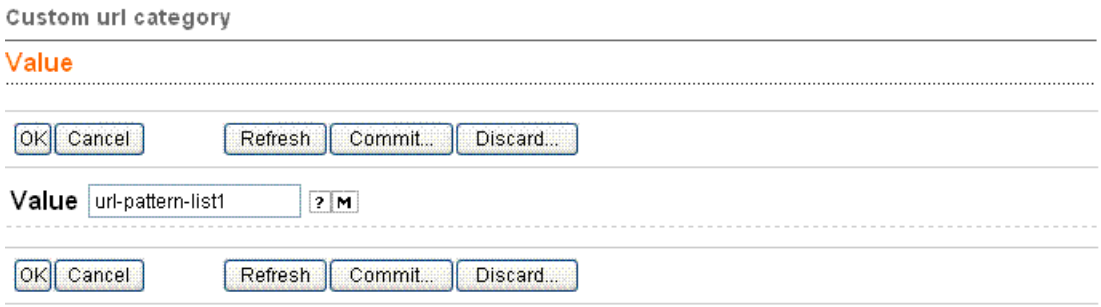

Now that you have created your white list and black list custom objects, you can configure antispam filtering to use those lists:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature profile, click **Configure**.
5. Next to Anti spam, click **Configure**.
6. Next to Address blacklist, enter the name of the local list you created earlier as a custom object.
7. Next to Address whitelist, enter the name of the local list you created earlier as a custom object.
8. Next to sbl, select the check box and click **Edit**.
9. Next to Profile, click **Add new entry**.
10. In the Name box, enter a unique name for the antispam profile you are creating.
11. Next to Custom tag string, enter a custom string for identifying a message as spam. By default, the devices uses ***SPAM***.
12. Next to Sbl default server, if you are using the default server, select the **Yes** check box. If you are not using the default server, select the No check box. You would

select No if you are using only local lists or if you do not have a license for server-based spam filtering.

13. In the Spam action list, select one of the following: **tag subject** (of e-mail), **block** (e-mail), or **tag header** (of e-mail). Here you are selecting the action to be taken by the device when spam is detected.
14. Click **OK**.

Next, you configure a UTM policy for SMTP to which you attach the antispam profile you have configured.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.
5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Anti spam, click **Configure**.
7. In the Smtp profile box, enter the name of the antispam profile you created earlier.
8. Click **OK**.
9. Click **OK** again to return to main UTM configuration page. Your UTM antispam policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policy, select the **Yes** check box click **Edit**.
4. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

5. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
6. Click **OK**.

CLI Configuration

1. Configure the local list spam blocking by first creating your global local spam lists (black list and white list).

```
user@host# set security utm custom-objects url-pattern as-black value
  [150.61.8.134 150.61.8.206 67.216.192.16 12.154.55.44 194.39.243.98
  12.154.55.42 12.154.55.41 12.154.55.43]
user@host# set security utm custom-objects url-pattern as-white value [150.1.2.3
  150.1.2.4]
```

2. Configure a custom URL category list custom object using the URL pattern list you created as follows:

```
user@host# set security utm custom-objects custom-url-category whilecusturl1
  value as-white
```

3. Configure the local list antispam feature-profile by first attaching your custom-object black list or white list or both. When both the white list and the black list are in use, the white list is checked first. If there is no match, then the black list is checked.

```
user@host# set security utm feature-profile anti-spam address-whitelist
  whiltecusturl1
```

4. Configure a profile for your local list spam blocking. Although you are not using the sbl for local list spam blocking, you configure your profile from within that command similar to the server-based spam blocking procedure.

```
user@host# set security utm feature-profile anti-spam sbl profile localprofile1
```

5. Configure the action to be taken by the device when spam is detected (block, tag-header, tag-subject).

```
user@host# set security utm feature-profile anti-spam sbl profile localprofile1
  spam-action block
```

6. Configure a custom string for identifying a message as spam.

```
user@host# set security utm feature-profile anti-spam sbl profile localprofile1
  custom-tag-string ***spam***
```

7. Configure a UTM policy for SMTP to which you attach the spam feature profile you have configured.

```
user@host# set security utm utm-policy spampolicy2
```

8. Attach the spam feature profile to the UTM policy you have created.

```
user@host# set security utm utm-policy spampolicy2 anti-spam smtp-profile
  localprofile1
```

9. Configure a security policy for UTM and attach the UTM policy you created to the security policy.

```
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 match source-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 match application junos-smtp
```

```
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 then permit application-services utm-policy spampolicy2
```

Understanding Spam Message Handling

There are two possible actions the device can take when spam is detected. It can perform a drop action or a tag action.

This topic contains:

- Blocking Detected Spam on page 636
- Tagging Detected Spam on page 636

Blocking Detected Spam

The device can block and drop detected spam at either the connection level or the e-mail level:

- Blocking spam at the connection level

  When the SMTP sender is identified as a spam sender based on its IP address , the SMTP connection is rejected and dropped. An error message with a proper error code from the firewall is sent out on behalf of the SMTP server. An example of such an error message is:

  554 Transaction failed due to anti spam setting

- Blocking spam at the e-mail level

  When a particular e-mail sender is identified as spam sender based on its sender address, the e-mail is rejected and dropped. An error message with a proper error code from the firewall is sent back to the sender on behalf of the server. An example of such an error message is:

  550 Requested action not taken: mailbox unavailable

Tagging Detected Spam

The device can allow and tag the e-mail if the message sender is detected as a spammer. This tagging can occur at the connection level so that all the e-mails for the connection in question are tagged. Otherwise, you can tag only an individual e-mail. Two tagging methods are supported:

- Tag the subject: A user-defined string is added at the beginning of the subject of the e-mail.
- Tag the header: A user-defined string is added to the e-mail header.

Verifying Antispam Configurations

J-Web monitoring tools consist of the options that appear when you select Monitor in the task bar. The Monitor options display diagnostic information about the feature in question. Alternatively, you can enter show commands from the CLI to display the same information. CLI show commands display the current configuration and information about the feature in question.

This topic contains:

- Using J-Web for Antispam Monitoring on page 637
- Using the CLI for Antispam Monitoring on page 637

Using J-Web for Antispam Monitoring

View antispam statistics using J-Web as follows:

1. Select **Monitor > Security > UTM > Anti Spam**.

   The following information becomes viewable in the right pane.

   ```
   user@host > show security utm anti-spam status

   SBL Whitelist Server:
   SBL Blacklist Server:
   server.juniper.net
   DNS Server:
   Primary : 1.2.3.4, Src Interface: ge-0/0/0
   Secondary: 2.3.4.5, Src Interface: ge-0/0/1
   Ternary : 0.0.0.0, Src Interface: fe-0/0/2

   Total connections: #
   Denied connections: #
   Total greetings: #
   Denied greetings: #
   Total e-mail scanned: #
   Spam total: #
   Spam tagged: #
   Spam dropped: #
   DNS errors: #
   Timeout errors: #
   Return errors: #
   Invalid parameter errors: #
   Statistics start time:
   Statistics for the last 10 days.
   ```

2. You can click the **Clear Antispam statistics** button to clear all current viewable statistics and begin collecting new statistics.

Using the CLI for Antispam Monitoring

Use the following commands to view antispam status and statistics

The antispam show status command displays the following information:

```
user@host > show security utm anti-spam status

SBL Whitelist Server:
SBL Blacklist Server:
server.juniper.net
DNS Server:
Primary : 1.2.3.4, Src Interface: ge-0/0/0
Secondary: 2.3.4.5, Src Interface: ge-0/0/1
Ternary : 0.0.0.0, Src Interface: fe-0/0/2
```

The antispam show statistics command display the following information:

```
user@host > show security utm anti-spam statistics

Total connections: #
Denied connections: #
Total greetings: #
Denied greetings: #
Total e-mail scanned: #
Spam total: #
Spam tagged: #
Spam dropped: #
DNS errors: #
Timeout errors: #
Return errors: #
Invalid parameter errors: #
```

Chapter 23

Full File-based Antivirus Protection

A virus is executable code that infects or attaches itself to other executable code in order to reproduce itself. Some malicious viruses erase files or lock up systems, while other viruses merely infect files and can overwhelm the target host or network with bogus data. The full file-based antivirus feature provides file-based scanning on specific Application Layer traffic checking for viruses against a virus signature database. It collects the received data packets until it has reconstructed the original application content, such as an e-mail file attachment, and then scans this content.

The full file-based antivirus scanning feature is a separately licensed subscription service. Kaspersky Lab provides the scan engine for full file-based antivirus. When your antivirus license key expires, you can continue to use locally stored antivirus signatures without any updates. But in that case, if the local database is deleted, antivirus scanning is disabled.

NOTE: For information on the express antivirus feature, which provides better performance but lower security, refer to that section in the manual. Note that if you switch from full file-based antivirus protection to express antivirus protection, you must reboot the device in order for express antivirus to begin working.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding the Internal Scan Engine and Scan Options on page 640
- Understanding Updating the Antivirus Scanner Pattern Database on page 645
- Understanding the Scanning of Application Protocols on page 647
- HTTP Scanning on page 648
- FTP Scanning on page 651
- SMTP Scanning on page 651
- POP3 Scanning on page 653
- IMAP Scanning on page 655
- Configuring Full Antivirus Protection on page 657
- Understanding Virus-Detected Notification Options on page 676

- Understanding Scan Result Handing on page 679
- Verifying Antivirus Configurations on page 679

Understanding the Internal Scan Engine and Scan Options

The full file-based antivirus module is the software subsystem on the gateway device that scans specific Application Layer traffic to protect users from virus attacks and to prevent viruses from spreading. The antivirus software subsystem consists of a virus signature database, an application proxy, the scan manager, and the scan engine.

Kaspersky Lab provides the scan engine and it works in the following manner. A client establishes a TCP connection with a server and then starts a transaction. If the application protocol in question is marked for antivirus scanning, the traffic is forwarded to an application proxy for parsing. When the scan request is sent, the scan engine scans the data by querying a virus pattern database. The scan manager monitors antivirus scanning sessions, checking the properties of the data content against the existing antivirus settings. After scanning has occurred, the result is then handled by the scan manager.

The following topics provide information on the scan engine and various configurable scanning options:

This chapter contains the following topics:

- Scan Type Support on page 640
- Scan Mode Support on page 641
- Enable and Disable Scanning on a Per-Protocol Basis on page 641
- File Extension Scanning on page 641
- Intelligent Prescreening on page 642
- Signature Database Support on page 642
- Content Size Limits on page 642
- Decompression Layer Limit on page 643
- Scanning Timeout on page 643
- Global, Profile-Based, and Policy-Based Scan Settings on page 644
- Scan Session Throttling on page 645

Scan Type Support

The Kaspersky Lab scan engine supports two types of file scanning.

- Regular file scan — In this case, the input object is a regular file. The engine matches the input content with all possible signatures.
- Scan script files — In this case, the input object is a script file. It can be JavaScript, VBScript, mIRC script, bat scripts (DOS bat files), and other text scripts. The engine matches the input content only with signatures for script files. Script scanning is only applicable for HTML content over the HTTP protocol. There are two criteria for this scan type. First, the content-type field of this HTML document

must be text or HTML. Second, there is no content encoding in the HTTP header. If those two criteria are met, an HTML parser is used to parse the HTML document for scripts.

Scan Mode Support

The Kaspersky Lab scan engine supports two modes of scanning.

- scan-all

  This option tells the scan engine to scan all the data it receives.

- scan-by-extension

  This option bases all scanning decisions on the file extensions found in the traffic in question. See "File Extension Scanning" on page 641.

Enable and Disable Scanning on a Per-Protocol Basis

You can turn antivirus scanning on and off on a per protocol basis. If scanning for a protocol is disabled in an antivirus profile, there is no application intelligence for this protocol. Therefore, in most cases, traffic using this protocol is not scanned. But if the protocol in question is based on another protocol for which scanning is enabled in an antivirus profile, then the traffic is scanned as that enabled protocol.

File Extension Scanning

When scanning content, you can use a file extension list to define a set of file extensions that are used in file extension scan mode (scan-by-extension). The antivirus module can then scan files with extensions on the scan-extension list. If an extension is not defined in an extension list, the file with that extension is not scanned in scan-by-extension mode. If there is no extension present, the file in question is scanned.

There are some requirements to note when using a file extension list to scan content.

- File extension entries are case-insensitive.
- The maximum length of the file extension list name is 29 bytes.
- The maximum length of each file extension entry is 15 bytes.
- The maximum entry number in a file extension list is 255.

```
edit security utm {
    custom-objects {
        filename-extension { ; set of list
            name <extension-list-name>; #mandatory
                value <windows-extension-string>;
            }
        }
    }

edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
```

```
scan-options {
   scan-extension <ext-list>
   }
}
```

Intelligent Prescreening

By default, intelligent prescreening is enabled to improve antivirus scanning performance. The antivirus module generally begins to scan data after the gateway device has received all the packets of a file. Intelligent prescreening tells the antivirus module to begin scanning a file much earlier. In this case, the scan engine uses the first packet or the first several packets to determine if a file could possibly contain malicious code. The scan engine does a quick check on these first packets and if it finds that it is unlikely that the file is infected, it then decides that it is safe to bypass the normal scanning procedure.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for MIME encoded traffic, mail protocols (SMTP, POP3, IMAP) and HTTP POST.

Signature Database Support

The full file-based antivirus protection signature database is called the Juniper Full antivirus database (downloaded by the pattern-update command). This database is different from the database used by express antivirus. It detects all destructive malicious code, including viruses (polymorphic and other advanced virus types), worms, Trojans, and malware.

Content Size Limits

Due to resource constraints, there is a default, device-dependent limit on maximum content size for the database. The content size value is configurable. There is also a lower and upper limit for maximum content size. (This range is device dependent and is not configurable.)

The content size check occurs before the scan request is sent. The exact timing of this is protocol dependent. If the protocol header contains an accurate content length field, the content size check takes place when the content length field is extracted during header parsing. The content size usually refers to file size. If there is no content length field, the size is checked while the antivirus module is receiving packets. The content size, in this case, refers to accumulated TCP payload size.

NOTE: This setting can be used in all protocols.

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
   scan-options {
      content-size-limit <KB>;
   }
```

```
}
```

Decompression Layer Limit

The decompression layer limit specifies how many layers of nested compressed files and files with internal extractable objects, such as archive files (tar), MS Word and PowerPoint files, the internal antivirus scanner can decompress before it executes the virus scan. For example, if a message contains a compressed .zip file that contains another compressed .zip file, there are two compression layers. Decompressing both files requires a decompress layer setting of 2.

It is worth noting that during the transfer of data, some protocols use content encoding. The antivirus scan engine must decode this layer, which is considered a decompression level, before it scans for viruses.

There are three kinds of compressed data:

- compressed file (zip, rar, gzip)
- encoded data (MIME)
- packaged data (OLE, .CAP, .MSI, .TAR, .EML)

A Decompression Layer could be a layer of a zipped file or an embedded object in packaged data. The antivirus engine scans each layer before unpacking the next layer, until it either reaches the user-configured decompress limit, reaches the device decompress layer limit, finds a virus or other malware, or decompresses the data completely, whichever comes first.

As the virus signature database becomes larger and the scan algorithms become more sophisticated, the scan engine has the ability to look deeper into the data for embedded malware. As a result, it can uncover more layers of compressed data. The Juniper device's level of security is limited by decompress limit, which is based on the memory allocated to the security service. If a virus is not found within the decompress limit, the user has an option to either pass or drop the data.

NOTE: This setting can be used in all protocols.

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
    scan-options {
        decompress-layer-limit <num>
    }
}
```

Scanning Timeout

The scanning timeout value includes the time frame from when the scan request is generated to when the scan result is returned by the scan engine. The time range can be 1 to 1800 seconds. By default, it is 180 seconds.

NOTE: This timeout parameter is used by all supported protocols. Each protocol can have a different timeout value.

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
    scan-options {
        timeout-value <sec> {
        }
    }
}
```

Global, Profile-Based, and Policy-Based Scan Settings

The antivirus module allows you to configure scanning options on a global level, on a UTM profile level, or on a firewall policy level. Each configuration level has the following implications:

- Global antivirus settings — Settings are applied to all antivirus sessions.
- Profile-based settings — Antivirus settings are different for different protocols within the same policy.
- Policy-based settings — Antivirus settings are different for different policies. Policy-based antivirus settings are applied to all scan-specified traffic defined in a firewall policy.

The majority of antivirus settings are configured within an antivirus profile, bound to specified protocols, and used by designated policies. These UTM policies are then applied to the traffic according to firewall policies. If a firewall policy with an antivirus setting matches the properties of a traffic flow, the antivirus setting is applied to the traffic session. Therefore, you can apply different antivirus settings for different protocols and for different traffic sessions.

Global settings are general overall configurations for the antivirus module or settings that are not specific for profiles.

Global antivirus setting CLI example:

```
edit security utm feature-profile anti-virus kaspersky-lab-engine
```

Profile-based antivirus setting CLI example:

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name>
```

Policy-based antivirus setting CLI example:

```
edit security utm utm-policy <name>
```

The following is an example of different antivirus settings applied to different protocols configured as profiles within a designated UTM policy:

```
edit security utm
    utm-policy <name> {
        anti-virus {
            http-profile <av profile>
```

```
        ftp {
            upload-profile <av profile>
            download-profile <av profile>
        }
        smtp-profile <av profile>
        pop3-profile <av profile>
        imap-profile <av profile>
    }
}
```

Scan Session Throttling

In an attempt to consume all available resources and hinder the ability of the scan engine to scan other traffic, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, a session throttle is imposed for antivirus resources , thereby restricting the amount of traffic a single source can consume at one time. The limit is an integer with 100 as the default setting. This integer refers to the maximum allowed sessions from a single source. You may change this default limit, but understand that if this limit is set high, that is comparable to no limit.

Over-limit is a fallback setting for the connection-per-client limit. The default behavior of over-limit is to block sessions. This is a per-policy setting. You can specify different settings for different UTM policies.

```
edit security utm utm-policy <name>
    traffic-options {
        sessions-per-client {
        limit <number>;
            over-limit { log-and-permit | block}
        }
    }
```

Understanding Updating the Antivirus Scanner Pattern Database

When Kaspersky Lab updates the signatures in its pattern database, the security device downloads these updates so that the antivirus scanner is using the latest, most up-to-date signatures when scanning traffic. The security device can perform these updates automatically (the default), or you can perform pattern update downloads manually.

There are several prerequisites that must be met in order to perform a successful pattern database update:

- You must have a valid antivirus scanner license.
- You must have network connectivity and access to the pattern database server.
- Your DNS settings and port settings (port 80) must be correct.

Updating Antivirus Patterns

Update the patterns for the antivirus signature database as follows:

1. On the security device, specify the URL address of the pattern-update server.

   By default, the Juniper-Kaspersky URL for full antivirus is http://update.juniper-updates.net/AV/SRX210. "SRX210" in the URL is the platform name. This part of the URL is different and platform specific for each platform. (Other than the platform name, you should not change this URL unless you are experiencing problems with it and have called for support.)

2. After the security device downloads the server-initialization file, the device checks that the pattern file is valid. The device then parses the file to obtain information about it, including the file version, size, and location of the pattern file server.

   If the pattern file on the security device is out-of-date (or nonexistent because this is the first time you are loading it), and, if the antivirus pattern-update service subscription is still valid, the device automatically retrieves an updated pattern file from the pattern file server.

NOTE: Once your subscription expires, you have a 30 day grace period during which you can continue to update the antivirus pattern file. Once that grace period expires, the update server no longer permits antivirus pattern file updates.

Updates to the pattern file are added as new viruses are discovered. You can configure the security device to regularly update the pattern file automatically, or you can update the file manually.

Example: Automatic Update

In this example, you configure the security device to update the pattern file automatically every 120 minutes. (The default antivirus pattern-update interval is 60 minutes.)

- J-Web

  1. Select **Configure > UTM > Anti-Virus**.
  2. Next to Interval, in the Kaspersky Lab Engine section, enter **120** in the box.
  3. Click **OK**.

- CLI

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
   pattern-update interval 120
```

Example: Manual Update

In this example, you use the request command to update the pattern file manually.

- J-Web

  At this time, you cannot perform manual database updates using J-Web.

- CLI

  user@host> request security utm anti-virus kaspersky-lab-engine **pattern-update**

  You can also manually delete the pattern database (pattern-delete) and reload the pattern database (pattern-reload) using the request command.

Database Download Process

The database pattern server is accessible through HTTP or HTTPS. By default, the antivirus module checks for database updates automatically every 60 minutes. You can change this interval and you can trigger updates manually, as well. The number of files that are downloaded during an update and the duration of the download process can vary.

A local copy of the pattern database is saved in persistent data storage (that is, the flash disk). If the device is rebooted, the local copy remains available for the antivirus scan engine to use during the antivirus scan engine initialization time, without the need for network access to the pattern database server.

The following is an example of the CLI for configuring the database update feature:

```
utm {
    feature-profile {
        anti-virus {
            type
                kaspersky-lab-engine {
                    pattern-update
                        url < url>
                    interval <min>
                }
            }
        }
    }
```

NOTE: If the auto-update fails, the updater automatically retries to update three more times. If the database download continues to fail, the updater stops trying and waits for the next periodic update before trying again.

Understanding the Scanning of Application Protocols

The internal antivirus scan engine supports scanning for specific Application Layer transactions allowing you to select the content (HTTP, FTP, SMTP, POP3, or IMAP traffic) to scan. For each content type that you are scanning, you have different configuration options.

Profile-based settings, including enable/disable, scan-mode, and scan result handling settings, may not be applicable to all supported protocols. The following table lists profile-based settings and their protocol support.

Table 99: Supported Profile-based Settings By Protocol

Profile Setting	Protocol Support
"Enable and Disable Scanning on a Per-Protocol Basis" on page 641	All protocols support this feature
"Scan Mode Support" on page 641	All protocols support this feature
"File Extension Scanning" on page 641	All protocols support this feature
"Content Size Limits" on page 642	All protocols support this feature
"Decompression Layer Limit" on page 643	All protocols support this feature
"Scanning Timeout" on page 643	All protocols support this feature
"HTTP Trickling" on page 649	HTTP only
"Fallback Options" on page 677	All protocols support this feature
"Protocol-Only Notifications" on page 676	All protocols support this feature
"E-mail Notification" on page 676	SMTP, POP3, and IMAP only
"Custom Message Notification" on page 677	All protocols support this feature

HTTP Scanning

If antivirus scanning is enabled for HTTP (Hypertext Transfer Protocol) traffic in a content security profile, TCP traffic to defined HTTP service ports (generally port 80) is monitored. For HTTP traffic, the security device scans both HTTP responses and requests (get, post, and put commands).

NOTE: For HTTP antivirus scanning, both HTTP 1.0 and 1.1 are supported. If the protocol version is HTTP 0.x , the antivirus scanner attempts to scan the traffic. Unknown protocols are bypassed. For example, some application protocols use HTTP as the transport but do not comply with HTTP 1.0 or 1.1. These are considered unknown protocols and are not scanned.

This topic contains:

- HTTP Scanning: Process Overview on page 649
- HTTP Trickling on page 649
- MIME White List on page 649
- URL White List on page 650
- Script-only Scanning on page 650

HTTP Scanning: Process Overview

This is a general description of how HTTP traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. An HTTP client sends an HTTP request to a Web server or a Web server responds to an HTTP request.
2. The security device intercepts the request and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the device follows one of two courses:
   - If there is no virus, the device forwards the request to the Web server.
   - If there is a virus, the device drops the request and sends an HTTP message reporting the infection to the client.

HTTP Trickling

HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning. On some slow link transferring, a large file could timeout if too much time is taken for the antivirus scanner to scan a complex file.

HTTP trickling is the forwarding of specified amounts of unscanned HTTP traffic to the requesting HTTP client to prevent the browser window from timing out while the scan manager examines downloaded HTTP files. (The security device forwards small amounts of data in advance of transferring an entire scanned file.)

HTTP Trickling is time-based and there is only one parameter to configure for this feature. That is the **Time-out Interval**. By default, trickling is disabled.

NOTE: The timeout based trickling is packet driven. This means, if no packet is received within a certain time frame, HTTP trickling is discontinued. This setting is only supported for HTTP connections.

The following is a trickling configuration example:

```
edit security utm feature-profile anti-virus kaspersky-lab-engine {
    profile <name> {
        trickling timeout <sec>;
    }
}
```

MIME White List

The gateway device uses MIME (Multipurpose Internet Mail Extension) types to decide which traffic may bypass antivirus scanning. The MIME white list defines a list of MIME types and can contain one or many MIME entries.

A MIME entry is case-insensitive. An empty MIME is an invalid entry and should never appear in the MIME list. If the MIME entry ends with a / character, prefix matching takes place. Otherwise, exact matching occurs.

There are two types of MIME lists used to configure MIME type antivirus scan bypassing:

- mime-whitelist list — This is the comprehensive list for those MIME types that can bypass antivirus scanning.
- exception list — The exception list is a list for excluding some MIME types from the mime-whitelist list. This list is a subset of MIME types found in the mime-whitelist.

  For example, if the mime-whitelist includes the entry,video/ and the exception list includes the entry video/x-shockwave-flash, by using these two lists, you can bypass objects with “video/” MIME type but not bypass “video/x-shockwave-flash” MIME type.

  You should note that there are limits for mime-whitelist entries as follows:

  - The maximum number of MIME items in a MIME list is 50.
  - The maximum length of each MIME entry is restricted to 40 bytes.
  - The maximum length of a MIME list name string is restricted to 40 bytes.

URL White List

A URL white list is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for scanning. Because antivirus scanning is CPU and memory intensive action, if there are URLs or IP addresses that you are sure do not require scanning, you may want to create this custom list add them to it.

The following is a white list configuration example:

```
set security utm custom-objects {
   custom-url-category { ; set of list
      name <url-category-name>; #mandatory
      value <url-pattern-name>;
   }
}
```

Script-only Scanning

With script-only scanning, the input object is a script file. It can be JavaScript, VBScript, mIRC script, bat scripts (DOS bat files) and other text scripts. The engine matches the input content only with signatures for script files. Script scanning is applicable only for HTML content over the HTTP protocol. There are two criteria for this scan-type. First, the content-type field of this HTML document must be text or HTML. Second, there is no content encoding in the HTTP header. If those two criteria are met, an HTML parser is used to parse the HTML document.

FTP Scanning

If antivirus scanning is enabled for FTP (File Transfer Protocol) traffic in a content security profile, the security device monitors the control channel and, when it detects one of the FTP commands for transferring data, it scans the data sent over the data channel.

NOTE: In order to scan FTP traffic, the FTP ALG must be enabled.

FTP Scanning: Process Overview

This is a general description of how FTP traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. A local FTP client opens an FTP control channel to an FTP server and requests the transfer of some data.
2. The FTP client and server negotiate a data channel over which the server sends the requested data. The security device intercepts the data and passes it to the antivirus scan engine, which scans it for viruses.
3. After completing the scan, the device follows one of two courses:
   - If there is no virus, the device forwards the data to the client.
   - If there is a virus, the device replaces the data with a drop message in the data channel and sends a message reporting the infection in the control channel.

SMTP Scanning

If SMTP (Simple Mail Transfer Protocol) antivirus scanning is enabled in a content security profile, the security device redirects traffic from local SMTP clients to the antivirus scanner before sending it to the local mail server.

NOTE: Chunking is an alternative to the data command. It provides a mechanism to transmit a large message in small chunks. It is not supported. Messages using chunking are bypassed and are not scanned.

This topic contains:

- SMTP Scanning: Process Overview on page 652
- Mail Message Replacement on page 652
- Sender Notification on page 652
- Subject Tagging on page 653

SMTP Scanning: Process Overview

This is a general description of how SMTP traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. An SMTP client sends an e-mail message to a local mail server or a remote mail server forwards an e-mail message via SMTP to the local mail server.
2. The security device intercepts the e-mail message and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the device follows one of two courses:
   - If there is no virus, the device forwards the message to the local server.
   - If there is a virus, the device sends a replacement message to the client.

Mail Message Replacement

If the antivirus scanner finds a virus in an e-mail message, the original message is dropped, the message body is truncated, and the content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> contains contaminated file
   <filename> with virus <virusname>, so it is dropped.
```

If a scan error is returned and the fail mode is set to drop, the original message is dropped and the entire message body is truncated. The content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> is dropped for <reason>.
```

Sender Notification

If notify-sender-on-virus is set and the message is dropped due to a detected virus, an e-mail is sent to the mail sender. The content of the notification may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you
   sent could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <ENVID> contaminated file <filename>
   with virus <virusname>.
e-mail Header is:
<header of scanned e-mail>
```

If notify-sender-on-error-drop is set and the message is dropped due to a scan error, an e-mail is sent to the mail sender of the scanned message. The content of the e-mail may appear as follows:

From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <ENVID> <reason>.
e-mail Header is:
<header of scanned e-mail>

NOTE: For information on the ENVID parameter, refer to RFC 3461.

Subject Tagging

If a scan error is returned and the fail mode is set to pass, the antivirus module passes the message through to the server. If notify-recipient-on-error-pass is set, the following string is appended to the end of the subject field:

(No virus check: <reason>)

POP3 Scanning

If Post Office Protocol 3 (POP3) antivirus scanning is enabled in a content security profile, the security device redirects traffic from a local mail server to antivirus scanner before sending it to the local POP3 client.

This topic contains:

- POP3 Scanning: Process Overview on page 653
- Mail Message Replacement on page 654
- Sender Notification on page 654
- Subject Tagging on page 654

POP3 Scanning: Process Overview

This is a general description of how POP3 traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. The POP3 client downloads an e-mail message from the local mail server.
2. The security device intercepts the e-mail message and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the security device follows one of two courses:
   - If there is no virus, the device forwards the message to the client.
   - If there is a virus, the device sends a message reporting the infection to the client.

NOTE: See "Protocol-Only Notifications" on page 676 for information on protocol-only notifications for IMAP.

Mail Message Replacement

If the antivirus scanner finds a virus in an e-mail message, the original message is dropped, the message body is truncated, and the content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> contains contaminated file
   <filename> with virus <virusname>, so it is dropped.
```

Sender Notification

If notify-sender-on-virus is set and the message is dropped due to a detected virus, an e-mail is sent to the mail sender.

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you
   sent could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> contaminated file <filename> with
   virus <virusname>.
e-mail Header is:
<header of scanned e-mail>
```

If notify-sender-on-error-drop is set and the message is dropped due to a scan error, an e-mail is sent to the mail sender of the scanned message. The content of the e-mail may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you
   sent could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <reason>.
e-mail Header is:
<header of scanned e-mail>
```

Subject Tagging

If a scan error is returned and the fail mode is set to pass, the antivirus module passes the message through to the server. If notify-recipient-on-error-pass is set, the following string is appended to the end of subject field:

```
(No virus check: <reason>)
```

IMAP Scanning

If IMAP (Internet Message Access Protocol) antivirus scanning is enabled in a content security profile, the security device redirects traffic from a local mail server to the internal antivirus scanner before sending it to the local IMAP client.

This topic contains:

- IMAP Scanning: Process Overview on page 655
- Mail Message Replacement on page 655
- Sender Notification on page 656
- Subject Tagging on page 656
- IMAP Antivirus Scanning Limitations on page 656

IMAP Scanning: Process Overview

This is a general description of how IMAP traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. The IMAP client downloads an e-mail message from the local mail server.
2. The security device intercepts the e-mail message and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the security device follows one of two courses:
   - If there is no virus, the device forwards the message to the client.
   - If there is a virus, the device sends a message reporting the infection to the client.

NOTE: See “Protocol-Only Notifications” on page 676 for information on protocol-only notifications for IMAP.

Mail Message Replacement

If the antivirus scanner finds a virus in an e-mail message, the original message is dropped, the message body is truncated, and the content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> contains contaminated file
  <filename> with virus <virusname>, so it is dropped.
```

Sender Notification

If notify-sender-on-virus is set and the message is dropped due to a detected virus, an e-mail is sent to the mail sender.

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you
   sent could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> contaminated file <filename> with
   virus <virusname>.
e-mail Header is:
<header of scanned e-mail>
```

If notify-sender-on-error-drop is set and the message is dropped due to a scan error, an e-mail is sent to the mail sender of the scanned message. The content of the e-mail may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you
   sent could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <reason>.
e-mail Header is:
<header of scanned e-mail>
```

Subject Tagging

If a scan error is returned and the fail mode is set to pass, the antivirus module passes the message through to the server. If notify-recipient-on-error-pass is set, the following string is appended to the end of subject field:

```
(No virus check: <reason>)
```

IMAP Antivirus Scanning Limitations

Mail Fragments — It is possible to chop one e-mail into multiple parts and to send each part through a different response. This is called mail fragmenting and most popular mail clients support it in order to send and receive large e-mails. Scanning of mail fragments is not supported by the antivirus scanner and in such cases, the message body is not scanned.

Partial Content — Some mail clients treat e-mail of different sizes differently. For example, small e-mails (less than 10 KB) are downloaded as a whole. Large e-mails (e.g. less than 1 MB) are chopped into 10 KB pieces upon request from the IMAP server. Scanning of any partial content requests is not supported by the antivirus scanner.

IMAP Uploads — Only antivirus scanning of IMAP downloads is supported. IMAP upload traffic is not scanned.

Configuring Full Antivirus Protection

You can use J-Web or the CLI to configure the full antivirus feature. When configuring antivirus protection, you must first create the antivirus custom objects you are using. Those custom objects may include the MIME pattern list, MIME exception list, and the filename extension list. Once you have created your custom objects, you can configure full antivirus protection, including intelligent prescreening, and content size limits.

This topic contains:

- Configuration Overview on page 657
- J-Web Configuration on page 658
- J-Web Point and Click CLI Configuration on page 664
- CLI Configuration on page 671

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any), for the feature in question. Custom objects are global parameters for UTM features. That is, configured custom objects apply to all UTM policies where applicable, rather than only to individual policies.

   The CLI commands for setting antivirus custom objects are:

   ```
   user@host# set security utm custom-objects mime-pattern
   user@host# set security utm custom-objects filename-extension
   user@host# set security utm custom-objects url-pattern
   user@host# set security utm custom-objects custom-url-category
   ```

2. Configure main feature parameters, called feature profiles.

   The CLI command for setting antivirus feature profiles is :

   ```
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
   ```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for HTTP (for example) and attaching that policy to a profile are:

   ```
   user@host# set security utm utm-policy <name>
   user@host# set security utm utm-policy utmp2 anti-virus http-profile http1
   ```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

   ```
   user@host# set security policies
   ```

```
user@host# set security policies from-zone trust to-zone untrust policy p2 then
    permit application-services utm-policy utmp2
```

J-Web Configuration

To configure antivirus protection using the J-Web configuration editor, you must first create your custom objects (MIME Pattern List, Filename Extension List, URL Pattern List, and Custom URL Category List).

Configure a MIME Pattern List Custom Object as follows :

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the MIME Pattern List tab, click the **Add** button to create MIME pattern lists.
3. In the Add MIME Pattern pop-up window, next to MIME Pattern Name, enter a unique name for the list you are creating.

   Keep in mind that you are creating a MIME white list and a MIME exception list (if necessary). Both MIME lists appear in the MIME Whitelist and Exception MIME Whitelist fields when you configure antivirus. Therefore, the MIME list names you create should be as descriptive as possible.
4. Next to MIME Pattern Value, enter the MIME pattern.
5. Click **Add** to add your MIME pattern to the Values list box.

   Within this box, you can also select an entry and use the Delete button to delete it from the list. Continue to add MIME patterns in this manner.
6. Optionally, create a new MIME list to act as an exception list.

   The exception list is generally a subset of the main MIME list.
7. Click **OK** to save the selected values as part of the MIME list you have created.
8. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a Filename Extension List Custom Object as follows :

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the Filename Extension List tab, click the **Add** button to create filename extension lists.
3. Next to File Extension Name, enter a unique name for the list you are creating.

   This name appears in the Scan Option By Extension list when you configure an antivirus profile.
4. In the Available Values box, select one or more default values (press Shift to select multiple concurrent items or press Ctrl to select multiple separate items) and click the right arrow button to move the value or values to the Selected Values box.

5. Click **OK** to save the selected values as part of the extension list you have created.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a URL Pattern List Custom Object as follows:

NOTE: Because you use URL Pattern Lists to create Custom URL Category Lists, you must configure URL Pattern List Custom Objects before you configure a Custom URL Category List.

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the URL Pattern List tab, click the **Add** button to create URL pattern lists.
3. Next to URL Pattern Name, enter a unique name for the list you are creating. This name appears in the Custom URL Category List Custom Object page for selection.
4. Next to URL Pattern Value, enter the URL or IP address you want added to the list for bypassing scanning.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

5. Click **Add** to add your URL pattern to the Values list box.

   The list can contain up to 8192 items. You can also select an entry and use the Delete button to delete it from the list. Continue to add URLs or IP addresses in this manner.
6. Click **OK** to save the selected values as part of the URL pattern list you have created.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a Custom URL Category List Custom Object as follows .

NOTE: Because you use URL Pattern Lists to create Custom URL Category Lists, you must configure URL Pattern List Custom Objects before you configure a Custom URL Category List. URL Pattern List Custom Objects are described in the previous section.

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **URL Category List** tab, click **Add** to create URL category lists.
3. Next to URL Category Name, enter a unique name for the list you are creating. This name appears in the URL Whitelist list when you configure antivirus global options.
4. In the Available Values box, select a URL Pattern List name from the list for bypassing scanning and click the right arrow button to move it to the Selected Values box.
5. Click **OK** to save the selected values as part of the custom URL list you have created.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Now that your custom objects have been created, you can configure the antivirus feature profile.

1. Select **Configure > Security > UTM > Global options**.
2. In the Anti-Virus tab, next to MIME whitelist, select the custom object you created from the list.
3. Next to Exception MIME whitelist, select the custom object you created from the list.
4. Next to URL Whitelist, select the custom object you created from the list.
5. In the Engine Type section, select the type of engine you are using.

   For full antivirus protection, you should select **Kaspersky Lab**.
6. In the Kaspersky Lab Engine Option section, enter the URL for the pattern database in the Pattern update URL box.

   Note that the URL is http://update.juniper-updates.net/AV/<device version> and you should not change it.
7. Next to Pattern update interval, enter the time interval, in seconds, for automatically updating the pattern database in the box.

   The default interval is 60.
8. Select whether you want the pattern file to update automatically (Auto update) or not (No Auto update).
9. Click **OK** to save the selected values.

10. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in a pop-up window that appears to discover why.
11. Select **Anti-Virus**, under Security, in the left pane.
12. Click **Add** in the right window to create a profile for the antivirus Kaspersky Lab Engine. (To edit an existing item, select it and click the **Edit** button.)
13. Next to Profile name, enter a unique name for this antivirus profile.
14. Select the Profile Type.

    In this case, select **Kaspersky**.
15. Next to Trickling timeout, enter timeout parameters. Note that trickling applies only to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing out during a file transfer or during antivirus scanning.
16. Next to Intelligent prescreening, select **Yes** or **No**.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP, and HTTP POST).

17. In the Scan Options section, next to Intelligent prescreening, select **Yes** if you are using it. .

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP, and HTTP POST).

18. Next to Content Size Limit, enter content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.
19. Next to Scan engine timeout, enter scanning timeout parameters.
20. Next to Decompress Layer Limit, enter decompression layer limit parameters.
21. In the Scan mode section, select either **Scan all files**, if you are scanning all content, or **Scan files with specified extension**, if you are scanning by file extensions.

    If you select Scan files with specified extension, you must select a filename extension list custom object from the Scan engine filename extention list that appears .
22. Select the **Fallback settings** tab.
23. Next to Default (fallback option), select **Log and permit** or **Block** from the list.

    Note that in most cases, Block is the default fallback option.
24. Next to Corrupt File (fallback option), select **Log and permit** or **Block** from the list.

25. Next to Password File (fallback option), select **Log and permit** or **Block** from the list.
26. Next to Decompress Layer (fallback option), select **Log and permit** or **Block** from the list.
27. Next to Content Size (fallback option), select **Log and permit** or **Block** from the list.
28. Next to Engine Not Ready (fallback option), select **Log and permit** or **Block** from the list.
29. Next to Timeout (fallback option), select **Log and permit** or **Block** from the list.
30. Next to Out Of Resources (fallback option), select **Log and permit** or **Block** from the list.
31. Next to Too Many Request (fallback option), select **Log and permit** or **Block** from the list.
32. Select the **Notification options** tab.
33. In the Fallback block section, next to Notification type, select **Protocol Only** or **Message** to select the type of notification that is sent when a fallback option of block is triggered.
34. Next to Notify mail sender, select **Yes** or **No**.
35. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
36. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
37. In the Fallback non block section, next to Notify mail recipient, select **Yes** or **No**.
38. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
39. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
40. Select the **Notification options cont** tab.
41. In the Virus detection section, next to Notification type, select **Protocol Only** or **Message** to select the type of notification that is sent when a fallback option of block is triggered.
42. Next to Notify mail sender, select **Yes** or **No**.
43. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
44. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).

    The limit is 255 characters.
45. Click **OK** .
46. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

NOTE: You create a separate antivirus profile for each antivirus protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for an antivirus profile, the UTM policy configuration page provides separate antivirus profile selection fields for each supported protocol.

Next, you configure a UTM policy for antivirus to which you attach the antivirus profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy.

   This takes you to the policy configuration pop-up window.
3. Select the **Main** tab in pop-up window.
4. In the Policy name box, enter a unique name for the UTM policy you are creating.
5. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.
6. For Session per client over limit, select one of the following: **Log and permit**, **Block**.

   This is the action the device takes when the session per client limit for this UTM policy is exceeded.
7. Select the **Anti-Virus profiles** tab in the pop-up window.
8. Select the appropriate profile you have configured from the list for the corresponding protocol listed.
9. Click **OK**.
10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.
2. From the Security Policy window , click **Add** to configure a security policy with UTM.

   This takes you to the policy configuration pop-up window.
3. In the Policy tab, enter a name in the Policy Name box.
4. Next to From Zone, select a zone from the list.
5. Next to To Zone, select a zone from the list.
6. Choose a **Source Address**.
7. Choose a **Destination Address**.
8. Choose an **Application**.

Do this by selecting junos-<protocol> (for all protocols that support antivirus scanning) in the Application Sets box and clicking the right arrow —> button to move them to the Matched box.

9. Next to Policy Action, select **Permit**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

10. Select the **Application Services** tab in the pop-up window.
11. Next to UTM Policy, select the appropriate policy from the list.

    This attaches your UTM policy to the security policy.

NOTE: There are several fields on this page that are not described in this section. See the Security Policies section for detailed information on configuring security policies and all the available fields.

12. Click **OK** to save your policy.
13. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your new policy to apply it.

J-Web Point and Click CLI Configuration

To configure antivirus protection using the J-Web Point and Click CLI, you must first create your custom objects. (MIME Pattern List, Filename Extension List, URL Pattern List, and Custom URL Category List)

Configure a MIME Pattern List Custom Object as follows (see “MIME White List” on page 649 for overview information on MIME white lists):

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Mime pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the MIME list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the MIME pattern.
9. Click **OK**. See Figure 107 on page 665.

Figure 107: Custom Object, Mime Pattern Configuration, Point and Click CLI Configuration

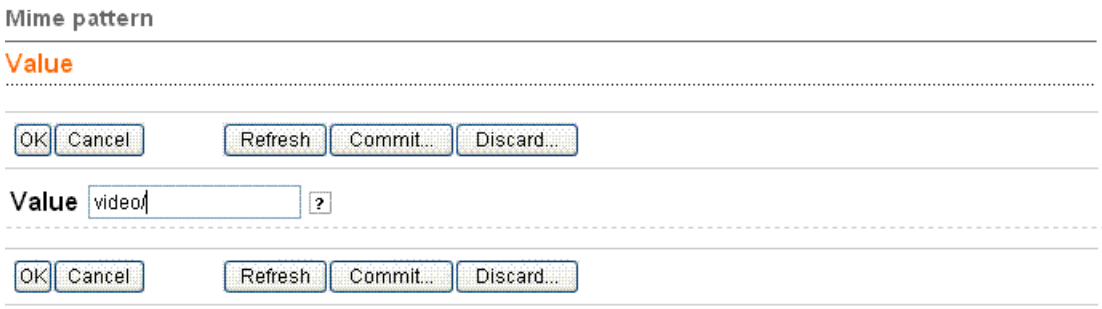

Configure a Filename Extension List Custom Object as follows (see "File Extension Scanning" on page 641):

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Filename extension, click **Add new entry**.
6. Next to Name, enter a unique name for the extension list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the extensions.
9. Click **OK**. See Figure 108 on page 665.

Figure 108: Custom Object, Filename Extension Configuration, Point and Click CLI Configuration

Configure a URL Pattern List Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.

4. Next to Custom objects, click **Configure**.
5. Next to Url pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the URLs or IP addresses you want added to list for bypassing scanning.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

9. Click **OK**. See Figure 109 on page 666.

Figure 109: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

Configure a Custom URL Category List Custom Object as follows (see "URL White List" on page 650 for overview information):

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Custom url category, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.

8. Next to Value, enter the name of the Url pattern list you created for bypassing scanning.
9. Click **OK**.

Figure 110: Custom Object, Custom URL Category Configuration, Point and Click CLI Configuration

Now that your custom objects have been created, you can configure the antivirus feature profile.

1. Select the **Configure > CLI Tools > Point and Click CLI** .
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature profile, click **Configure**.
5. Next to Anti virus, click **Configure**.
6. Next to Kaspersky lab engine, click **Configure**.
7. Next to Pattern update, select the **Yes** check box and click **Edit** to set the interval.
8. Next to Admin email, enter the e-mail addresses of the administrators who should receive e-mail notifications when updates are made to the pattern file.
9. Next to Custom message, enter the text to appear in the body of the notification e-mail.
10. Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
11. Next to Interval, enter the time interval for automatically updating the pattern database in the box. The default interval is 60. See “Updating Antivirus Patterns” on page 645 for more information.
12. Next to No autoupdate, select the No checkbox if you want to disable automatic updates and update the pattern database manually.
13. Next to URL, if it is not already entered, enter the URL for the pattern database in the box. Note that the URL is http://update.juniper-updates.net/AV/ < device version > and you should not change it.
14. Click **OK**.

15. Next to Profile, click **Add new entry** to create a profile for the Kaspersky Lab Engine.
16. Next to Name, enter a unique name for the profile you are creating.
17. Next to Fallback options, select the **Yes** check box and click **Edit** to configure.
    - Next to Content Size, select Log and Permit or Block from the list.
    - Next to Corrupt File (fallback option), select Log and Permit or Block from the list.
    - Next to Password File (fallback option), select Log and Permit or Block from the list.
    - Next to Decompress Layer (fallback option), select Log and Permit or Block from the list.
    - Next to Engine Not Ready (fallback option), select Log and Permit or Block from the list.
    - Next to Timeout (fallback option), select Log and Permit or Block from the list.
    - Next to Out Of Resources (fallback option), select Log and Permit or Block from the list.
    - Next to Too Many Request (fallback option), select Log and Permit or Block from the list.
18. Click **OK**.
19. Next to Notification options, select the **Yes** check box and click **Configure**.
    - Next to Fallback block, select the **Yes** check box and click **Configure**.
      - Next to Custom message, enter text for your custom message for this notification.
      - Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
      - Next to Notify mail sender, select the **Yes** check box to enable this notification.
      - Next to Type, select **protocol-only** or **message** from the list to select the type of notification that is sent when a fallback option of block is triggered
      - Click **OK**.
    - Next to Fallback non block, select the **Yes** check box and click **Configure**.
      - Next to Custom message, enter text for your custom message for this notification.
      - Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)

- Next to Notify mail sender, select the **Yes** check box to enable this notification.
- Click **OK**.

- Next to Virus detection, select the **Yes** check box and click **Configure**.
  - Next to Custom message, enter text for your custom message for this notification.
  - Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
  - Next to Notify mail sender, select the **Yes** check box to enable this notification.
  - Next to Type, select **protocol-only** or **message** from the list to select the type of notification that is sent when a fallback option of block is triggered
  - Click **OK**.

20. Next to Scan options, select the **Yes** check box and click **Configure**.
    - Next to Content size limit, enter content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.
    - Next to Decompress layer limit, enter decompression layer limit parameters. See "Decompression Layer Limit" on page 643.
    - Next to Intelligent prescreening, select the **Yes** check box to enable intelligent prescreening if you are using it. See "Intelligent Prescreening" on page 642.
    - Next to Scan extension, enter the name of the Filename extension list custom object you created.
    - Next to Scan mode, select **all** or **by-extension** from the list.
    - Next to Timeout, enter scanning timeout parameters. See "Scanning Timeout" on page 643.
21. Next to Trickling, select the **Yes** check box and click **Configure**.
    - Next to Timeout, enter trickling timeout parameters. See "HTTP Trickling" on page 649.
22. Click **OK**.
23. Click **OK** again to save the Kaspersky lab engine profile.
24. Back on the main Anti virus page, next to Mime whitelist, select the **Yes** check box and click **Configure** or **Edit**.
25. Next to Name, enter the name of the Mime list custom object you created.
26. Click **OK**.
27. Next to Type, select kaspersky-lab-engine from the list.

28. Next to Url whitelist, enter the name of the URL whitelist custom object you created.
29. Click **OK**. See Figure 111 on page 670 for main Kaspersky lab engine profile page.

Figure 111: Antivirus Full Profile Configuration, Point and Click CLI Configuration

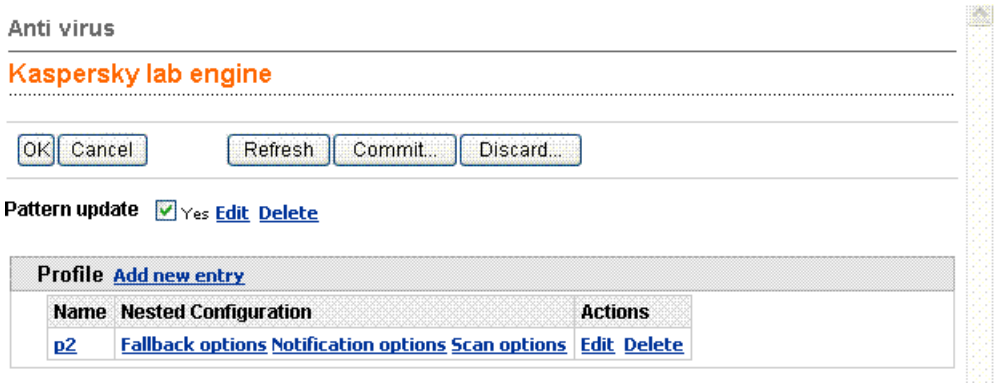

Next, you configure a UTM policy for full file-based antivirus to which you attach the Kaspersky lab engine profile you have configured.

1. Select the **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.
5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Anti virus, click **Configure**.
7. In the Http, Imap, Pop3, or Smtp profile boxes, enter the name of the profile you created earlier. For Ftp, click **Configure** or **Edit** to enter Upload and Download profiles.

NOTE: You create a separate antivirus profile for each antivirus protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for an antivirus profile, the UTM antivirus policy configuration page provides separate antivirus profile selection fields for each supported protocol.

8. Click **OK**.
9. Click **OK** again to return to main UTM configuration page. Your UTM antivirus policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select the **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Security, click **Configure**.
4. Next to Policy, select the **Yes** check box click **Edit**.
5. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

6. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
7. Click **OK**.

CLI Configuration

To configure antivirus protection using the CLI, you must first create your custom objects.

1. Configure the filename-extension custom object by first creating a name for the list. (Note that the Kaspersky scan engine ships with a read-only default extension list that you can use. See "File Extension Scanning" on page 641.)

```
user@host# set security utm custom-objects filename-extension extlist1
```

2. Next, add extensions to the list.

```
user@host# set security utm custom-objects filename-extension extlist1 value
  [zip js vbs]
```

3. Configure MIME lists. This includes creating a MIME whitelist and a MIME exception list for antivirus scanning. First create names for MIME lists and then add values to the lists. See "MIME White List" on page 649 for overview information on MIME white lists.

```
user@host# set security utm custom-objects mime-pattern avmime1
user@host# set security utm custom-objects mime-pattern ex-avmime1
```

4. Next, add MIME patterns to the lists.

```
user@host# set security utm custom-objects mime-pattern avmime1 value
  [video/quicktime image/x-portable-anymap x-world/x-vrml]
user@host# set security utm custom-objects mime-pattern ex-avmime1 value
  [video/quicktime-inappropriate]
```

Configure URL white lists for a list of URLs or addresses in a specified that you want to be bypassed by antivirus scanning. First create names for the URL list and then add values to the list. See “URL White List” on page 650 for overview information on URL white lists.

1. Configure a URL pattern list custom object by creating the list name and adding values to it as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

```
user@host# set security utm custom-objects url-pattern urllist1 value
  [http://www.url.com 5.6.7.8]
```

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use “*” if it is at the beginning of the URL and is followed by a “.”. You can only use “?” at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

2. Configure a custom URL category list custom object using the URL pattern list you created as follows:

```
user@host# set security utm custom-objects custom-url-category custurl1 value
  urllist1
```

Now that your custom objects have been created, you can configure the antivirus feature profile as follows:

1. Select and configure the engine type. Because you are configuring “full antivirus,” you select the Kaspersky-Lab-Engine and then designate the pattern update interval. The default full file-based antivirus pattern-update interval is 60 minutes. You can choose to leave this default as is or you can change it. You can also force a manual update, if necessary. See “Updating Antivirus Patterns” on page 645 for more information.

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
  pattern-update interval 20
```

NOTE: The command for changing the URL for the pattern database is:

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
  pattern-update url http://..
```

The default URL is http://update.juniper-update.net/AV/<device version>. You should not change this URL unless you are experiencing problems with it and have called for support.

2. You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
  pattern-update email-notify admin-email administrator@juniper.net
  custom-message "pattern file was updated" custom-message-subject "AV
  pattern file updated"
```

3. Configure a profile for the Kaspersky Lab engine. This profile includes configuring fallback, notification, and scanning options. It also includes HTTP trickling configuration options. First you create the profile name.

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
  kasprof1
```

4. Configure a list of fallback options as "block" or "log and permit." In most cases, the default is to block. You can use the default settings or you can change them. See "Fallback Options" on page 677 for a detailed overview of this feature and each fallback category.

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options content-size block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options corrupt-file block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options decompress-layer block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options default block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options engine-not-ready block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options out-of-resources block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options password-file block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options timeout block
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1
  fallback-options too-many-requests block
```

5. Configure the notification options. You can configure notifications for both fallback blocking and fallback nonblocking actions and for virus detection. See "Understanding Virus-Detected Notification Options" on page 676 for overview information on notification options.

You configure a custom message for the fallback blocking action and send a notification. See “Custom Message Notification” on page 677 for overview information.

user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** notification-options fallback-block custom-message ***virus-found*** **notify-mail-sender**

6. You can also configure a custom subject line for the custom message notification for both the sender and the recipient.

   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine kasprof1 notification-options fallback-block custom-message-subject “Antivirus Alert” notify-mail-sender

7. Configure a notification for protocol-only virus detection and send a notification. See “Protocol-Only Notifications” on page 676 for overview information.

   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** notification-options virus-detection type **protocol-only notify-mail-sender**

8. Configure scan options. See“Understanding the Internal Scan Engine and Scan Options” on page 640 for overview information.

   You configure content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size. See “Content Size Limits” on page 642 for details.

   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options content-size-limit 20000

9. Configure the decompression layer limit. See “Decompression Layer Limit” on page 643 for details.

   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options decompress-layer-limit **3**

10. Configure intelligent prescreening. It is either on or off. See “Intelligent Prescreening” on page 642 for details. (Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP) and HTTP POST.)

    user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options **intelligent-prescreening**

    user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options **no-intelligent-prescreening**

11. Configure scan extension settings. You can select the default list (junos-default-extension) or you can select an extension list you created as a custom object. See “File Extension Scanning” on page 641 for overview details.

    user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options scan-extension **extlist1**

12. Configure the scan mode setting. You can choose to scan all files or only files with extensions you specify. If you select to scan by-extension, the device uses the extension file lists you create.

    user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options scan-mode **by-extension**

13. Configure the timeout settings. See "Scanning Timeout" on page 643 for overview details.

    user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** scan-options scan-mode timeout **1800**

14. Configure trickling settings. If you use trickling, you can also set timeout parameters. Trickling applies only to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning. See "HTTP Trickling" on page 649 for overview details.

    user@host# set security utm feature-profile anti-virus kaspersky-lab-engine **kasprof1** **trickling** timeout **600**

15. Configure the antivirus scanner to use MIME bypass lists and exception lists. You can use your own custom object lists, or you can use the default list that ships with the device called junos-default-bypass-mime. See "MIME White List" on page 649 for overview information.)

    user@host# set security utm feature-profile anti-virus mime-whitelist list **avmime1**
    user@host# set security utm feature-profile anti-virus mime-whitelist list avmime1 exception **ex-avmime1**

16. Configure the antivirus module to use URL bypass lists. If you are using a URL white list, this is a custom URL category you have previously configured as custom object. URL white lists are valid only for HTTP traffic. See "URL White List" on page 650 for overview information.)

    user@host# set security utm feature-profile anti-virus url-whitelist **custurl1**

1. Configure a UTM policy for an antivirus scanning protocol and attach this policy to a profile. See "Global, Profile-Based, and Policy-Based Scan Settings" on page 644 for information on policy and profile-based scanning.

   CLI commands for configuring a UTM policy for HTTP antivirus scanning, and attaching that policy to a profile we created earlier for antivirus scanning, are:

   user@host# set security utm utm-policy <name>
   user@host# set security utm utm-policy **utmp2** anti-virus http-profile **kasprofile1**

2. Attach the UTM policy to a firewall security policy.

   user@host# set security policies from-zone trust to-zone untrust policy **p2** match source-address **any**
   user@host# set security policies from-zone trust to-zone untrust policy **p2** match destination-address **any**
   user@host# set security policies from-zone trust to-zone untrust policy **p2** match application **junos-http**

```
user@host# set security policies from-zone trust to-zone untrust policy p2  then
    permit application-services utm-policy utmp2
```

Understanding Virus-Detected Notification Options

You have multiple notifications options for when a virus is detected.

This topic contains:

- Protocol-Only Notifications on page 676
- E-mail Notification on page 676
- Custom Message Notification on page 677
- Fallback Options on page 677

Protocol-Only Notifications

When content is blocked because a virus is found or a scan error occurs, the client generally still receives a successful response code but with modified content (file replacement) containing a warning message. But with protocol-only notifications, a protocol-specific error code may be returned to the client. This way, the client determines that a virus was detected rather than interpreting that a file transfer succeeded.

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
    notification-options {
        virus-detection {
            type { protocol-only | message }
        }
        fallback-block {
            type { protocol-only | message }
                }
        }
    }
}
```

E-mail Notification

For mail protocols (SMTP, POP3, IMAP), e-mail notification is used to notify the sender or the recipient about the detected viruses or the scanning errors. There are three settings for e-mail notifications:

- virus-detection/notify-mail-sender — This setting is used when a virus is detected. If it is enabled, an e-mail is sent to the sender upon virus detection.
- fallback-block/notify-mail-sender — This setting is used when other scan codes or scanning errors are returned and the message is dropped. If it is enabled, an e-mail is sent to the sender when an error code is returned.
- fallback-non-block/notify-mail-recipient — This setting is used when other scan codes or scanning errors are returned and the message is passed. If it is enabled, the e-mail sent to the recipient is tagged when an error code is returned.

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
   notification-options {
      virus-detection {
         notify-mail-sender
         }
   fallback-block {
      notify-mail-sender
      }
   fallback-non-block {
   notify-mail-recipient
   }
   }
}
}
```

Custom Message Notification

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. When using custom messages, you can provide a customized message in the message content you can define customized subject tags

NOTE: Custom-message in fallback-nonblock is used only by mail protocols.

```
edit security utm feature-profile anti-virus kaspersky-lab-engine profile <name> {
   notification-options {
      virus-detection {
         custom-message <msg>
         custom-message-subject <subject msg>
      }
   fallback-block {
      custom-message <msg>
      custom-message-subject <subject msg>
   }
   fallback-non-block {
      custom-message <msg>
      custom-message-subject <subject msg>
      }
      }
   }
}
```

Fallback Options

Fallback options tell the system how to handle the errors returned by either the scan engine or the scan manager. The following is a list of possible errors and the default fallback actions for those error types.

- Scan engine is not ready (engine-not-ready)

  The scan engine is initializing itself, for example, loading the signature database. During this phase, it is not ready to scan a file. A file could either pass or be blocked according to this setting. The default action is BLOCK.

- Corrupt file (corrupt-file)

  Corrupt file is the error returned by the scan engine when engine detects a corrupted file. The default action is PASS.

- Decompression layer (decompress-layer)

  Decompress layer error is the error returned by the scan engine when the scanned file has too many compression layers. The default action is BLOCK.

- Password protected file (password-file)

  Password protected file is the error returned by the scan engine when the scanned file is protected by a password. The default action is PASS.

- Max content size (content-size)

  If the content size exceeds a set limit, the content is passed or blocked depending on the max-content-size fallback option. The default action is BLOCK.

- Too many requests (too-many-requests)

  If the total number of messages received concurrently exceeds the device limits, the content is passed or blocked depending on the too-many-request fallback option. The default action is BLOCK. (The allowed request limit is not configurable.)

- Timeout

  Scanning a complex file could consume resources and time. If the time it is taking to scan exceeds the timeout setting in the antivirus profile, the processing is aborted and the content is passed or blocked without completing the virus checking. The decision is made based on the timeout fallback option. The default action is BLOCK.

- Out of resources (out-of-resources)

  Virus scanning requires a great deal of memory and CPU resources. Due to resource constraints, memory allocation requests can be denied by the system. This failure could be returned by either scan engine (as a scan-code) or scan manager. When out-of-resources occurs, scanning is aborted. The default action is BLOCK.

- Default

  All the errors other than those in the above list fall into this category. This could include either unhandled system exceptions (internal errors) or other unknown errors. The default action is BLOCK.

Understanding Scan Result Handing

Different antivirus scan results are handled in different manners. For example, if a scan result is clean, the traffic is forwarded to the receiver. If the scan result is infected, the traffic is dropped. If the scan results in an error, the result handling depends on the cause of the failure and the configuration (fallback settings).

The following is a list of actions based on scan results:

- Scan Result = Pass

  The scan result handling action is to pass the message. In this case, no virus is detected and no error code is returned. Or, an error code is returned, but the fallback option for this error code is set to log-and-permit.

- Scan Result = Block

  The scan result handling action is to block the message. In this case, either a virus is detected or an error code is returned and the fallback option for this error code is BLOCK.

Verifying Antivirus Configurations

The antivirus module provides functions which allow you to use the CLI to check the system settings and the status of scan engine. It also provides functions to check the ongoing antivirus sessions and antivirus statistics.

This topic contains:

- General Scan Engine Status on page 679
- Session Status on page 680
- Verifying Antivirus Scan Results Using J-Web on page 680
- Verifying Antivirus Scan Results Using the CLI on page 682

General Scan Engine Status

Using the CLI, you can view the following scan engine status items:

Antivirus license key status

- View license expiration dates.

Scan engine status and settings

- View last action result.
- View default file extension list.

Antivirus pattern update server settings

- View update URL (HTTP or HTTPS-based).
- View update interval.

Antivirus pattern database status

- View auto update status.
- View last result of database loading.
- If the download completes, view database version timestamp virus record number.
- If the download fails, view failure reason.

Example status result:

CLI: show security utm anti-virus status

```
AV Key Expire Date: 03/01/2010 00:00:00
Update Server: http://update.juniper-updates.net/AV/SRX210
interval: 60 minutes
auto update status: next update in 12 minutes
last result: new database loaded
AV signature version: 12/21/2008 00:35 GMT, virus records: 154018
Scan Engine Info: last action result: No error(0x00000000)
```

Session Status

Using the CLI, you can view the following session status items:

Antivirus session status displays a snapshot of current antivirus sessions. It includes

- Maximum supported antivirus session numbers.
- Total allocated antivirus session numbers.
- Total freed antivirus session numbers.
- Current active antivirus session numbers.

CLI: show security utm session

Verifying Antivirus Scan Results Using J-Web

View antivirus scan results using J-Web as follows:

1. Select **Monitor > UTM > Anti-Virus**.

   The following information becomes viewable in the right pane.

Antivirus license key status

- View license expiration dates.

Antivirus pattern update server settings

- View update URL (HTTP or HTTPS-based).
- View update interval.

Antivirus pattern database status

- View auto update status.
- View last result of database loading.
- If the download completes, view database version timestamp virus record number.
- If the download fails, view failure reason.

Antivirus statistics provide

- The number of scan request being pre-windowed.
- The total number of scan request forwarded to the engine.
- The number of scan requests using scan-all mode.
- The number of scan requests using scan-by-extension mode.

Scan code counters provide

- Number of clean files.
- Number of infected files.
- Number of password protected files.
- Number of decompress layers.
- Number of corrupt files.
- When the engine is out of resources.
- When there is an internal error.

Fallback applied status provides either a log-and-permit or block result when the following has occurred

- Scan engine not ready.
- Password protected file found.
- Decompress layer too large.
- Corrupt file found.
- Out of resources.

- Timeout occurred.
- Maximum content size reached.
- Too many requests.
- Other.

2. You can click the **Clear Anti-Virus Statistics** button to clear all current viewable statistics and begin collecting new statistics.

Verifying Antivirus Scan Results Using the CLI

Using the CLI, you can view statistics for antivirus requests, scan results, and fallback counters.

Scan requests provide

- The total number of scan request forwarded to the engine.
- The number of scan request being pre-windowed.
- The number of scan requests using scan-all mode.
- The number of scan requests using scan-by-extension mode.

Scan code counters provide

- Number of clean files.
- Number of infected files.
- Number of password protected files.
- Number of decompress layers.
- Number of corrupt files.
- When the engine is out of resources.
- When there is an internal error.

Fallback applied status provides either a log-and-permit or block result when the following has occurred

- Scan engine not ready.
- Maximum content size reached.
- Too many requests.
- Password protected file found.
- Decompress layer too large.
- Corrupt file found.
- Timeout occurred.
- Out of resources.
- Other.

CLI: show security utm anti-virus statistics

Chapter 24

Express Antivirus Protection

Express antivirus scanning is offered as a less CPU intensive alternative to the full file-based antivirus feature.

The express Antivirus scanning feature is a separately licensed subscription service. When your antivirus license key expires, you can continue to use locally stored antivirus signatures. But in that case, if the local database is deleted, antivirus scanning is disabled.

NOTE: For information on the full antivirus feature, refer to that section in the manual. Note that if you switch from express antivirus protection to full file-based antivirus protection, you must reboot the device in order for full file-based antivirus to begin working.

Express antivirus supports the same protocols as full antivirus and functions in much the same manner. The main differences between express antivirus and full antivirus are as follows:

- Express antivirus uses a different antivirus scan engine and a different back-end hardware engine to accelerate pattern matching for higher data throughput.
- Express antivirus uses a different signature database, called Juniper Express antivirus database, compatible with the hardware engine, to target to only critical viruses and malware, including worms, Trojans, and spyware. This is a smaller sized database, providing less coverage than the full antivirus signature database.
- Express antivirus has a smaller memory footprint, compatible with the smaller system memory present on lower end devices.
- Express antivirus supports packet-based antivirus scanning. See “Packet-Based Scanning Versus File-Based Scanning” on page 686
- Express antivirus is not supported for J Series devices.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Packet-Based Scanning Versus File-Based Scanning on page 686
- Express Antivirus Limitations on page 686
- Expanded MIME Decoding Support on page 686

- Result Handling on page 687
- Updating Antivirus Patterns on page 687
- Intelligent prescreening Support on page 688
- Configuring Express Antivirus on page 688

Packet-Based Scanning Versus File-Based Scanning

The packet based scanning done by express antivirus provides virus scanning data buffers without waiting for entire file to be received by the firewall, whereas the file-based scanning done by full antivirus can only start virus scanning when entire file is received.

Express Antivirus Limitations

Express antivirus has the following limitations when compared to full antivirus functionality:

- Express antivirus provides limited support for the scanning of file archives and compressed file formats. Express antivirus can only support gzip, deflate and compressed compressing formats.
- Express antivirus provides limited support for decompression. Decompression is only supported with HTTP (supports only gzip, deflate, and compress for HTTP and only supports one layer of compression) and POP3 (supports only gzip for POP3 and only supports one layer of compression).
- Express antivirus does not support scanning by extension.
- Express antivirus scanning is interrupted when the scanning database is loading.
- Express antivirus may truncate a warning message if a virus has been detected and the replacement warning message that is sent is longer than the original content it is replacing.

NOTE: Because express antivirus does only packet based string matching, if you use the standard EICAR file to test express antivirus, you will see false positives. To avoid these false positives, Juniper has disabled scanning on the standard EICAR file to create a modified EICAR file for testing express antivirus. You can download this modified EICAR file from the following links:

http://www.juniper.net/security/avtest/ss-eicar.txt

http://www.juniper.net/security/avtest/ss-eicar.com

http://www.juniper.net/security/avtest/ss-eicar.zip

Expanded MIME Decoding Support

Express antivirus offers MIME decoding support for HTTP, POP3, SMTP, and IMAP.

MIME decoding support includes the following for each supported protocol:

- Multi-part and nested header decoding
- Base64 decoding, printed quote decoding, and encoded word decoding (in the subject field)

Result Handling

With express antivirus, the TCP traffic is closed gracefully when a virus is found and the data content is dropped.

NOTE: Express antivirus supports the following fail mode options: default, engine-not-ready, out-of-resource, and too-many-requests. Fail mode handling of supported options with express antivirus is much the same as with full antivirus.

Updating Antivirus Patterns

The express antivirus pattern database is updated over HTTP or HTTPS and can occur automatically or manually. This is similar functionality to that found in full antivirus with some minor differences.

The following list explains the differences found in express antivirus for database pattern updating functionality:

- With express antivirus, the signature database auto-update interval, is once a day.
- With express antivirus, there is no support for the downloading of multiple database types.
- With express antivirus, during database loading, all scan operations are interrupted. Scan operations for existing traffic flows are stopped and no new scan operations are initiated for newly established traffic flows. You can specify the desired action for this interruption period using the fall-back parameter for engine-busy-loading-database. The available actions are block or log-and-permit.
- By default, the URL for express antivirus is http://update.juniper-updates.net/EAV/SRX210. “SRX210” in the URL is the platform name. This part of the URL is different and platform specific for each platform. (Other than the platform name, you should not change this URL unless you are experiencing problems with it and have called for support.)

NOTE: Once your subscription expires, you have a 30 day grace period during which you can continue to update the antivirus pattern file. Once that grace period expires, the update server no longer permits antivirus pattern file updates.

Example: Automatic Update

In this example, you configure the security device to update the pattern file automatically every 120 minutes. (The default antivirus pattern-update interval is once a day.)

- J-Web

  1. Select **Configure > Security > UTM > Anti-Virus**.
  2. Next to Interval, in the Juniper Express Engine section, enter **120** in the box.
  3. Click **OK**.

- CLI

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  pattern-update interval 120
```

Example: Manual Update

In this example, you use the request command to update the pattern file manually.

- J-Web

  At this time, you cannot perform manual database updates using J-Web.

- CLI

```
user@host> request security utm anti-virus juniper-express-engine pattern-update
```

  You can also manually delete the pattern database (pattern-delete) and reload the pattern database (pattern-reload) using the request command.

Intelligent prescreening Support

This functionality is identical in both express antivirus and full antivirus. Refer to “Intelligent Prescreening” on page 642 for information.

Configuring Express Antivirus

You can use J-Web or the CLI to configure the express antivirus feature. When configuring express antivirus protection, you must first create the antivirus custom objects you are using. Those custom objects may include the MIME pattern list, MIME exception list, and the filename extension list. Once you have created your custom objects, you can configure express antivirus protection, including fallback options, intelligent prescreening, and content size limits.

This topic contains:

- Configuration Overview on page 689
- J-Web Configuration on page 689
- J-Web Point and Click CLI Configuration on page 695
- CLI Configuration on page 701

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any), for the feature in question. Custom objects are global parameters for UTM features. This means that configured custom objects apply to all UTM policies where applicable, rather than only to individual policies.

   The CLI commands for setting antivirus custom objects are:

   ```
   user@host# set security utm custom-objects mime-pattern
   user@host# set security utm custom-objects url-pattern
   user@host# set security utm custom-objects custom-url-category
   ```

2. Configure main feature parameters, called feature profiles.

   The CLI command for setting antivirus feature profiles is:

   ```
   user@host# set security utm feature-profile anti-virus juniper-exress-engine
   ```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for HTTP (for example) and attaching that policy to a profile are:

   ```
   user@host# set security utm utm-policy <name>
   user@host# set security utm utm-policy utmp3 anti-virus http-profile http1
   ```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

   ```
   user@host# set security policies
   user@host# set security policies from-zone trust to-zone untrust policy p3 then
     permit application-services utm-policy utmp3
   ```

J-Web Configuration

To configure express antivirus protection using the J-Web configuration editor, you must first create your custom objects (MIME pattern list, URL pattern list, and custom URL category list).

Configure a MIME pattern list custom object as follows (see “MIME White List” on page 649 for overview information on MIME white lists):

1. Select **Configure > Security > UTMCustom Objects**.
2. From the **MIME Pattern List** tab, click **Add** to create MIME pattern lists.
3. In the Add MIME Pattern pop-up window, next to MIME Pattern Name, enter a unique name for the list you are creating.

   Keep in mind that you are creating a MIME white list and a MIME exception list (if necessary). Both MIME lists appear in the MIME Whitelist and Exception MIME Whitelist fields when you configure antivirus. Therefore, the MIME list names you create should be as descriptive as possible.
4. Next to MIME Pattern Value, enter the MIME pattern.
5. Click **Add** to add your MIME pattern to the Values list box.

   Within this box, you can also select an entry and use the Delete button to delete it from the list. Continue to add MIME patterns in this manner.
6. Optionally, create a new MIME list to act as an exception list.

   The exception list is generally a subset of the main MIME list.
7. Click **OK** to save the selected values as part of the MIME list you have created.
8. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a URL Pattern List Custom Object as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **URL Pattern List** tab, click **Add** to create URL pattern lists.
3. Next to URL Pattern Name, enter a unique name for the list you are creating. This name appears in the Custom URL Category List Custom Object page for selection.
4. Next to URL Pattern Value, enter the URL or IP address you want added to list for bypassing scanning.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

5. Click **Add** to add your URL pattern to the Values list box. The list can contain up to 8192 items. You can also select an entry and use the Delete button to delete it from the list. Continue to add URLs or IP addresses in this manner.
6. Click **OK** to save the selected values as part of the URL pattern list you have created.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a custom URL category list custom object using the URL pattern list you created as follows (see "URL White List" on page 650 for overview information on URL white lists):

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **URL Category List** tab, click **Add** to create URL category lists.
3. Next to URL Category Name, enter a unique name for the list you are creating. This name appears in the URL Whitelist list when you configure antivirus global options.
4. In the Available Values box, select a URL Pattern List name from the list for bypassing scanning and click the right arrow button to move it to the Selected Values box.
5. Click **OK** to save the selected values as part of the custom URL list you have created.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Now that your custom objects have been created, you can configure the antivirus feature profile.

1. Select **Configure > Security > UTM > Global options**.
2. In the **Anti-Virus** tab, next to MIME whitelist, select the custom object you created from the list.
3. Next to Exception MIME whitelist, select the custom object you created from the list.
4. Next to URL Whitelist, select the custom object you created from the list.

5. In the Engine Type section, select the type of engine you are using.

   For express antivirus protection, you should select **Juniper Express**.

6. Next to Pattern update URL, enter the URL for the pattern database in the box. Note that the URL is http://update.juniper-updates.net/EAV/<device version> and you should not change it. See "Updating Antivirus Patterns" on page 687 for more information.

7. Next to Pattern update interval, enter the time interval for automatically updating the pattern database in the box.

   The default for express antivirus checking is once per day. See "Updating Antivirus Patterns" on page 687 for more information.

8. Select whether you want the pattern file to update automatically (Auto update) or not (No Auto update).

9. Click **OK** to save the selected values.

10. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

11. Select **Anti-Virus**, under Security, in the left pane.

12. Click **Add** in the right window to create a profile for the antivirus Juniper Express Engine.

    To edit an existing item, select it and click **Edit**.

13. In the Main tab, next to Profile name, enter a unique name for this antivirus profile.

14. Select the Profile Type.

    In this case, select **Juniper Express**.

15. Next to Trickling timeout, enter timeout parameters.

    Note that trickling only applies to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning. See "HTTP Trickling" on page 649 for overview details.

16. Next to Intelligent prescreening, select **Yes** or **No**.

    See "Intelligent Prescreening" on page 642 for details.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP, and HTTP POST).

17. Next to Content Size Limit, enter content size parameters.

    The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size. See "Content Size Limits" on page 642 for details.

18. Next to Scan engine timeout, enter scanning timeout parameters.

    See “Scanning Timeout” on page 643 for overview details.

19. Select the **Fallback settings** tab.
20. Next to Default (fallback option), select **Log and permit** or **Block** from the list.

    Note that in most cases, Block is the default fallback option. See “Fallback Options” on page 677 for a detailed overview of this feature and each fallback category.

21. Next to Decompress Layer (fallback option), select **Log and permit** or **Block** from the list.
22. Next to Content Size (fallback option), select **Log and permit** or **Block** from the list.
23. Next to Engine Not Ready (fallback option), select **Log and permit** or **Block** from the list.
24. Next to Timeout (fallback option), select **Log and permit** or **Block** from the list.
25. Next to Out of Resource (fallback option), select **Log and permit** or **Block** from the list.
26. Next to Too Many Requests (fallback option), select **Log and permit** or **Block** from the list.
27. Select the **Notification options** tab.
28. In the Fallback block section, next to Notification type, select **Protocol Only** or **Message** to select the type of notification that is sent when a fallback option of block is triggered.
29. Next to Notify mail sender, select **Yes** or **No**.
30. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
31. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
32. In the Fallback non block section, next to Notify mail recipient, select **Yes** or **No**.
33. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
34. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
35. Select the **Notification options cont** tab.
36. In the Virus detection section, next to Notification type, select **Protocol Only** or **Message** to select the type of notification that is sent when a fallback option of block is triggered.

    See “Protocol-Only Notifications” on page 676 for overview information.

37. Next to Notify mail sender, select **Yes** or **No**.
38. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).

39. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).

    The limit is 255 characters.

40. Click **OK**.

41. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up that appears window to discover why.

NOTE: You create a separate antivirus profile for each antivirus protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for antivirus, the UTM policy configuration page provides separate antivirus profile selection fields for each supported protocol.

Next, you configure a UTM policy for express antivirus to which you attach the antivirus profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy.

   The policy configuration pop-up window appears.

3. Select the **Main** tab in pop-up window.
4. In the Policy name box, enter a unique name for the UTM policy you are creating.
5. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.
6. For Session per client over limit , select one of the following: **Log and permit** or **Block**.

   This is the action the device takes when the session per client limit for this UTM policy is exceeded.

7. Select the **Anti-Virus profiles** tab in the pop-up window.
8. Select the appropriate profile you have configured from the list for the corresponding protocol listed.
9. Click **OK**.
10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.
2. From the Security Policy window, click **Add** to configure a security policy with UTM.

   The policy configuration pop-up window appears.

3. In the **Policy** tab, enter a name in the Policy Name box.
4. Next to Default Policy Action, select one of the following: **Deny-All** or **Permit-All**.
5. Next to From Zone, select a zone from the list.
6. Next to To Zone, select a zone from the list.
7. Under Zone Direction, click **Add a Policy**.
8. Choose a **Source Address**.
9. Choose a **Destination Address**.
10. Choose an Application. Do this by selecting **junos- < protocol >** (for all protocols that support antivirus scanning) in the Application Sets box and clicking the right arrow **— >** button to move them to the Matched box.
11. Next to Policy Action, select **Permit**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

12. Select the **Application Services** tab in the pop-up window.
13. Next to UTM Policy, select the appropriate policy from the list. This attaches your UTM policy to the security policy.

NOTE: There are several fields on this page that are not described in this section. See the Security Policies section for detailed information on configuring security policies and all the available fields.

14. Click **OK**.
15. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your new policy to apply it.

J-Web Point and Click CLI Configuration

To configure antivirus protection using the J-Web Point and Click CLI, you must first create your custom objects (MIME Pattern List, URL Pattern List, and Custom URL Category List).

Configure a MIME Pattern List Custom Object as follows (see “MIME White List” on page 649 for overview information on MIME white lists):

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Utm, click **Configure**.

4. Next to Custom objects, click **Configure**.
5. Next to Mime pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the MIME list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the MIME pattern.
9. Click **OK**. See Figure 112 on page 696.

Figure 112: Custom Object, MIME Pattern Configuration, Point and Click CLI Configuration

Configure a URL Pattern List Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Url pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the URLs or IP addresses you want added to list for bypassing scanning.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

9. Click **OK**. See Figure 113 on page 697.

Figure 113: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

Configure a Custom URL Category List Custom Object as follows (see "URL White List" on page 650 for overview information).

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Custom url category, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the name of the Url pattern list you created for bypassing scanning.
9. Click **OK**. See Figure 114 on page 697.

Figure 114: Custom Object, Custom URL Category Configuration, Point and Click CLI Configuration

Now that your custom objects have been created, you can configure the antivirus feature profile.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature profile, click **Configure**.
5. Next to Anti virus, click **Configure**.
6. Next to Juniper express engine, click **Configure**.
7. Next to Pattern update, select the **Yes** check box and click **Edit** to set the interval.
8. Next to Admin email, enter the e-mail addresses of the administrators who should receive e-mail notifications when updates are made to the pattern file.
9. Next to Custom message, enter the text to appear in the body of the notification e-mail.
10. Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
11. Next to Interval, enter the time interval for automatically updating the pattern database in the box. The default interval is 60. See "Updating Antivirus Patterns" on page 645 for more information.
12. Next to No autoupdate, select the **No** checkbox if you want to disable automatic updates and update the pattern database manually.
13. Next to URL, if it is not already entered, enter the URL for the pattern database in the box. Note that the URL is http://update.juniper-updates.net/AV/ < device version > and you should not change it.
14. Click **OK**.
15. Next to Profile, click **Add new entry** to create a profile for the Juniper Express Engine.
16. Next to Name, enter a unique name for the profile you are creating.
17. Next to Fallback options, select the **Yes** check box and click **Edit** to configure.
    - Next to Content Size, select**Log and Permit** or **Block**from the list.
    - Next to Default (fallback option), select **Log and Permit**or **Block**from the list.
    - Next to Engine Not Ready (fallback option), select **Log and Permit**or **Block**from the list.
    - Next to Out Of Resources (fallback option), select **Log and Permit**or **Block** from the list.
    - Next to Timeout (fallback option), select **Log and Permit**t or **Block**from the list.
    - Next to Too Many Request (fallback option), select **Log and Permit**or **Block**from the list.
18. Click **OK**.

19. Next to Notification options, select the **Yes** check box and click **Configure** or **Edit**.
    - Next to Fallback block, select the **Yes** check box and click **Configure** or **Edit**.
        - Next to Custom message, enter the text to appear in the body of the notification e-mail.
        - Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
        - Next to Notify mail sender, select the **Yes** check box to enable this notification.
        - Next to Type, select **protocol-only** or **message** from the list to select the type of notification that is sent when a fallback option of block is triggered
        - Click **OK**.
    - Next to Fallback non block, select the **Yes** check box and click **Configure** or **Edit**.
        - Next to Custom message, enter the text to appear in the body of the notification e-mail.
        - Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
        - Next to Notify mail sender, select the **Yes** check box to enable this notification.
        - Click **OK**.
    - Next to Virus detection, select the **Yes** check box and click **Configure** or **Edit**.
        - Next to Custom message, enter the text to appear in the body of the notification e-mail.
        - Next to Custom message subject, enter the text to appear in the subject line of the notification e-mail. (The limit is 255 characters.)
        - Next to Notify mail sender, select the **Yes** check box to enable this notification.
        - Next to Type, select **protocol-only** or **message** from the list to select the type of notification that is sent when a fallback option of block is triggered
        - Click **OK**.
20. Next to Scan options, select the **Yes** check box and click **Configure** or **Edit**.
    - Next to Content size limit, enter content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.
    - Next to Intelligent prescreening, select the **Yes** check box to enable intelligent prescreening if you are using it. See “Intelligent Prescreening” on page 642.

- Next to Timeout, enter scanning timeout parameters. See “Scanning Timeout” on page 643.

21. Next to Trickling, select the **Yes** check box and click **Configure** or **Edit**.
    - Next to Timeout, enter trickling timeout parameters. See “HTTP Trickling” on page 649.
22. Click **OK**.
23. Click **OK** again to save the Juniper express engine profile.
24. Back on the main Anti virus page, next to Mime whitelist, select the **Yes** check box and click **Configure** or **Edit**.
25. Next to Name, enter the name of the Mime list custom object you created.
26. Click **OK**.
27. Next to Type, select juniper-express-engine from the list.
28. Next to Url whitelist, enter the name of the URL whitelist custom object you created.
29. Click **OK**. See Figure 115 on page 700 for the Juniper express engine profile main window.

Figure 115: Antivirus Express Profile, Point and Click CLI Configuration

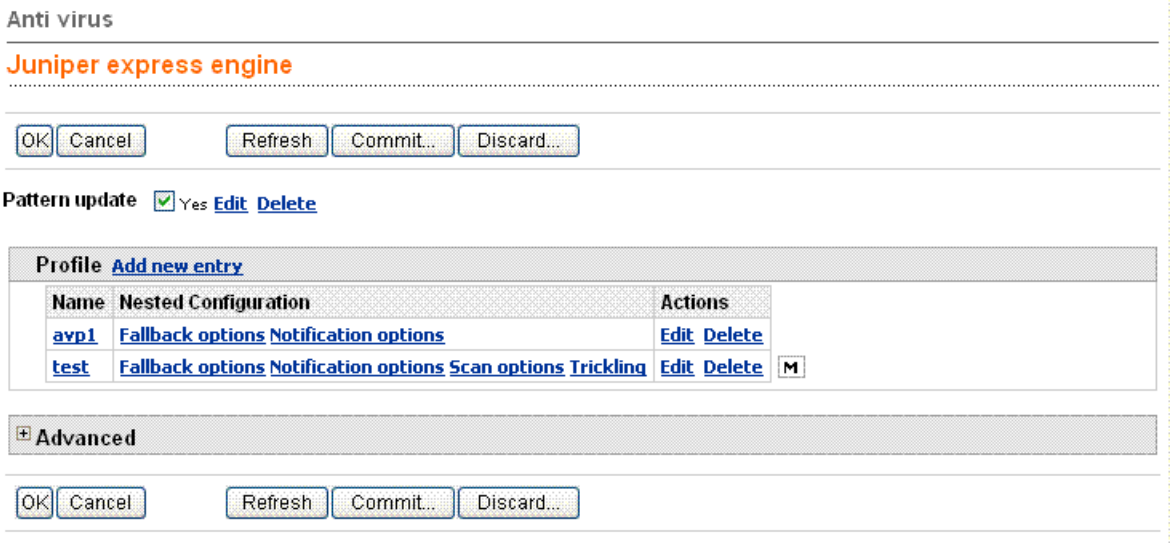

Next, you configure a UTM policy for express antivirus to which you attach the Juniper express engine profile you have configured.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.

5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Anti virus, click **Configure**.
7. In the Http, Imap, Pop3, or Smtp profile boxes, enter the name of the profile you created earlier. For Ftp, click **Configure** or **Edit** to enter Upload and Download profiles.

NOTE: You create a separate antivirus profile for each antivirus protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for an antivirus profile, the UTM antivirus policy configuration page provides separate antivirus profile selection fields for each supported protocol.

8. Click **OK**.
9. Click **OK** again to return to main UTM configuration page. Your UTM antivirus policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Policy, select the **Yes** check box click **Edit**.
5. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

6. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
7. Click **OK**.

CLI Configuration

To configure antivirus protection using the CLI, you must first create your custom objects.

1. Configure MIME lists. This includes creating a MIME whitelist and a MIME exception list for antivirus scanning. First create names for MIME lists and then add values to the lists. See “MIME White List” on page 649 for overview information on MIME white lists.

```
user@host# set security utm custom-objects mime-pattern avmime2
```

```
user@host# set security utm custom-objects mime-pattern ex-avmime2
```

2. Add MIME patterns to the lists.

```
user@host# set security utm custom-objects mime-pattern avmime2 value
  [video/quicktime image/x-portable-anymap x-world/x-vrml]
user@host# set security utm custom-objects mime-pattern ex-avmime2 value
  [video/quicktime-inappropriate]
```

Configure URL white lists for a list of URLs or addresses in a specified list that you want to be bypassed by antivirus scanning. First create names for the URL list and then add values to the list. See "URL White List" on page 650 for overview information on URL white lists.

1. Configure a URL pattern list custom object by creating the list name and adding values to it as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

```
user@host# set security utm custom-objects url-pattern urllist2 value
  [http://www.juniper.net 1.2.3.4]
```

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

2. Configure a custom URL category list custom object by using the URL pattern list you created as follows:

```
user@host# set security utm custom-objects custom-url-category custurl2 value
  urllist2
```

1. Now that your custom objects have been created, you can configure the antivirus feature profile. First, select and configure the engine type. Because you are configuring express antivirus, you select the juniper-express-engine.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
```

2. Select a time interval for updating the pattern database. The default antivirus pattern-update interval is once a day. You can choose to leave this default as is or you can change it. You can also force a manual update, if necessary. See "Updating Antivirus Patterns" on page 687 for more information.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  pattern-update interval 12
```

3. The command for changing the URL for the pattern database is:

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  pattern-update urlhttp://...
```

NOTE: Under most circumstances, you should not need to change the default URL.

4. You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  pattern-update email-notify admin-email administrator@juniper.net
  custom-message “pattern file was updated” custom-message-subject “AV
  pattern file updated”
```

5. Configure a profile for the Juniper-Express-Engine. This profile includes configuring fallback, notification, and scanning options. It also includes HTTP trickling configuration options. First, you create the profile name.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine profile
  junexprof1
```

6. Configure a list of fallback options as “block” or “log and permit.” In most cases, the default is to block. You can use the default settings or you can change them. See “Fallback Options” on page 677 for a detailed overview of this feature and each fallback category.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 fallback-options content-size block
user@host# set security utm feature-profile anti-virus juniper-express-engine
  jjunexprof1 fallback-options default block
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 fallback-options engine-not-ready block
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 fallback-options out-of-resources block
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 fallback-options timeout block
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 fallback-options too-many-requests block
```

7. Configure the notification options. You can configure notifications for both fallback blocking and fallback nonblocking actions and for virus detection. See “Understanding Virus-Detected Notification Options” on page 676 for overview information on notification options.

   In this step, configure a custom message for the fallback blocking action and send a notification. See “Custom Message Notification” on page 677 for overview information.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  jjunexprof1 notification-options fallback-block custom-message
  ***virus-found*** notify-mail-sender
```

8. Configure a notification for protocol-only virus detection and send a notification. See "Protocol-Only Notifications" on page 676 for overview information.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 notification-options virus-detection type protocol-only
  notify-mail-sender
```

9. You can also configure a custom subject line for the custom message notification for both the sender and the recipient.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  jjunexprof1 notification-options fallback-block custom-message-subject "Antivirus
  Alert" notify-mail-sender
```

10. Configure content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size. See "Content Size Limits" on page 642 for details.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 scan-options content-size-limit 20000
```

11. Configure intelligent prescreening. It is either on or off. See "Intelligent Prescreening" on page 642 for details. (Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP) and HTTP POST.)

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 scan-options intelligent-prescreening
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 scan-options no-intelligent-prescreening
```

12. Configure the time-out setting. See "Scanning Timeout" on page 643 for overview details.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 scan-options timeout 1800
```

13. Configure trickling settings. If you use trickling, you can also set timeout parameters. Trickling applies only to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing out during a file transfer or during antivirus scanning. See "HTTP Trickling" on page 649 for overview details.

```
user@host# set security utm feature-profile anti-virus juniper-express-engine
  junexprof1 trickling timeout 600
```

14. Configure the antivirus module to use MIME bypass lists and exception lists. You can use your own custom object lists or you can use the default list that ships with the device called "junos-default-bypass-mime." See "MIME White List" on page 649 for overview information.)

```
user@host# set security utm feature-profile anti-virus mime-whitelist list avmime2
```

```
user@host# set security utm feature-profile anti-virus mime-whitelist list avmime2
  exceptionex-avmime2
```

15. Configure the antivirus module to use URL bypass lists. If you are using a URL white list, this is a custom URL category you have previously configured as a custom object. URL white lists are only valid for HTTP traffic. See “URL White List” on page 650 for overview information.)

```
user@host# set security utm feature-profile anti-virus url-whitelist custurl2
```

1. Configure a UTM policy for an antivirus scanning protocol and attach this policy to a profile. See “Global, Profile-Based, and Policy-Based Scan Settings” on page 644 for information on policy and profile-based scanning.

   CLI commands for configuring a UTM policy for HTTP antivirus scanning and attaching that policy to a profile we created earlier for antivirus scanning are:

```
user@host# set security utm utm-policy <name>
user@host# set security utm utm-policy utmp3 anti-virus http-profile junexprof1
```

2. Attach the UTM policy to a firewall security policy.

```
user@host# set security policies from-zone trust to-zone untrust policy p3 match
  source-address any
user@host# set security policies from-zone trust to-zone untrust policy p3  match
  destination-address any
user@host# set security policies from-zone trust to-zone untrust policy p3  match
  application junos-http
user@host# set security policies from-zone trust to-zone untrust policy p3  then
  permit application-services utm-policy utmp3
```

NOTE: For information on Express Antivirus Notifications and Verification, refer to“Understanding Virus-Detected Notification Options” on page 676, and “Verifying Antivirus Configurations” on page 679.

Chapter 25

Content Filtering

Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, and protocol command. The content filter controls file transfers across the gateway by checking traffic against configured filter lists.

The content filter module evaluates traffic before all other UTM modules, except Web Filtering. Therefore, if traffic meets criteria configured in the content-filter, the content-filter acts first upon this traffic.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Types of Content Filters on page 707
- Content Filtering Protocol Support on page 708
- Configuring Content Filtering on page 710
- Verifying Content Filtering Configurations on page 722

Types of Content Filters

You can configure the following types of content filters:

- MIME Pattern Filter — MIME patterns are used to identify the type of traffic in HTTP and MAIL protocols. There are two lists of MIME patterns that are used by the content filter to determine the action to be taken. The block MIME list contains a list of MIME type traffic that is to be blocked by the content filter. The MIME exception list contains MIME patterns that are not to be blocked by the content filter and are generally subsets of items on the block list. Note that the exception list has a higher priority than the block list. If you have MIME entries that appear on both lists, those MIME types are not blocked by the content filter because the exception list takes priority. Therefore, when adding items to the exception list, it is to your advantage to be specific. See “MIME White List” on page 649 for detailed information on MIME lists.
- Block Extension List — Because the name of a file is available during file transfers, using file extensions is a highly practical way to block or allow file transfers. The content filter list contains a list of file extensions to be blocked. All protocols support the use of the block extension list. See “File Extension Scanning” on page 641 for information on file extension scanning.

- Protocol Command Block and Permit Lists — Different protocols use different commands to communicate between servers and clients. By blocking or allowing certain commands, traffic can be controlled on the protocol command level.

  The block and permit command lists are intended to be used in combination, with the permit list acting as an exception list to the block list.

NOTE: If a protocol command appears on the both the permit list and the block list, that command is permitted.

Because not all harmful files or components can be controlled by the MIME type or by the file extension, you can also use the content filter module to block ActiveX, Java Applets, and other types of content. The following types of content blocking are supported only for HTTP:

- Block ActiveX
- Block Java applets
- Block cookies
- Block EXE files
- Block ZIP files

Content Filtering Protocol Support

Each supported protocol may implement available content filters differently. Not all filtering capabilities are supported for each protocol.

This topic contains:

- HTTP Support on page 708
- FTP Support on page 709
- E-mail Support on page 709
- Example: Content Filter CLI on page 709

HTTP Support

The HTTP protocol supports all content filtering features. With HTTP, the content filter remains in the gateway, checking every request and response between the HTTP client and server.

If an HTTP request is dropped due to content filtering, the client receives a response such as:

```
<custom drop message/user-configured drop
  message>.<src_port><dst_ip>:<dst_port>Download request was dropped due to
  <reason>
```

Therefore, a message may appear as follows:

Juniper Networks Firewall Content Filtering blocked request. 5.5.5.1:80->4.4.4.1:55247 Download request was dropped due to file extension block list

FTP Support

The FTP protocol does not support all content filtering features. It supports only the following: Block Extension List and Protocol Command Block List.

When content filtering blocks an FTP request, the following response is sent through the control channel:

550 <src_ip>:<src_port>-<dst_ip>:<dst_port><custom drop message/user-configured drop message> for Content Filtering file extension block list.>

Therefore, a message may appear as follows:

550 5.5.5.1:21->4.4.4.1:45237 Requested action not taken and the request is dropped for Content Filtering file extension block list

E-mail Support

E-mail protocols (SMTP, IMAP, POP3) have limited content filtering support for the following features: Block Extension List, Protocol Command Block List, and MIME Pattern Filtering. Support is limited for e-mail protocols for the following reasons:

- The content filter scans only one level of an e-mail header. Therefore recursive e-mail headers and encrypted attachments are not scanned.
- If an entire e-mail is MIME encoded, the content filter can only scan for the MIME type.
- If any part of an e-mail is blocked due to content filtering, the original e-mail is dropped and replaced by a text file with an explanation for why the e-mail was blocked.

Example: Content Filter CLI

```
content-filtering {
   profile <name> {
      permit-command <cmd-list>
         block-command <cmd-list>
      block-extension <file-ext-list>
         block-mime {
            list <mime-list>
            exception <ex-mime-list>
         }
   block-content-type {
      activex
         java-applet
         exe
         zip
      http-cookie
```

```
    }
    notification-options {
        type { message }
            notify-mail-sender
        custom-message <msg>
    }
}
traceoptions {
    flag {
        all
            basic
            detail
        }
    }
}
```

Configuring Content Filtering

You can use J-Web or the CLI to configure the content filtering feature. When configuring content filtering, you must first create the custom objects you are using. Those custom objects may include protocol command lists, file extension lists, MIME pattern lists, and MIME exception lists. Once you have created your custom objects, you can configure content filtering,

This topic contains:

- Configuration Overview on page 710
- J-Web Configuration on page 711
- J-Web Point and Click CLI Configuration on page 715
- CLI Configuration on page 719

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any), for the feature in question. Custom objects are global parameters for UTM features. This means that configured custom objects apply to all UTM policies where applicable, rather than only to individual policies.

   The CLI commands for setting content filtering custom objects are:

   ```
   user@host# set security utm custom-objects protocol-command
   user@host# set security utm custom-objects filename-extension
   user@host# set security utm custom-objects mime-pattern
   ```

2. Configure main feature parameters, called feature profiles.

   The CLI command for setting content filtering feature profiles is:

   ```
   user@host# set security utm feature-profile content-filtering
   ```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for HTTP (for example) and attaching that policy to a profile are:

   ```
   user@host# set security utm utm-policy <name>
   user@host# set security utm utm-policy utmp4 content-filtering ftp download-profile
     ftp1
   ```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

   ```
   user@host# set security policies
   user@host# set security policies from-zone trust to-zone untrust policy p4 then
     permit application-services utm-policy utmp4
   ```

J-Web Configuration

To configure content filtering using the J-Web configuration editor, you must first create your custom objects (Protocol Command List, Filename Extension List, MIME Pattern List).

Configure a Protocol Command Custom Object as follows (See "Types of Content Filters" on page 707 for information on protocol commands.):

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **Protocol command List** tab, click **Add** to create command lists. (To edit an existing item, select it and click **Edit**.)
3. Next to Protocol Command Name, enter a unique name for the protocol list you are creating. (This name appears in the Permit command and Block command lists when you configure a content filter profile.)
4. Next to Protocol Command Value, enter the command for the protocol in question.
5. Click **Add** to add your protocol command to the Values list box.

   Within this box, you can also select an item and click Delete to remove it. Continue to add protocol commands in this manner.
6. Click **OK** to save the selected values as part of the protocol command list you have created.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a Filename Extension List Custom Object as follows :

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **Filename Extension List** tab, click **Add** to create extension lists.

3. Next to File Extension Name, enter a unique name for the list you are creating. (This name appears in the Block extension list when you configure a content filter profile.)
4. In the Available Values box, select one or more default values (press Shift to select multiple concurrent items or press Ctrl to select multiple separate items) and click the —> right arrow button to move the value or values to the Selected Values box.
5. Click **OK** to save the selected values as part of the extension list you have created.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a MIME Pattern List Custom Object as follows :

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **MIME Pattern List** tab, click **Add** to create MIME pattern lists.
3. In the Add MIME Pattern pop-up window, next to MIME Pattern Name, enter a unique name for the list you are creating.

   Keep in mind that you are creating a MIME block list and a MIME block exception list (if necessary). Both MIME lists appear in the Block MIME list and the Block MIME exception list fields when you configure content filtering. Therefore, the MIME list names you create should be as descriptive as possible.
4. Next to MIME Pattern Value, enter the MIME pattern.
5. Click **Add** to add your MIME pattern to the Values list box.

   Within this box, you can also select an entry and use the Delete button to delete it from the list. Continue to add MIME patterns in this manner.
6. Optionally, create a new MIME list to act as an exception list.

   The exception list is generally a subset of the main MIME list.
7. Click **OK** to save the selected values as part of the MIME list you have created.
8. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Now that your custom objects have been created, you can configure the content-filtering feature profile.

1. Select **Configure > Security > UTM > Content Filtering**.
2. Click **Add** to create a profile for content filtering. (To edit an existing item, select it and click **Edit**.)
3. Next to Profile name, enter a unique name in the box.
4. Next to Permit command list, select the protocol command custom object you created for permitting commands from the list.

The permit protocol command list is intended to act as an exception list for the block protocol command list.

NOTE: Protocol command lists, both permit and block, are created using the same custom object.

5. Next to Block command list, select the protocol command custom object you created for blocking commands from the list. (See "Types of Content Filters" on page 707 for information on protocol commands.)
6. Next to Block extension list, select the file extension list custom object you created for blocking extensions from the list.
7. Next to Block MIME list, select the MIME pattern list custom object you created for blocking MIME patterns from the list.
8. In the Block content type section, select content types in the Available content types box on the left and click the right arrow button — > to move items to the Selected content types box. (Press Shift to select multiple concurrent items or press Ctrl to select multiple separate items.)

NOTE: Block content type applies blocks to other available content such exe, http cookie, Java applet, and so on. The list of content types available from the Block content type box are supported only for HTTP blocking.

9. Select the **Notifications Options** tab.
10. Next to Notification type, select **Protocol** or **Message**.
11. Next to Notify mail sender, select **Yes** or **No**.
12. If you selected Yes, in the Custom notification message box, enter text for your custom message for this notification in the box (if you are using a custom message).
13. Click **OK**.
14. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

NOTE: You create a separate content filtering profile for each supported content filtering protocol. When you are creating your UTM policy for content filtering , the UTM policy configuration page provides separate content filtering profile selection fields for each supported protocol.

Next, configure a UTM policy for content filtering to which you attach the content filtering profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy.

   The policy configuration pop-up window appears.
3. In the **Main** tab, next to the Policy Name box, enter a unique name for the UTM policy you are creating.
4. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.
5. In Session per client over limit, select one of the following: **Log and permit** or **Block**.

   This is the action the device takes when the session per client limit for this UTM policy is exceeded.
6. Select the **Content filtering profiles** tab in the pop-up window.
7. Select the appropriate profile you have configured from the list for the corresponding protocol listed.
8. Click **OK**.
9. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.
2. From the Security Policy window, click **Add** to configure a security policy with UTM.

   The policy configuration pop-up window appears.
3. In the Policy tab, enter a name in the Policy Name box.
4. Next to From Zone, select a zone from the list.
5. Next to To Zone, select a zone from the list.
6. Choose a **Source Address.**
7. Choose a **Destination Address**.
8. Choose an Application. Do this by selecting junos- < protocol > (for all protocols that support content filtering) in the Application Sets box and clicking the — > button to move them to the Matched box.
9. Next to Default Policy Action, select one of the following: **Deny-All** or **Permit-All**.
10. Next to Policy Action, select one of the following: **Permit**, **Deny**, or **Reject**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

11. Select the **Application Services** tab in the pop-up window.
12. Next to UTM Policy, select the appropriate policy from the list.

    This attaches your UTM policy to the security policy.

NOTE: There are several fields on this page that are not described in this section. See the Security Policies section for detailed information on configuring security policies and all the available fields.

13. Click **OK** to save your policy.
14. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your new policy to apply it.

J-Web Point and Click CLI Configuration

To configure content filtering using the J-Web Point and Click CLI, you must first create your custom objects (Protocol Command List, Filename Extension List, MIME Pattern List).

Configure a Protocol Command Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Protocol command, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating. (This name appears in the Permit command and Block command lists when you configure a content filter profile.)
7. Next to Values, enter the commands for the protocol in question.
8. Click **OK**. See Figure 116 on page 716.

Figure 116: Custom Object, Protocol Command Configuration, Point and Click CLI Configuration

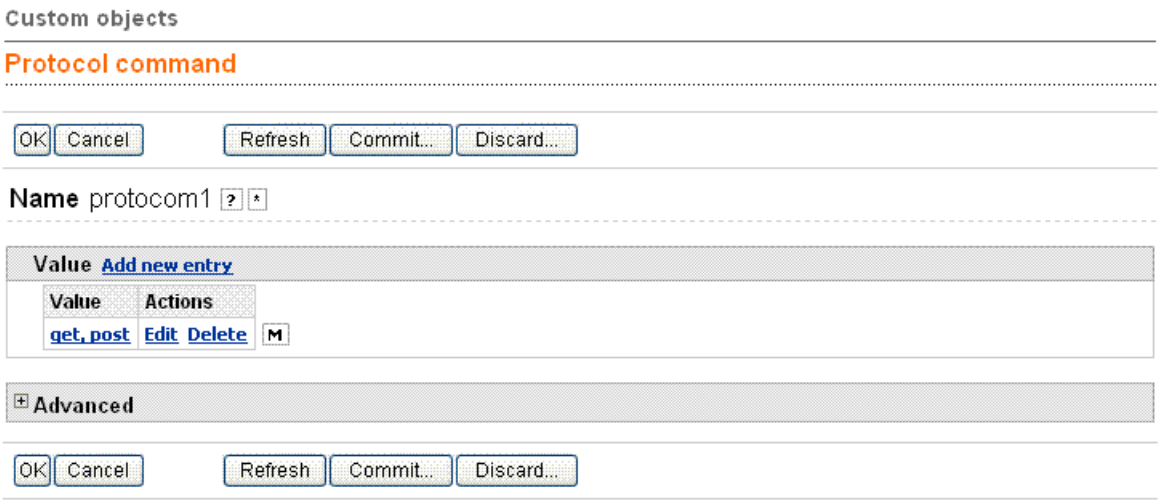

Configure a Filename Extension List Custom Object as follows :

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Filename extension, click **Add new entry**.
6. Next to Name, enter a unique name for the extension list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the extensions.
9. Click **OK**. See Figure 117 on page 716.

Figure 117: Custom Object, Filename Extension Configuration, Point and Click CLI Configuration

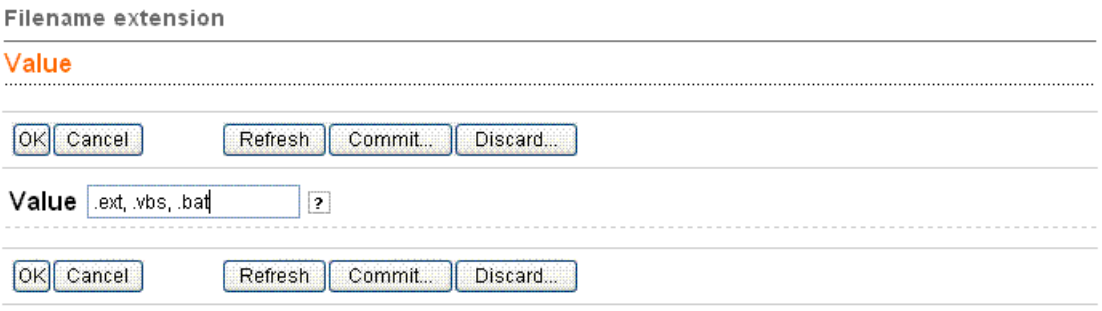

Configure a MIME Pattern List Custom Object as follows (see “MIME White List” on page 649 for overview information on MIME white lists):

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Mime pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the MIME list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the MIME pattern.
9. Click **OK**. See Figure 118 on page 717.

Figure 118: Custom Object, MIME Pattern Configuration, Point and Click CLI Configuration

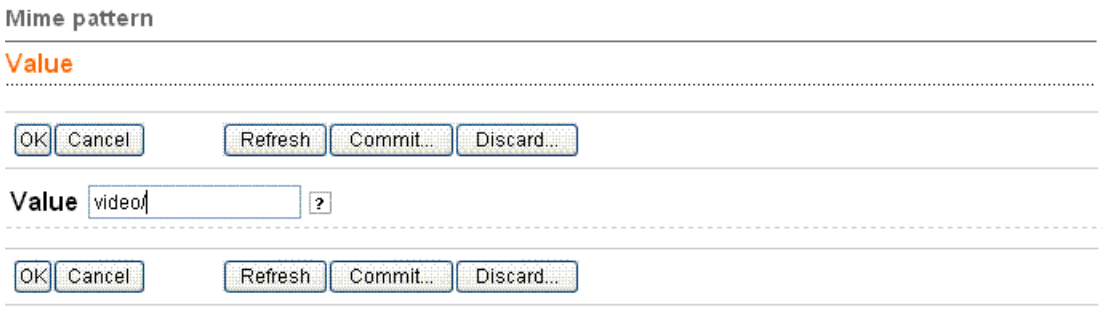

Now that your custom objects have been created, you can configure the content-filtering feature profile.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature profile, click **Configure**.
5. Next to Content filtering, click **Configure**.
6. Next to Profile, click **Add new entry**.
7. Next to Name, enter a unique name for this profile.
8. Next to Block command, enter the protocol command custom object you created for blocking commands from the list.
9. Next to Block content type, click **Configure** or **Edit**.
10. Select one or more of the available check boxes to block ActiveX, exe, Http cookie, Java applet, and zip file content.
11. Click **OK**.

12. Next to Block extension, enter the file extension list custom object you created for blocking extensions from the list.
13. Next to Block mime, select the **Yes** check box and click **Configure** or **Edit**.
14. Next to Exception, enter the exception Mime list custom object you created for Mime patterns that will not be blocked.
15. Next to List, enter the Mime list custom object you created for blocking Mime patterns.
16. Next to Notification options, select the **Yes** checkbox and click **Configure** or **Edit**.
    - Next to Custom message, enter text for your custom message for this notification.
    - Next to Notify mail sender, select the **Yes** check box to enable this notification.
    - Next to Type, select **message** as the type of notification that is sent when a fallback option of block is triggered
    - Click **OK**.
17. Next to Permit command, enter the protocol command custom object you created for permitting commands from the list. (The permit protocol command list is intended to act as an exception list for the block protocol command list.)
18. Click **OK**. See Figure 119 on page 718 for the content filtering profile main window.

Figure 119: Content Filtering Profile, Point and Click CLI Configuration

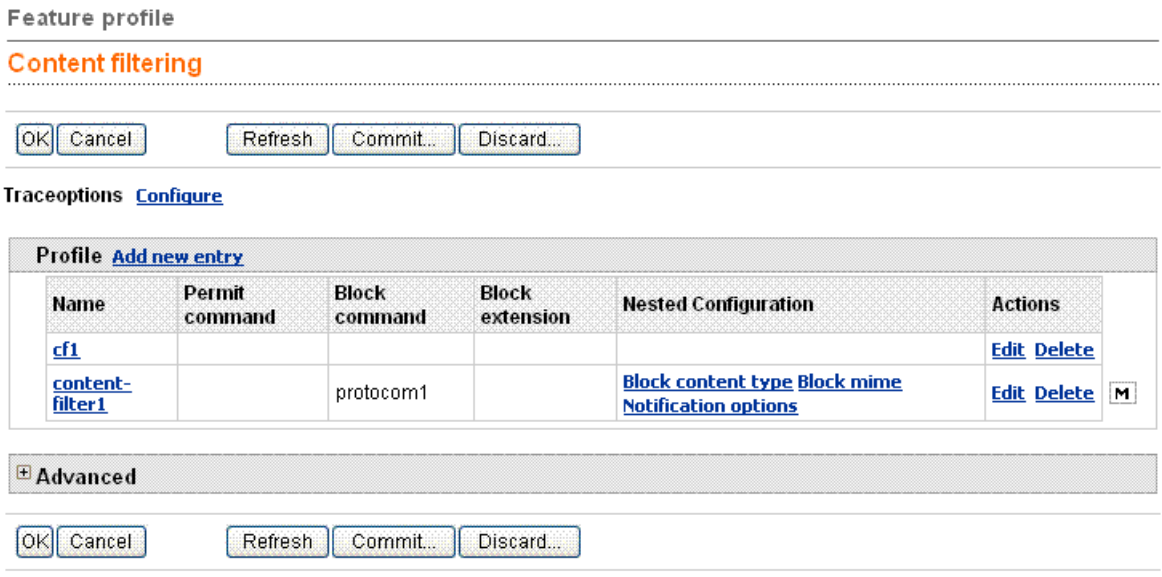

Next, you configure a UTM policy for content-filtering to which you attach the profile you have configured.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.
5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Content filtering, click **Configure**.
7. In the Http, Imap, Pop3, or Smtp profile boxes, enter the name of the profile you created earlier. For Ftp, click **Configure** or **Edit** to enter Upload and Download profiles.

NOTE: You create a separate content-filtering profile for each protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for a content-filtering profile, the UTM content-filter policy configuration page provides separate profile selection fields for each supported protocol.

8. Click **OK**.
9. Click **OK** again to return to main UTM configuration page. Your UTM content-filtering policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policy, select the **Yes** check box and click **Edit**.
4. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

5. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
6. Click **OK**.

CLI Configuration

To configure content filtering using the CLI, you must first create your custom objects.

1. Configure the protocol command list custom object by first creating a name for the list. See “Types of Content Filters” on page 707 for information on protocol commands.

   ```
   user@host# set security utm custom-objects protocol-command ftpprotocom1
   ```

2. Add commands to the list.

   ```
   user@host# set security utm custom-objects custom-objects protocol-command
     ftpprotocom1 value [user pass port type]
   ```

3. Configure the filename-extension custom object by first creating a name for the list. See “File Extension Scanning” on page 641 for information on file extension lists.

   ```
   user@host# set security utm custom-objects filename-extension extlist2
   ```

4. Add extensions to the list.

   ```
   user@host# set security utm custom-objects filename-extension extlist2 value
     [zip js vbs]
   ```

5. Configure MIME pattern lists. This includes creating a main MIME list and a MIME exception list for antivirus scanning. First create names for MIME lists and then add values to the lists. See “MIME White List” on page 649 for overview information on MIME pattern lists.

   ```
   user@host# set security utm custom-objects mime-pattern cfmime1
   user@host# set security utm custom-objects mime-pattern ex-cfmime1
   ```

6. Add MIME patterns to the lists.

   ```
   user@host# set security utm custom-objects mime-pattern cfmime1 value
     [video/quicktime image/x-portable-anymap x-world/x-vrml]
   user@host# set security utm custom-objects mime-pattern ex-cfmime1 value
     [video/quicktime-inappropriate]
   ```

Now that your custom objects have been created, you can configure the content-filtering feature profile.

1. Create a profile as follows.

   ```
   user@host# set security utm feature-profile content-filtering profile confilter1
   ```

2. Apply protocol block command custom objects to the content-filtering profile.

   ```
   user@host# set security utm feature-profile content-filtering profile confilter1
     block-command ftpprotocom1
   ```

3. Apply blocks to other available content such as exe, http-cookie, java-applet, and so on. The list of content types available from the “block-content-type” command are only supported for HTTP blocking.

   ```
   user@host# set security utm feature-profile content-filtering profile confilter1
     block-content-type java-applet exe http-cookie
   ```

4. Apply extension list custom objects to the content-filtering profile for blocking extensions.

   user@host# set security utm feature-profile content-filtering profile **confilter1** block-extension **extlist2**

5. Apply MIME pattern list custom objects to the content-filtering profile for blocking MIME types. If configured, you can also apply a MIME exception list.

   user@host# set security utm feature-profile content-filtering profile **confilter1** block-mime list **cfmime1** exception **ex-cfmime1**

6. Apply protocol permit command custom objects to the content-filtering profile. (The permit protocol command list is intended to act as an exception list for the block protocol command list.)

NOTE: Protocol command lists, both permit and block, are created by using the same custom object.

   user@host# set security utm feature-profile content-filtering profile **confilter1** permit-command **ftpprotocom2**

7. Next you configure the notification options. You can configure notifications with custom messages or configure no notification to be sent, as well. In this example, you configure a custom message and send a notification message. (You can configure a message notification for content-filtering.

   user@host# set security utm feature-profile content-filtering profile **confilter1** notification-options custom-message **"the action is not taken"** notify-mail-sender type **message**

1. Configure a UTM policy for a content filtering protocol and attach this policy to a profile. CLI commands for configuring a UTM policy for HTTP content filtering and attaching that policy to a profile you created earlier for content filtering are:

   user@host# set security utm utm-policy <name>
   user@host# set security utm utm-policy **utmp4** content-filtering http-profile **contentfilter1**

2. Attach the UTM policy to a firewall security policy.

   user@host# set security policies from-zone trust to-zone untrust policy **p4** match source-address **any**
   user@host# set security policies from-zone trust to-zone untrust policy **p4** match destination-address **any**
   user@host# set security policies from-zone trust to-zone untrust policy **p4** match application **junos-htttp**
   user@host# set security from-zone trust to-zone untrust policy **p4** then permit application-services utm-policy **utmp4**

Verifying Content Filtering Configurations

J-Web monitoring tools consist of the options that appear when you select Monitor in the task bar. The Monitor options display diagnostic information about the feature in question. Alternatively, you can enter show commands from the CLI to display the same information. CLI show commands display the current configuration and information about the feature in question.

This topic contains:

- Using J-Web to Monitor Content Filtering on page 722
- Using the CLI to Monitor Content Filtering on page 722

Using J-Web to Monitor Content Filtering

View content filtering statistics using J-Web as follows:

1. Select **Monitor > Security > UTM > Content Filtering**.

   The following statistics becomes viewable in the right pane.

   ```
   Base on command list: # Passed # Blocked
   Base on mime list: # Passed # Blocked
   Base on extension list: # Passed # Blocked
   ActiveX plugin: # Passed # Blocked
   Java applet: # Passed # Blocked
   EXE files: # Passed # Blocked
   ZIP files: # Passed # Blocked
   HTTP cookie: # Passed # Blocked
   ```

2. You can click **Clear Content filtering statistics** to clear all current viewable statistics and begin collecting new statistics.

Using the CLI to Monitor Content Filtering

Use the following commands to view content filtering statistics

```
user@host > show security utm content-filtering statistics
```

The content filtering show statistics command displays the following information:

```
Base on command list: # Blocked
Base on mime list: # Blocked
Base on extension list: # Blocked
ActiveX plugin: # Blocked
Java applet: # Blocked
EXE files: # Blocked
ZIP files: # Blocked
HTTP cookie: # Blocked
```

Chapter 26

Web Filtering

Web filtering lets you manage Internet usage by preventing access to inappropriate Web content.

There are three types of Web filtering solutions.

- The integrated Web filtering solution intercepts every HTTP request in a TCP connection. In this case, the decision making is done on the device after it identifies the category for a URL either from user-defined categories or from a category server (SurfControl Content Portal Authority provided by Websense).

NOTE: The integrated Web filtering feature is a separately licensed subscription service. When the license key for Web filtering has expired, no URLs are sent to the category server for checking, only local user-defined categories are checked.

- The redirect Web filtering solution intercepts HTTP requests and sends them to an external URL filtering server, provided by Websense, to determine whether to block the requests.

NOTE: Redirect Web filtering does not require a license.

- The local Web filtering solution intercepts every HTTP request in a TCP connection. In this case, the decision making is done on the device after it looks up a URL to determine if it is in the whitelist or blacklist based on its user-defined category.

NOTE: Local Web filtering does not require a license or a remote category server.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding Integrated Web Filtering on page 724
- Configuring Integrated Web Filtering on page 727

- Understanding Redirect Web Filtering on page 739
- Configuring Redirect Web Filtering on page 740
- Understanding Local Web Filtering on page 751
- Configuring Local Web Filtering on page 753
- Verifying Web Filtering Configurations on page 755
- What Happens When Both Web Filtering and Antivirus Scanning Are Employed? on page 756
- WebTrends Enhanced Log File Format (WELF) for UTM Features on page 756

Understanding Integrated Web Filtering

With integrated Web filtering, the firewall intercepts every HTTP request in a TCP connection and extracts the URL from the HTTP request. The decision making is done on the device after it identifies a category for a URL either from user-defined categories or from the SurfControl category server. Each individual HTTP request is blocked or permitted based on URL filtering profiles defined by you.

You can permit or block access to a requested site by binding a Web filtering profile to a firewall policy. A Web filtering profile specifies URL categories and the action the device takes (permit or block) when it receives a request to access a URL in each category. A URL category is a list of URLs grouped by content. URL categories are predefined and maintained by SurfControl or are defined by you.

- **SurfControl URL Categories**

  SurfControl maintains about 40 predefined categories. To view SurfControl global categories, type the following command:

  ```
  user@host# show groups junos-defaults
  ```

  You can also look for custom-url-category

- **User-Defined URL Categories**

  When defining your own URL categories, you can group URLs and create categories specific to your needs. Each category can have a maximum of 20 URLs. When you create a category, you can add either the URL or the IP address of a site. When you add a URL to a user-defined category, the device performs DNS lookup, resolves the host name into IP addresses, and caches this information. When a user tries to access a site with the IP address of the site, the device checks the cached list of IP addresses and tries to resolve the hostname. Many sites have dynamic IP addresses, meaning that their IP addresses change periodically. A user attempting to access a site can type an IP address that is not in the cached list on the device. Therefore, if you know the IP addresses of sites you are adding to a category, enter both the URL and the IP address(es) of the site.

  You define your own categories using URL pattern list and custom URL category list custom objects. Once defined, you can select your categories when you configure your Web filtering profile as follows:

```
user@host# set security utm feature-profile Web-filtering surf-control-integrated
    profile <profile_name> category
```

NOTE: If a URL appears in both a user-defined category and a predefined category, the device matches the URL to the user-defined category.

NOTE: Web filtering is performed on all the methods defined in HTTP 1.0 and HTTP 1.1.

Integrated Web Filtering: Process Overview

This is a general description of how Web traffic is intercepted and acted upon by the Web filtering module.

1. The device intercepts a TCP connection.
2. The device intercepts each HTTP request in the TCP connection.
3. The device extracts each URL in the HTTP request and checks its URL filter cache.
4. Global Web filtering white and blacklists are checked first for block or permit.
5. If the HTTP request URL is allowed based on cached parameters, it is forwarded to the Web server. If there is no cache match, a request for categorization is sent to the SurfControl server. (If the HTTP request URL is blocked, the request is not forwarded and a notification message is logged.)
6. In the allowed case, the SurfControl server responds with the corresponding category.
7. Based on the identified category, if the URL is permitted, the device forwards the HTTP request to the Web server. If the URL is not permitted, then a deny page is sent to the HTTP client.

Integrated Web Filtering Cache

By default, the device retrieves and caches the URL categories from the SurfControl CPA server. This process reduces the overhead of accessing the SurfControl CPA server each time the device receives a new request for previously requested URLs. You can configure the size and duration of the cache, according to the performance and memory requirements of your networking environment. The lifetime of cached items is configurable between 1 and 1800 seconds with a default value of 300 seconds.

NOTE: Caches are not preserved across device reboots or power losses.

Web Filtering Profiles

You configure Web filtering profiles that permit or block URLs according to defined categories. A Web filtering profile consists of a group of URL categories assigned one of the following actions:

- Permit — The device always allows access to the websites in this category.
- Block — The device blocks access to the websites in this category. When the device blocks access to this category of websites, it displays a message in your browser indicating the URL category.
- Blacklist — The device always blocks access to the websites in this list. You can create a user-defined category or use a predefined category.
- Whitelist — The device always allows access to the websites in this list. You can create a user-defined category or use a predefined category.

NOTE: A predefined profile is provided and can be used if you choose not to define your own profile.

A Web filtering profile may contain one blacklist or one whitelist, multiple user-defined and/or predefined categories each with a permit or block action, and an *Other* category with a permit or block action. You can define an action for all *Other* categories in a profile to specify what to do when the incoming URL does not belong to any of the categories defined in the profile. If the action for the *Other* category is block, the incoming URL is blocked if it does not match any of the categories explicitly defined in the profile. If an action for the *Other* category is not specified, the default action of permit is applied to the incoming URL not matching any category.

Profile Matching Precedence

When a profile employs several categories for URL matching, those categories are checked for matches in the following order:

1. If present, the global blacklist is checked first. If a match is made, the URL is blocked. If no match is found...
2. The global whitelist is checked next. If a match is made, the URL is permitted. If no match is found...
3. User-defined categories are checked next. If a match is made, the URL is blocked or permitted as specified. If no match is found...
4. Predefined categories are checked next. If a match is made, the URL is blocked or permitted as specified. If no match is found...
5. The Other category is checked next. If a match is made, the URL is blocked or permitted as specified.

Configuring Integrated Web Filtering

You can use J-Web or the CLI to configure the integrated Web filtering feature. When configuring Web filtering, if you are using custom objects, you must first create those. Custom objects may include URL Pattern Lists and Custom URL Category Lists. Once you have created your custom objects, you can configure Web filtering.

This topic contains:

- Configuration Overview on page 727
- J-Web Configuration on page 728
- J-Web Point and Click CLI Configuration on page 732
- CLI Configuration on page 736

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any), for the feature in question. Custom objects are global parameters for UTM features. This means that configured custom objects apply to all UTM policies where applicable, rather than only to individual policies.

   The CLI commands for setting integrated Web filtering custom objects are:

   ```
   user@host# set security utm custom-objects url-pattern
   user@host# set security utm custom-objects custom-url-category
   ```

2. Configure main feature parameters, called feature profiles.

   The CLI command for setting Web filtering feature profiles is:

   ```
   user@host# set security utm feature-profile Web-filtering
   ```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for Web filtering and attaching that policy to a profile are:

   ```
   user@host# set security utm utm-policy <name>
   user@host# set security utm utm-policy utmp5 web-filtering http-profile webhttp1
   ```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

   ```
   user@host# set security policies
   user@host# set security policies from-zone trust to-zone untrust policy p5 then
     permit application-services utm-policy utmp5
   ```

J-Web Configuration

To configure Web filtering using the J-Web Configuration editor, if you are using custom objects, you must first create those custom objects (URL pattern list, custom URL category list).

NOTE: Rather than or in addition to custom object lists, you can use included default lists and included whitelist and blacklist categories. See "Web Filtering Profiles" on page 726 for profile list information.

Configure a URL Pattern List Custom Object as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **URL Pattern List** tab , click **Add** to create URL pattern lists.
3. Next to URL Pattern Name, enter a unique name for the list you are creating. This name appears in the Custom URL Category List Custom Object page for selection.
4. Next to URL Pattern Value, enter the URL or IP address you want added to list for bypassing scanning.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

5. Click **Add** to add your URL pattern to the Values list box.

   The list can contain up to 8192 items. You can also select an entry and use the Delete button to delete it from the list. Continue to add URLs or IP addresses in this manner.
6. Click **OK** to save the selected values as part of the URL pattern list you have created.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a custom URL category list custom object as follows ():

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list. URL pattern list custom objects are described in "J-Web Configuration" on page 728.

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **URL Category List** tab, click **Add** to create URL category lists.
3. Next to URL Category Name, enter a unique name for the list you are creating. This name appears in the URL Whitelist, Blacklist, and Custom Category lists when you configure Web filtering global options.
4. In the Available Values box, select a URL Pattern List name from the list for bypassing scanning and click the right arrow button to move it to the Selected Values box.
5. Click **OK** to save the selected values as part of the custom URL list you have created.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Now that your custom objects have been created, you can configure the integrated Web filtering feature profile.

1. Select **Configure > Security > UTM > Global options**.
2. In the **Web Filtering**, next to URL whitelist, select the Custom URL list you created from the available options.

   This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the SurfControl server. See "Understanding Integrated Web Filtering" on page 724 for URL category information.
3. Next to URL blacklist, select the Custom URL list you created from the list.

   This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the SurfControl server.
4. In the Filtering Type section, select the type of Web filtering engine you are using.

   In this case, you would select **Surf Control Integrated**.
5. In the SurfControl Integrated options section, next to Cache timeout, enter a timeout limit, in minutes, for expiring cache entries (24 hours is the default and the maximum allowed life span). See "Integrated Web Filtering Cache" on page 725 for information on the cache.
6. Next to Cache Size, enter a size limit, in kilobytes, for the cache (500 KB is the default).

7. Next to Server Host, enter the Surf Control server name or IP address.
8. Next to Server Port, enter the port number for communicating with the Surf Control server (default ports are 80, 8080, and 8081).
9. Click **OK** to save these values.
10. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.
11. Select **Web Filtering**, under UTM, in the left pane.
12. In Web Filtering Profiles Configuration, click **Add** to create a profile for Surf Control Integrated Web filtering. (To edit an existing item, select it and click **Edit**.)
13. In Profile name, enter a unique name for this Web filtering profile.
14. Select the Profile Type. In this case, select **Surf Control**.
15. Next to Default action, select **Permit**, **Log and permit**, or **Block**.

    This is the default action for this profile for requests that experience errors.
16. Next to Custom Block Message, enter a custom message to be sent when HTTP requests are blocked.
17. Next to Timeout, enter a value, in seconds.

    Once this limit is reached, fail mode settings are applied. The default here is 10 seconds. You can enter a value from 10 to 240 seconds.
18. Next to Custom block message subject, enter text to appear in the subject line of your custom message for this block notification.
19. Select the **Fallback options** tab.
20. Next to Default, select **Log and Permit** or **Block** as the action to occur when a request fails for any reason not specifically called out.
21. Next to Server Connectivity, select **Log and Permit** or **Block** as the action to occur when a request fails for this reason.
22. Next to Timeout, select **Log and Permit** or **Block** as the action to occur when a request fails for this reason.
23. Next to Too Many Requests, select **Log and Permit** or **Block** as the action to occur when a request fails for this reason.
24. Click **Save**.
25. Select **Custom Objects**, under UTM, in the left pane.
26. Select the **URL category list** tab.
27. In the custom URL category list section, click **Add** to use a configured custom URL category list custom object in the profile.
28. Next to Categories, select a configured custom object from the list.
29. Next to Actions, select **Permit**, **Block**, or **Log and Permit** from the list.
30. Click **Add**.

32. Click **OK**.
33. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you configure a UTM policy for Web filtering to which you attach the content filtering profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy.

   The policy configuration pop-up window appears.
3. Select the **Main** tab in pop-up window.
4. In the Policy Name box, enter a unique name for the UTM policy you are creating.
5. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.
6. For Session per client over limit, select one of the following: **Log and Permit** or **Block**. This is the action the device takes when the session per client limit for this UTM policy is exceeded.
7. Select the **Web Filtering profiles** tab in the pop-up window.
8. Next to HTTP profile, select the profile you have configured from the list.
9. Click **OK**.
10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.
2. From the Security Policy window, click **Add** to configure a security policy with UTM.

   The policy configuration pop-up window appears.
3. In the **Policy** tab, enter a name in the Policy Name box.
4. Next to From Zone, select a zone from the list.
5. Next to To Zone, select a zone from the list.
6. Choose a **Source Address**.
7. Choose a **Destination Address**.
8. Choose an **Application**. Do this by selecting junos- < protocol > (for all protocols that support Web filtering, http in this case) in the Application Sets box and clicking the — > button to move them to the Matched box.
9. Next to Policy Action, select one of the following: **Permit**, **Deny**, or **Reject**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

10. Select the **Application Services** tab in the pop-up window.
11. Next to UTM Policy, select the appropriate policy from the list.

    This attaches your UTM policy to the security policy.

NOTE: There are several fields on this page that are not described in this section. See the section on Security Policies for detailed information on configuring security policies and all the available fields.

12. Click **OK**.
13. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your policy in order to apply it.

J-Web Point and Click CLI Configuration

To configure Web filtering using the J-Web Point and Click CLI, if you are using custom objects, you must first create those custom objects. (URL pattern list, custom URL category list).

Configure a URL Pattern List Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Url pattern, click **Add new entry**.
6. In the Name box, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. In the Value box, enter the URL pattern.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

9. Click **OK**. See Figure 120 on page 733.

Figure 120: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

NOTE: Continue to create as many Whitelist and Blacklist entries as necessary in your lists. You cannot use URL patterns directly. You must configure Custom URL category lists to which you apply your URL patterns.

Configure a Custom URL Category List Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Custom url category, click **Add new entry**.
6. In the Name box, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. In the Value box, enter the name of the Url pattern list you created for bypassing filtering.
9. Click **OK**. See Figure 121 on page 734.

Figure 121: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

Now that your custom objects have been created, you can configure the integrated Web filtering feature profile.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature profile, click **Configure**.
5. Next to Web filtering, click **Configure**.
6. Next to Surf control integrated, click **Configure**.
7. Next to Cache, select the **Yes** check box and click **Configure**.
8. In the Size box, enter a size limit, in kilobytes, for the cache. (500 KB is the default.)
9. In the Timeout box, enter a timeout limit in minutes for expiring cache entries. (24 hours is the default and the maximum allowed life span.) See “Integrated Web Filtering Cache” on page 725 for information on the cache.
10. Click **OK**.
11. Next to Server, select the **Yes** check box and click **Configure** or **Edit**.
12. In the Host box, enter the Surf Control server name or IP address.
13. In the Port box, enter the port number for communicating with the Surf Control server. (Default ports are 80, 8080, and 8081.)
14. Click **OK**.
15. Next to Profile, click **Add new entry**.
16. In the Name box, enter a unique name for this profile.
17. In the Custom block message box, enter a custom message to be sent when HTTP requests are blocked.
18. Next to Default, select **Permit**, **Block**, or **Log and Permit** from the list. This is the default action for this profile for requests that experience errors.
19. Next to Fallback settings, select the **Yes** check box and click **Configure**.

- Next to Default, select **block** or **log and permit** from the list as the action to occur when a request fails for any reason not specifically called out.
- Next to Server connectivity, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.
- Next to Timeout, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.
- Next to Too many requests, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.

20. Click **OK**.
21. Next to Timeout, enter a value, in seconds. Once this limit is reached, fail mode settings are applied. The default here is 10 seconds.
22. Next to Category, click **Add new entry**.
23. Next to Name, enter the name of the custom URL category list custom object you created.
24. Next to Action, select **permit**, **block**, or **log and permit** from the list.
25. Click **OK**.
26. Click **OK** again to return to the Web filtering main page.
27. Next to Type, select **surf-control-integrated** from the list.
28. Next to Url blacklist, enter the name of the custom URL list you created. This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the SurfControl server.
29. Next to Url whitelist, enter the name of the custom URL list you created. This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the SurfControl server.
30. Click **OK**.

Next, you configure a UTM policy for Web filtering to which you attach the profile you have configured.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.
5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Web filtering, click **Configure**.
7. In the Http profile box, enter the name of the Web filtering profile you created earlier.
8. Click **OK**. See Figure 122 on page 736 for Web filtering main profile window.

Figure 122: Web Filtering Integrated Profile Selection, Point and Click CLI Configuration

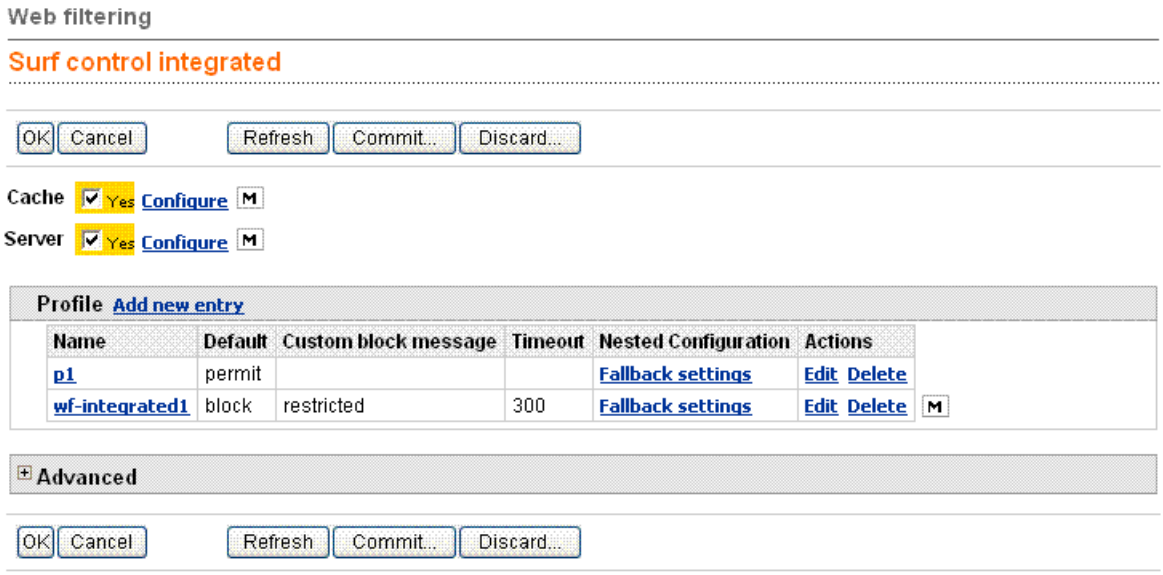

9. Click **OK** again to return to main UTM configuration page. Your UTM Web filtering policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Security, click **Configure**.
4. Next to Policy, select the **Yes** check box click **Edit**.
5. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

6. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
7. Click **OK**.

CLI Configuration

To configure integrated Web filtering using the CLI, you must first create your custom objects.

NOTE: Rather than or in addition to custom object lists, you can use included default lists and included whitelist and blacklist categories. See "Web Filtering Profiles" on page 726 for profile list information.

1. Configure a URL pattern list custom object by creating the list name and adding values to it as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

```
user@host# set security utm custom-objects url-pattern urllist3 value
    [http://www.juniper.net 1.2.3.4]
```

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

1. Configure a custom URL category list custom object by using the URL pattern list you created as follows:

```
user@host# set security utm custom-objects custom-url-category custurl3 value
    urllist3
```

Now that your custom objects have been created, you can configure the surf-control-integrated Web filtering feature profile.

1. If you are using included global whitelist and blacklist categories, select those global categories. This is the first filtering category that both integrated and redirect Web filtering use. If no match is made, the URL is forwarded to the SurfControl server. See "Understanding Integrated Web Filtering" on page 724 for URL category information.

```
user@host# set security utm feature-profile web-filtering url-blacklist
    Drugs_Alcohol_Tobacco
user@host# set security utm feature-profile web-filtering url-whitelist
    Computing_Internet
```

2. Select surf-control-integrated as your Web filtering engine as follows:

```
user@host# set security utm feature-profile web-filtering surf-control-integrated
```

3. Set the cache size parameters for surf-control-integrated Web filtering (500 KB is the default). See “Integrated Web Filtering Cache” on page 725 for information on the cache.

   user@host# set security utm feature-profile web-filtering surf-control-integrated cache size **500**

4. Set the cache timeout parameters for surf-control-integrated Web filtering (24 hours is the default and the maximum allowed life span). See “Integrated Web Filtering Cache” on page 725 for information on the cache.

   user@host# set security utm feature-profile web-filtering surf-control-integrated cache timeout **1800**

5. Set Surf Control server parameters by first entering the Surf Control server name or IP address.

   user@host# set security utm feature-profile web-filtering surf-control-integrated server host **surfcontrolserver**

6. Enter the port number for communicating with the Surf Control server. (Default ports are 80, 8080, and 8081.)

   user@host# set security utm feature-profile web-filtering surf-control-integrated server port **8080**

7. Create a surf-control-integrated profile by first creating a profile name as follows:

   user@host# set security utm feature-profile web-filtering surf-control-integrated profile **surfprofile1**

8. Select a category from the included whitelist and blacklist categories or select a custom URL category list you created for filtering against. Then enter an action (permit, log and permit, block) to go with the filter as follows (do this as many time as necessary to compile your whitelists and blacklists and their accompanying actions):

   user@host# set security utm feature-profile web-filtering surf-control-integrated profile **surfprofile1** category **custurl2** action **block**

9. Enter a custom message to be sent when HTTP requests are blocked.

   user@host# set security utm feature-profile web-filtering surf-control-integrated profile **surfprofile1** custom-block-message ***access denied***

10. Select a default action (permit, log and permit, block) for this profile for requests that experience errors.

    user@host# set security utm feature-profile web-filtering surf-control-integrated profile **surfprofile1** default **block**

11. Select fallback settings (block or log and permit) for this profile. The fallback actions are taken when errors in each configured category occur.

    user@host# set security utm feature-profile web-filtering surf-control-integrated profile **surfprofile1** fallback-settings **default block**

```
user@host# set security utm feature-profile web-filtering surf-control-integrated
  profile surfprofile1 fallback-settings server-connectivity block
user@host# set security utm feature-profile web-filtering surf-control-integrated
  profile surfprofile1 fallback-settings timeout block
user@host# set security utm feature-profile web-filtering surf-control-integrated
  profile surfprofile1 fallback-settings too-many-requests block
```

12. Enter a timeout value, in seconds. Once this limit is reached, fail mode settings are applied. The default here is 10 seconds. You can enter a value from 10 to 240 seconds.

```
user@host# set security utm feature-profile web-filtering surf-control-integrated
  profile surfprofile1 timeout 10
```

1. Configure a UTM policy for the Web filtering protocol and attach this policy to a profile. CLI commands for configuring a UTM policy for HTTP Web filtering, and attaching that policy to a profile you created earlier for content filtering are:

```
user@host# set security utm utm-policy <name>
user@host# set security utm utm-policy utmp5 web-filtering http-profile surfprofile1
```

2. Attach the UTM policy to a firewall security policy.

```
user@host# set security policies from-zone trust to-zone untrust policy p5  match
  source-address any
user@host# set security policies from-zone trust to-zone untrust policy p5 match
  destination-address any
user@host# set security policies from-zone trust to-zone untrust policy p5 match
  application junos-http
user@host# set security policies from-zone trust to-zone untrust policy p5 then
  permit application-services utm-policy utmp5
```

Understanding Redirect Web Filtering

With redirect Web filtering, the Web filtering module intercepts an HTTP request. The URL in the request is then sent to the external Websense server, which makes a permit or a deny decision. If access is permitted to the URL in question, the original HTTP request and all the subsequent requests are sent to the intended HTTP server. But if access is denied to the URL in question, a blocking message is sent to the client.

NOTE: Web filtering is performed on all the methods defined in HTTP 1.0 and HTTP 1.1. However, redirect Web filtering does not support HTTPS traffic because it cannot be decrypted to obtain the URL.

NOTE: Decision making from real-time options provides a higher level of accuracy, therefore caching for redirect Web filtering is not supported.

NOTE: Redirect Web filtering does not require a subscription license.

Redirect Web Filtering: Process Overview

This is a general description of how Web traffic is intercepted, redirected, and acted upon by the Web filtering module.

1. A Web client establishes a TCP connection with the Web server.
2. The Web client then sends an HTTP request.
3. The device intercepts the requests and extract URL. The URL is checked against Global Web filtering white and blacklists. If no match is made, the Websense server configuration parameters are utilized. Otherwise go to step 6.
4. The URL is sent to the Websense server for checking,
5. The Websense server returns a response indicating whether or not the URL is to be permitted or blocked.
6. If access is allowed, then the original HTTP request is sent to the Web server. If access is denied, the device sends a blocking message to the client and tears down the TCP connection.

Configuring Redirect Web Filtering

You can use J-Web or the CLI to configure the redirect Web filtering feature. When configuring Web filtering, if you are using custom objects, you must first create those. Custom objects may include URL pattern lists and custom URL category lists. Once you have created your custom objects, you can configure Web filtering.

This topic contains:

- Configuration Overview on page 740
- J-Web Configuration on page 741
- J-Web Point and Click CLI Configuration on page 745
- CLI Configuration on page 749

Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. First configure UTM custom objects (if any), for the feature in question. Custom objects are global parameters for UTM features. Configured custom objects apply to all UTM policies where applicable, rather than only to individual policies.

   ```
   user@host# set security utm custom-objects url-pattern
   user@host# set security utm custom-objects custom-url-category
   ```

2. Configure main feature parameters, called feature profiles.

The CLI command for setting Web filtering feature profiles is:

```
user@host# set security utm feature-profile Web-filtering
```

3. Configure a UTM policy for each protocol and attach this policy to a profile.

   CLI commands for configuring a UTM policy for Web filtering and attaching that policy to a profile are:

```
user@host# set security utm utm-policy <name>
user@host# set security utm utm-policy utmp6 web-filtering http-profile webhttp2
```

4. Attach the UTM policy to a firewall security policy.

   The CLI command for attaching a UTM policy to a security policy is:

```
user@host# set security policies
user@host# set security policies from-zone trust to-zone untrust policy p6 then
    permit application-services utm-policy utmp6
```

J-Web Configuration

To configure Web filtering using the J-Web configuration editor, if you are using custom objects, you must first create those custom objects. (URL pattern list, custom URL category list).

NOTE: Rather than or in addition to custom object lists, you can use included default lists and included whitelist and blacklist categories. See "Web Filtering Profiles" on page 726 for profile list information.

Configure a URL Pattern List Custom Object as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.

1. Select **Configure > Security > UTM > Custom Objects**.
2. From the **URL Pattern List** tab, click **Add** to create URL pattern lists.
3. Next to URL Pattern Name, enter a unique name for the list you are creating. This name appears in the Custom URL Category List Custom Object page for selection.
4. Next to URL Pattern Value, enter the URL or IP address you want added to list for bypassing scanning.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

5. Click **Add** to add your URL pattern to the Values list box.

   The list can contain up to 8192 items. You can also select an entry and use the Delete button to delete it from the list. Continue to add URLs or IP addresses in this manner.

6. Click **OK** to save the selected values as part of the URL pattern list you have created.

7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a custom URL category list custom object as follows (see "URL White List" on page 650 for overview information on URL whitelists):

NOTE: Because you use URL Pattern Lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.

1. Select **Configure > Security > UTM > Custom Objects**.

2. From the **URL Category List** tab, click **Add** to create URL category lists.

3. Next to URL Category Name, enter a unique name for the list you are creating. This name appears in the URL Whitelist, Blacklist, and Custom Category lists when you configure Web filtering global options.

4. In the Available Values box, select a URL Pattern List name from the list for bypassing scanning and click the right arrow button to move it to the Selected Values box.

5. Click **Add** to add your entry to the Values list box.

   Within this box, you can select an entry and use the up and down arrows to change the order of the list. You can also select an entry and use the X button to delete it from the list. Continue to add URLs or IP addresses in this manner.

6. Click **OK** to save the selected values as part of the custom URL category list you have created.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Now that your custom objects have been created, you can configure the redirect Web filtering feature profile.

1. Select **Configure > Security > UTM > Global options**.
2. In the **Web Filtering** tab, next to URL whitelist, select the Custom URL list you created from the available options.

   This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the Websense server.
3. Next to URL blacklist, select the Custom URL list you created from the list. This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the Websense server.
4. In the Filtering Type section, select the type of Web filtering engine you are using.

   In this case, you would select **Websense Redirect**.
5. Click **OK** to save these values.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.
7. Select **Web Filtering**, under UTM, in the left pane.
8. Click **Add** o create a profile for redirect Web filtering. (To edit an existing item, select it and click the **Edit** button.)
9. In the **Main** tab, next to Profile name, enter a unique name for this Web filtering profile.
10. Select the Profile Type. In this case, select **Websense**.
11. Next to Account, enter the user account for which this profile is intended.
12. Next to Server, enter the Websense server name.
13. Next to Port, enter the port number for communicating with the Websense server (default ports are 80, 8080, and 8081).
14. Next to Sockets, enter the number of sockets used for communicating between the client and server (the default here is 1).
15. Next to Timeout, enter a timeout limit for requests.

    Once this limit is reached, fail mode settings are applied. The default here is 10 seconds. You can enter a value from 10 to 240 seconds.
16. Next to Custom Block Message, enter a custom message to be sent when HTTP requests are blocked.
17. Select the **Fallback options** tab.

18. Next to Default Action, select **Log and permit** or **Block**.
19. Next to Server Connectivity, select **Log and permit** or **Block** as the action to occur when a request fails for this reason.
20. Next to Timeout, select **Log and permit** or **Block** as the action to occur when a request fails for this reason.
21. Next to Too Many Requests, select **Log and permit** or **Block** as the action to occur when a request fails for this reason.
22. Click the **OK** button to save these values.
23. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you configure a UTM policy for Web filtering to which you attach the content filtering profile you have configured.

1. Select **Configure > Security > Policy > UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy.

   The policy configuration pop-up window appears.
3. Select the **Main** tab in the pop-up window.
4. In the Policy Name box, enter a unique name for the UTM policy you are creating.
5. In the Session per client limit box, enter a session per client limit from 0 to 20000 for this UTM policy.
6. For Session per client over limit, select one of the following: **Log and Permit** or **Block**.

   This is the action the device takes when the session per client limit for this UTM policy is exceeded.
7. Select the **Web Filtering profiles** tab in the pop-up window.
8. Next to HTTP profile, select the profile you have configured from the list.
9. Click **OK**.
10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > Security > Policy > FW Policies**.
2. From the Security Policy window, click **Add** to configure a security policy with UTM. This takes you to the policy configuration pop-up window.
3. In the **Policy** tab, enter a name in the Policy Name box.
4. Next to From Zone, select a zone from the list.
5. Next to To Zone, select a zone from the list.
6. Choose a **Source Address**.

7. Choose a **Destination Address**.
8. Choose an Application. Do this by selecting junos- < protocol > (for all protocols that support Web filtering, http in this case) in the Application Sets box and clicking the — > button to move them to the Matched box.
9. Next to Default Policy Action, select one of the following: **Deny-All** or **Permit-All**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

10. Select the **Application Services** tab in the pop-up window.
11. Next to UTM Policy, select the appropriate policy from the list.

    This attaches your UTM policy to the security policy.

NOTE: There are several fields on this page that are not described in this section. See the section on Security Policies for detailed information on configuring security policies and all the available fields.

12. Click **OK** .
13. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your policy in order to apply it.

J-Web Point and Click CLI Configuration

To configure Web filtering using the J-Web Point and Click CLI, if you are using custom objects, you must first create those custom objects. (URL pattern list, custom URL category list).

Configure a URL Pattern List Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to URL pattern, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the URL pattern.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

9. Click **OK**. See Figure 123 on page 746.

Figure 123: Custom Object, URL Pattern Configuration, Point and Click CLI Configuration

NOTE: Continue to create as many Whitelist and Blacklist entries as necessary in your lists. You cannot use URL patterns directly. You must configure Custom URL category lists to which you apply your URL patterns.

Configure a Custom URL Category List Custom Object as follows:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Custom objects, click **Configure**.
5. Next to Custom url category, click **Add new entry**.
6. Next to Name, enter a unique name for the list you are creating.
7. Next to Value, click **Add new entry**.
8. Next to Value, enter the name of the Url pattern list you created for bypassing filtering.
9. Click **OK**. See Figure 124 on page 747.

Figure 124: Custom Object, Custom URL Category Configuration, Point and Click CLI Configuration

Now that your custom objects have been created, you can configure the redirect Web filtering feature profile.

1. Select the **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Feature profile, click **Configure**.
5. Next to Web filtering, click **Configure**.
6. Next to Type, select **surf-control-integrated** from the list.
7. Next to Url blacklist, enter the name of the custom URL list you created. This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the SurfControl server.
8. Next to Url whitelist, enter the name of the custom URL list you created. This is the first filtering category that both integrated and redirect Web filtering use. If there is no match, the URL is sent to the SurfControl server.
9. Next to Websense redirect, click **Configure** or **Edit**.
10. Next to Profile, click **Add new entry**.
11. Next to Name, enter a unique name for this profile.
12. Next to Account, enter the user account for which this profile is intended.
13. Next to Custom block message, enter a custom message to be sent when HTTP requests are blocked.
14. Next to Fallback settings, select the **Yes** check box and click **Configure**.
15. Next to Default, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.
16. Next to Server connectivity, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.
17. Next to Timeout, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.
18. Next to Too many requests, select **block** or **log and permit** from the list as the action to occur when a request fails for this reason.

19. Click **OK**.
20. Next to Server, select the **Yes** check box and click **Configure**.
21. Next to Host, enter the Web server name or IP address.
22. Next to Port, enter the port number for communicating with the Web server. (Default ports are 80, 8080, and 8081.)
23. Click **OK**.
24. Next to Sockets, enter the number of sockets used for communicating between the client and server. The default here is 1.
25. Next to Timeout, enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied. The default here is 10 seconds.
26. Click **OK**. See Figure 125 on page 748 for the Web filtering main profile window.

Figure 125: Web Filtering Redirect Profile, Point and Click CLI Configuration

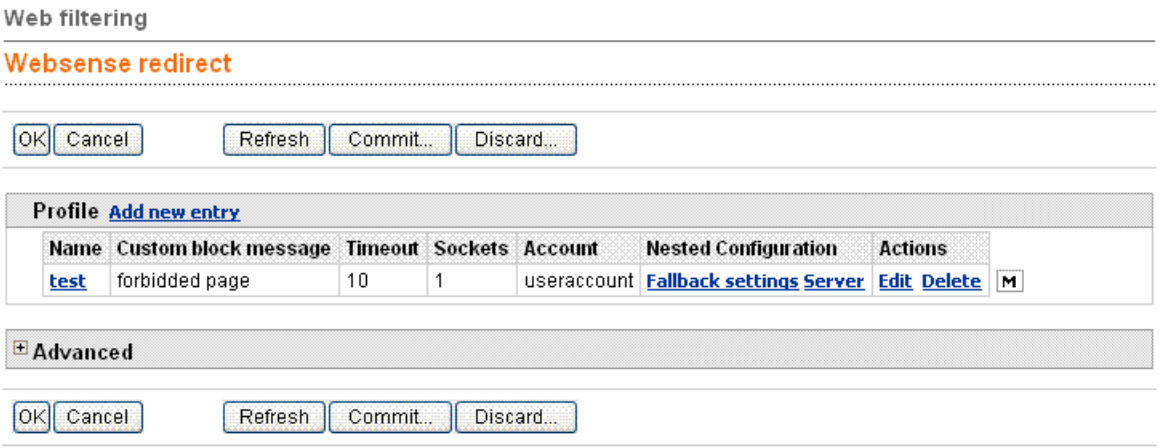

27. Click **OK** again to return to the Web filtering main page.

Next, you configure a UTM policy for Web filtering to which you attach the profile you have configured.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Utm, click **Configure**.
4. Next to Utm policy, click **Add new entry**.
5. In the Name box, enter a unique name for the UTM policy you are creating.
6. Next to Web filtering, click **Configure**.
7. In the Http profile box, enter the name of the Web filtering profile you created earlier.

8. Click **OK**.
9. Click **OK** again to return to main UTM configuration page. Your UTM Web filtering policy is now listed in the UTM policy table.

Next, you attach the UTM policy to a security policy that you create.

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policy, select the **Yes** check box and click **Configure** or **Edit**.
4. Next to Policy, click **Add new entry**.

NOTE: Refer to the section on security policy configuration for further details on configuring a policy. Note that when you configure the Then field as part of the policy, select Permit as the action, and then configure Application services, you are able to enter the Utm policy name as part of this security policy.

5. Next to Utm policy (in the Application services security policy window), enter the name of the appropriate policy. This attaches your UTM policy to the security policy.
6. Click **OK**.

CLI Configuration

To configure redirect Web filtering using the CLI, you must first create your custom objects.

NOTE: Rather than or in addition to custom object lists, you can use included default lists and included whitelist and blacklist categories. See "Web Filtering Profiles" on page 726 for profile list information.

1. Configure a URL pattern list custom object by creating the list name and adding values to it as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

```
user@host# set security utm custom-objects url-pattern urllist4 value
  [http://www.juniper.net 1.2.3.4]
```

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax IS supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

2. Configure a custom URL category list custom object by using the URL pattern list you created as follows:

   user@host# set security utm custom-objects custom-url-category **custurl4** value **urllist3**

1. Now that your custom objects have been created, you can configure the Websense-redirect Web filtering feature profile.

   If you are using included global whitelist and blacklist categories, select those global categories. This is the first filtering category that both integrated and redirect Web filtering use. If no match is made, the URL is forwarded to the Websense server. See "Understanding Integrated Web Filtering" on page 724 for URL category information.

   user@host# set security utm feature-profile web-filtering url-blacklist **Drugs_Alcohol_Tobacco**
   user@host# set security utm feature-profile web-filtering url-whitelist **Computing_Internet**

2. Select Websense-redirect as your Web filtering engine as follows:

   user@host# set security utm feature-profile web-filtering **websense-redirect**

3. Create a Websense-redirect profile by first creating a profile name as follows:

   user@host# set security utm feature-profile web-filtering websense-redirect profile **websenseprofile1**

4. Set Websense server parameters by entering the server name or IP address.

   user@host# set security utm feature-profile web-filtering websense-redirect profile **websenseprofile1** server host **Websenseserver**

5. Enter the port number for communicating with the Websense server. (Default ports are 80, 8080, and 8081.)

   user@host# set security utm feature-profile web-filtering websense-redirect profile **websenseprofile1** server port **8080**

6. Select fallback settings (block or log and permit) for this profile. The fallback actions are taken when errors in each category configured occur.

```
user@host# set security utm feature-profile web-filtering websense-redirect profile
  websenseprofile1 fallback-settings default block
user@host# set security utm feature-profile web-filtering websense-redirect profile
  websenseprofile1 fallback-settings server-connectivity block
user@host# set security utm feature-profile web-filtering websense-redirect profile
  websenseprofile1 fallback-settings timeout block
user@host# set security utm feature-profile web-filtering websense-redirect profile
  websenseprofile1 fallback-settings too-many-requests block
```

7. Enter the number of sockets used for communicating between the client and server. The default is 1.

```
user@host# set security utm feature-profile web-filtering websense-redirect profile
  websenseprofile1 sockets 1
```

8. Enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied. The default here is 10 seconds. You can enter a value from 10 to 240 seconds.

```
user@host# set security utm feature-profile web-filtering websense-redirect profile
  websenseprofile1 timeout 10
```

1. Next, configure a UTM policy for the Web filtering protocol and attach this policy to a profile. CLI commands for configuring a UTM policy for HTTP Web filtering, and attaching that policy to a profile you created earlier for Web filtering are:

```
user@host# set security utm utm-policy <name>
user@host# set security utm utm-policy utmp6 web-filtering http-profile
  websenseprofile1
```

2. Attach the UTM policy to a firewall security policy.

```
user@host# set security policies
user@host# set security policies from-zone trust to-zone untrust policy p6 then
  permit application-services utm-policy utmp6
```

Understanding Local Web Filtering

With local Web filtering, the firewall intercepts every HTTP request in a TCP connection and extracts the URL. The decision making is done on the device after it looks up a URL to determine if it is in the whitelist or blacklist based on its user-defined category. If the URL is in the url-blacklist, the request is blocked; if it's in the url-whitelist, the request is permitted. If the URL is not in either list, the defined default action will occur (block, log-and-permit, or permit). You can permit or block access to a requested site by binding a Web filtering profile to a firewall policy. Local Web filtering provides basic Web filtering without requiring an additional license or external category server.

User-Defined URL Categories

When defining your own URL categories, you can group URLs and create categories specific to your needs. Each category can have a maximum of 20 URLs. When you create a category, you can add either the URL or the IP address of a site. When you

add a URL to a user-defined category, the device performs DNS lookup, resolves the hostname into IP addresses, and caches this information. When a user tries to access a site with the IP address of the site, the device checks the cached list of IP addresses and tries to resolve the hostname. Many sites have dynamic IP addresses, meaning that their IP addresses change periodically. A user attempting to access a site can type an IP address that is not in the cached list on the device. Therefore, if you know the IP addresses of sites you are adding to a category, enter both the URL and the IP address(es) of the site.

You define your own categories using URL pattern list and custom URL category list custom objects. Once defined, you assign your categories to the global user-defined url-blacklist (block) or url-whitelist (permit) categories. You configure your Web filtering profile as follows:

```
user@host# set security utm feature-profile Web-filtering url-whitelist category <category
    name>
```

NOTE: In this example, the user-defined category is assigned to the global url-whitelist category. This will permit all URLs in that category.

Web filtering is performed on all the methods defined in HTTP 1.0 and HTTP 1.1.

Local Web Filtering: Process Overview

This is a general description of how Web traffic is intercepted and acted upon by the Web filtering module.

1. The device intercepts a TCP connection.
2. The device intercepts each HTTP request in the TCP connection.
3. The device extracts each URL in the HTTP request and checks its URL against the user-defined whitelist and blacklist.
4. If the URL is found in the blacklist, the request is not permitted and a deny page is sent to the http client. If the URL is found in the whitelist, the request is permitted.
5. If the URL is not found in the whitelist or blacklist, the configured default fallback action is applied. If no fallback action is defined, then the request is permitted.

Web Filtering Profiles

You configure Web filtering profiles that permit or block URLs according to defined custom categories. A Web filtering profile consists of a group of URL categories assigned one of the following actions:

- Blacklist — The device always blocks access to the websites in this list. Only user-defined categories are used with local Web filtering.
- Whitelist — The device always allows access to the websites in this list. Only user-defined categories are used with local Web filtering.

A Web filtering profile can contain one blacklist or one whitelist with multiple user-defined categories each with a permit or block action. You can define a default fallback action when the incoming URL does not belong to any of the categories defined in the profile. If the action for the default category is block, the incoming URL is blocked if it does not match any of the categories explicitly defined in the profile. If an action for the default action is not specified, the default action of permit is applied to the incoming URL not matching any category.

Profile Matching Precedence

When a profile employs several categories for URL matching, those categories are checked for matches in the following order:

1. If present, the global blacklist is checked first. If a match is made, the URL is blocked. If no match is found...
2. The global whitelist is checked next. If a match is made, the URL is permitted. If no match is found...
3. User-defined categories are checked next. If a match is made, the URL is blocked or permitted as specified.

Configuring Local Web Filtering

To configure local Web filtering using the CLI, you must first create your custom objects.

1. Configure a URL pattern list custom object by creating the list name and adding values to it as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists. The URL and IP address is added in this example.

```
user@host# set security utm custom-objects url-pattern urllist3 value
  [http://www.juniper.net 1.2.3.4]
user@host# set security utm custom-objects url-pattern urllist4 value
  [http://www.acmegizmo.com 1.2.3.4]
```

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax is supported: http://*.juniper.net, http://www.juniper.ne?, http://www.juniper.n??. The following wildcard syntax is NOT supported: *.juniper.net , www.juniper.ne?, http://*juniper.net, http://*.

1. Configure a custom URL category list custom object by using the URL pattern list you created as follows:

```
user@host# set security utm custom-objects custom-url-category custurl3 value
  urllist3
user@host# set security utm custom-objects custom-url-category custurl4 value
  urllist4
```

Now that your custom objects have been created, you can configure the juniper-local Web filtering feature profile.

1. If you are using included global whitelist and blacklist categories, select those global categories. This is the first filtering category that both integrated, redirect, and local Web filtering use. If no match is made, the configured default fallback action is performed. See "Understanding Local Web Filtering" on page 751 for URL category information.

```
user@host# set security utm feature-profile web-filtering url-blacklist custurl3
user@host# set security utm feature-profile web-filtering url-whitelist custurl4
```

2. Select juniper-local as your Web filtering engine as follows:

```
user@host# set security utm feature-profile web-filtering type juniper-local
```

3. Create a juniper-local profile by first creating a profile with a default action (permit, log and permit, block) for requests that experience errors as follows:

```
user@host# set security utm feature-profile web-filtering juniper–local profile
  localprofile1 default permit
```

4. Enter a custom message to be sent when HTTP requests are blocked.

```
user@host# set security utm feature-profile web-filtering juniper-local profile
  localprofile1 custom-block-message "Access to this site is not permitted"
```

5. Select fallback settings (block or log and permit) for this profile. The fallback actions are taken when errors in each configured category occur.

```
user@host# set security utm feature-profile web-filtering juniper–local profile
  localprofile1 fallback-settings default block
user@host# set security utm feature-profile web-filtering juniper–local profile
  localprofile1 fallback-settings too-many-requests block
```

1. Configure a UTM policy for the Web filtering protocol and attach this policy to a profile. CLI commands for configuring a UTM policy for HTTP Web filtering, and attaching that policy to a profile you created earlier for content filtering are:

```
user@host# set security utm utm-policy utmp5 web-filtering http-profile localprofile1
```

2. Attach the UTM policy to a firewall security policy.

```
user@host# set security policies from-zone trust to-zone untrust policy p5 match
  source-address any
user@host# set security policies from-zone trust to-zone untrust policy p5 match
  destination-address any
```

```
user@host# set security policies from-zone trust to-zone untrust policy p5 match
    application junos-http
user@host# set security policies from-zone trust to-zone untrust policy p5 then
    permit application-services utm-policy utmp5
```

Verifying Web Filtering Configurations

J-Web monitoring tools consist of the options that appear when you select Monitor in the task bar. The Monitor options display diagnostic information about the feature in question. Alternatively, you can enter show commands from the CLI to display the same information. CLI show commands display the current configuration and information about the feature in question.

This topic contains:

- Using J-Web to Monitor Web Filtering on page 755
- Using the CLI to Monitor Web Filtering on page 755

Using J-Web to Monitor Web Filtering

View Web filtering statistics using J-Web as follows:

1. Select **Monitor > UTM > Web Filtering**.

   The following information becomes viewable in the right pane.

   ```
   white list hit: #
   Black list hit: #
   Queries to server: #
   Server reply permit: #
   Server reply block: #
   Custom category permit: #
   Custom category block: #
   Cache hit permit: #
   Cache hit block: #
   Web-filtering sessions in total: #
   Web-filtering sessions in use: #
   Fall back: log-and-permit block
   Default # #
   Timeout # #
   Connectivity # #
   Too-many-requests # #
   ```

2. You can click the **Clear Web Filtering STAT** button to clear all current viewable statistics and begin collecting new statistics.

Using the CLI to Monitor Web Filtering

Use the following commands to view Web filtering statistics.

The Web filtering show status command displays the following information:

```
user@host > show security utm web-filtering status

UTM Web-filtering status:
Server status: SC-CPA server up
```

The Web filtering show statistics command displays the following information:

```
user@host > show security utm web-filtering statistics

white list hit: #
Black list hit: #
Queries to server: #
Server reply permit: #
Server reply block: #
Custom category permit: #
Custom category block: #
Cache hit permit: #
Cache hit block: #
Web-filtering sessions in total: #
Web-filtering sessions in use: #
Fall back: log-and-permit block
Default # #
Timeout # #
Connectivity # #
Too-many-requests # #
```

What Happens When Both Web Filtering and Antivirus Scanning Are Employed?

You can bind either Web filtering profiles or antivirus profiles, or both, to a firewall policy. When both are bound to a firewall policy, Web filtering is applied first, then antivirus is applied. If a URL is blocked by Web filtering, the TCP connection is closed and no antivirus scanning is necessary. If a URL is permitted, the content of the transaction is then passed to the antivirus scanning process.

WebTrends Enhanced Log File Format (WELF) for UTM Features

UTM features support the WELF standard. The WELF Reference defines the WebTrends industry standard log file exchange format. Any system logging to this format is compatible with Firewall Suite 2.0 and later, Firewall Reporting Center 1.0 and later, and Security Reporting Center 2.0 and later.

This section contains the following topics:

- WELF Overview on page 756
- CLI Configuration on page 757

WELF Overview

A WELF log file is made up of records. Each record makes up a single line of the file. Records are always in chronological order. The earliest record is the first record in the file; the most recent record is the last record in the file. The WebTrends Enhanced Log Format places no restrictions on log file names or log file rotation policies.

NOTE: Each WELF record is made up of fields. The record identifier field (id =) must be the first field in a record. All other fields can appear in any order.

The following is a sample WELF record:

```
id=firewall time="2000-2-4 12:01:01" fw=192.168.0.238 pri=6 rule=3 proto=http
src=192.168.0.23 dst=6.1.0.36 rg=www.webtrends.com/index.html op=GET result=0
rcvd=1426
```

The fields from the example WELF record include the following required elements (all other fields are optional):

- id (Record identifier)
- time (Date/time)
- fw (Firewall IP address or name)
- pri (Priority of the record)

CLI Configuration

1. From the set command, set the security log source IP address. You do this when you are exporting security logs.

NOTE: The WELF logging messages must be saved to a dedicated WebTrends server.

```
user@host# set security log source-address 1.2.3.4
```

2. Next, you set the name of the security log stream. In this case, you will create the name utm-welf.

```
user@host# set security log source-address 1.2.3.4 stream utm-welf
```

3. Next, set the format for the log messages. In this case, set it to welf.

```
user@host# set security log source-address 1.2.3.4 stream utm-welf format welf
```

4. Next, set the category of log messages that are sent. In this case, you want to send security messages. Selecting content-security here indicates UTM messages.

```
user@host# set security log source-address 1.2.3.4 stream utm-welf format welf
  category content-security
```

5. Next, set the severity level of log messages that are sent.

```
user@host# set security log source-address 1.2.3.4 stream utm-welf format welf
  category content-security severity critical emergency
```

6. Next, enter the host address of the dedicated WebTrends server to which the log messages are to be sent.

```
user@host# set security log source-address 1.2.3.4 stream utm-welf format welf
  category content-security severity critical emergency host 5.6.7.8
```

7. Your configuration takes effect when you commit your changes.

Part 9
Attack Detection and Prevention

■ Attack Detection and Prevention on page 761

Chapter 27

Attack Detection and Prevention

An Intrusion Detection and Prevention (IDP), also known as a *stateful firewall*, detects and prevents attacks in network traffic.

An exploit can be either an information-gathering probe or an attack to compromise, disable, or harm a network or network resource. In some cases, the distinction between the two objectives of an exploit can be unclear. For example, a barrage of TCP SYN segments might be an IP address sweep with the intent of triggering responses from active hosts, or it might be a SYN flood attack with the intent of overwhelming a network so that it can no longer function properly. Furthermore, because an attacker usually precedes an attack by performing reconnaissance on the target, we can consider information-gathering efforts as a precursor to an impending attack—that is, they constitute the first stage of an attack. Thus, the term *exploit* encompasses both reconnaissance and attack activities, and the distinction between the two is not always clear.

Juniper Networks provides various detection methods and defense mechanisms at the zone and policy levels to combat exploits at all stages of their execution:

- Screen options at the zone level
- Firewall policies at the inter-, intra-, and super-zone policy levels (*super-zone* here means in global policies, where no security zones are referenced).

To secure all connection attempts, JUNOS Softwareuses a dynamic packet-filtering method known as *stateful inspection*. Using this method, JUNOS Software identifies various components in the IP packet and TCP segment headers—source and destination IP addresses, source and destination port numbers, and packet sequence numbers—and maintains the state of each TCP session and pseudo UDP session traversing the firewall. (JUNOS Software also modifies session states based on changing elements such as dynamic port changes or session termination.) When a responding TCP packet arrives, JUNOS Software compares the information reported in its header with the state of its associated session stored in the inspection table. If they match, the responding packet is allowed to pass the firewall. If the two do not match, the packet is dropped.

JUNOS Software Screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone. JUNOS Software then applies firewall policies, which can contain content filtering and Intrusion Detection and Prevention (IDP) components, to the traffic that passes the Screen filters.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Reconnaissance Deterrence Overview on page 763
- Understanding IP Address Sweeps on page 764
- Blocking IP Address Sweeps on page 765
- Understanding Port Scanning on page 766
- Blocking Port Scans on page 768
- Understanding Network Reconnaissance Using IP Options on page 769
- Detecting Packets That Use IP Options for Reconnaissance on page 771
- Understanding Operating System Probes on page 773
- Blocking Packets with SYN and FIN Flags Set on page 776
- Blocking Packets with FIN Flag/No ACK Flag Set on page 777
- Blocking Packets with No Flags Set on page 778
- Understanding Attacker Evasion Techniques on page 779
- Thwarting a FIN Scan on page 784
- Setting TCP SYN Checking on page 785
- Setting Strict Syn Checking on page 786
- Blocking IP Spoofing on page 787
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790
- Suspicious Packet Attributes Overview on page 791
- Understanding ICMP Fragment Protection on page 792
- Blocking Fragmented ICMP Packets on page 793
- Understanding Large ICMP Packet Protection on page 794
- Blocking Large ICMP Packets on page 796
- Understanding Bad IP Option Protection on page 797
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798
- Understanding Unknown Protocol Protection on page 799
- Dropping Packets Using an Unknown Protocol on page 800
- Understanding IP Packet Fragment Protection on page 802
- Dropping Fragmented IP Packets on page 803
- Understanding SYN Fragment Protection on page 804
- Dropping IP Packets Containing SYN Fragments on page 805
- Denial-of-Service Attack Overview on page 807
- Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview on page 807

- Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview on page 808
- Enabling Protection Against Application-level Distributed Denial-of-Service (AppDDoS) Attack on page 812
- Firewall DoS Attacks Overview on page 815
- Understanding Session Table Flood Attacks on page 815
- Setting Source-Based Session Limits on page 818
- Setting Destination-Based Session Limits on page 819
- Understanding SYN-ACK-ACK Proxy Flood Attacks on page 821
- Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack on page 821
- Network DoS Attacks Overview on page 823
- Understanding SYN Flood Attacks on page 823
- Example: SYN Flood Protection on page 829
- Enabling SYN Flood Protection on page 835
- Understanding SYN Cookie Protection on page 836
- Enabling SYN Cookie Protection on page 837
- Understanding ICMP Flood Attacks on page 839
- Enabling ICMP Flood Protection on page 840
- Understanding UDP Flood Attacks on page 842
- Enabling UDP Flood Protection on page 843
- Understanding Land Attacks on page 844
- Enabling Protection Against a Land Attack on page 845
- OS-Specific DoS Attacks Overview on page 847
- Understanding Ping of Death Attacks on page 847
- Enabling Protection Against a Ping of Death Attack on page 848
- Understanding Teardrop Attacks on page 849
- Enabling Protection Against a Teardrop Attack on page 851
- Understanding WinNuke Attacks on page 852
- Enabling Protection Against a WinNuke Attack on page 853
- Configuring Firewall Screen Options—Quick Configuration on page 855
- Verifying Application Security Information Using Trace Options on page 860

Reconnaissance Deterrence Overview

Attackers can better plan their attack when they first know the layout of the targeted network (which IP addresses have active hosts), the possible entry points (which port numbers are active on the active hosts), and the constitution of their victims (which

operating system the active hosts are running). To gain this information, attackers must perform reconnaissance.

Before You Begin

For background information, read “Attack Detection and Prevention” on page 761.

Juniper Networks provides several screen options to deter attackers' reconnaissance efforts and thereby hinder them from obtaining valuable information about the protected network and network resources.

Related Topics

- Understanding IP Address Sweeps on page 764
- Understanding Port Scanning on page 766
- Understanding Network Reconnaissance Using IP Options on page 769
- Understanding Operating System Probes on page 773
- Understanding Attacker Evasion Techniques on page 779

Understanding IP Address Sweeps

An address sweep occurs when one source IP address sends a defined number of ICMP packets to different hosts within a defined interval (5000 microseconds is the default). The purpose of this attack is to send ICMP packets—typically echo requests—to various hosts in the hopes that at least one replies, thus uncovering an address to target.

Before You Begin

For background information, read:

- Attack Detection and Prevention on page 761
- Reconnaissance Deterrence Overview on page 763

The Juniper Networks device internally logs the number of ICMP packets to different addresses from one remote source. Using the default settings, if a remote host sends ICMP traffic to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as an address sweep attack and rejects the 11th and all further ICMP packets from that host for the remainder of the specified threshold time period. See Figure 126 on page 765.

Figure 126: Address Sweep

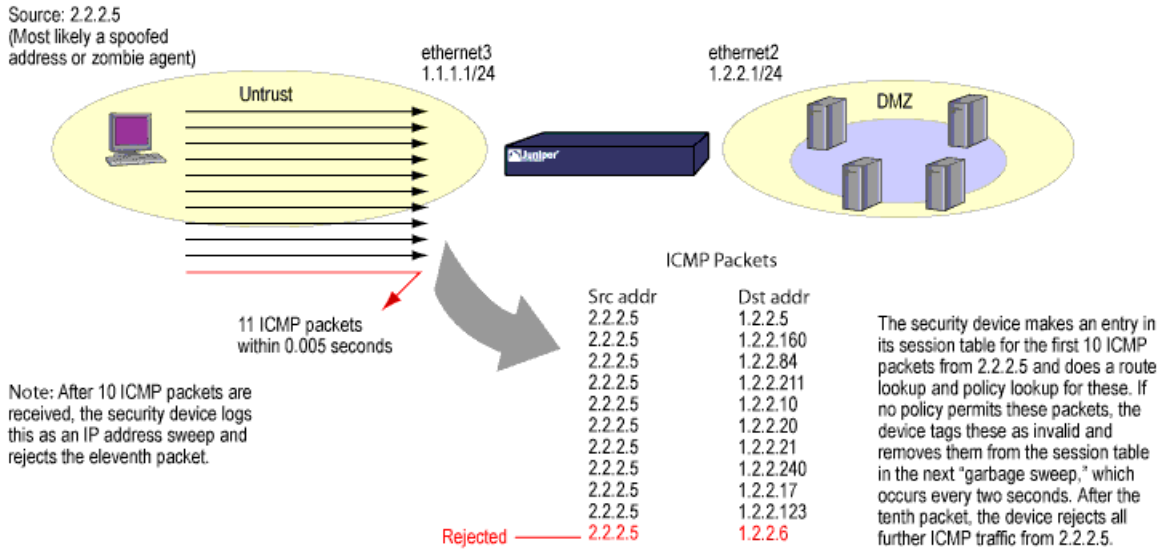

Consider enabling this screen option for a security zone only if there is a policy permitting ICMP traffic from that zone. Otherwise, you do not need to enable it. The lack of such a policy denies all ICMP traffic from that zone, precluding an attacker from successfully performing an IP address sweep anyway.

Related Topics

- Blocking IP Address Sweeps on page 765

Blocking IP Address Sweeps

Before You Begin
For background information, read "Understanding IP Address Sweeps" on page 764.

You can use either J-Web or the CLI configuration editor to block IP address sweeps originating in a particular security zone.

This topic covers:

- J-Web Configuration on page 765
- CLI Configuration on page 766
- Related Topics on page 766

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Edit**.
4. Next to Ids option, click **Add new entry**.
5. In the Name box, type **5000-ip-sweep**.
6. Next to Icmp, click **Configure**.
7. Next to Ip sweep, select the check box and click **Configure**.
8. In the Threshold box, type **5000** and click **OK**.
9. To save and commit the configuration, click **Commit**.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Zones, click **Edit**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type **zone**.
6. In the Screen box, type **5000-ip-sweep** and click **OK**.
7. To save and commit the configuration, click **Commit**.

CLI Configuration

```
user@host# set security screen ids-option 5000-ip-sweep icmp ip-sweep threshold
  5000
user@host# set security zones security-zone zone screen 5000-ip-sweep
```

NOTE: The value unit is in microseconds. The default value is 5000 microseconds.

Related Topics

- Attack Detection and Prevention on page 761
- Understanding Port Scanning on page 766

Understanding Port Scanning

A port scan occurs when one source IP address sends IP packets containing TCP SYN segments to a defined number of different ports at the same destination IP address within a defined interval (5000 microseconds is the default). The purpose of this

attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.

Before You Begin

For background information, read:

- Attack Detection and Prevention on page 761
- Reconnaissance Deterrence Overview on page 763

JUNOS Software internally logs the number of different ports scanned from one remote source. Using the default settings, if a remote host scans 10 ports in 0.005 seconds (5000 microseconds), then the device flags this as a port scan attack and rejects all further packets from the remote source, regardless of the destination IP address, for the remainder of the specified timeout period. See Figure 127 on page 767.

Figure 127: Port Scan

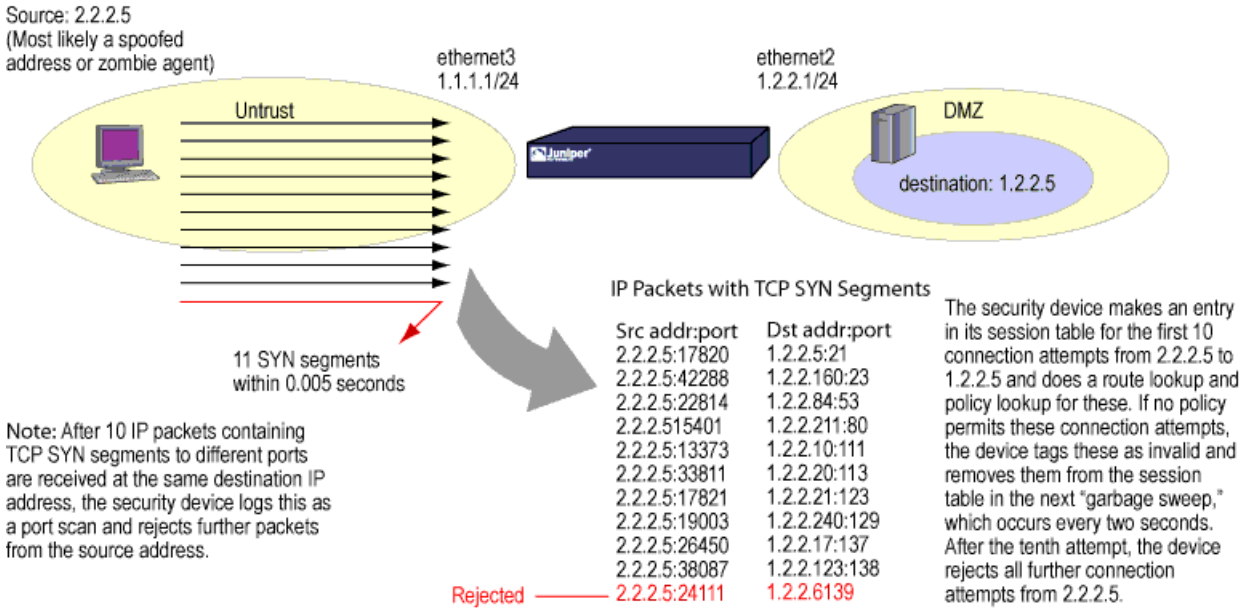

Related Topics

- Blocking Port Scans on page 768
- Blocking IP Address Sweeps on page 765

Blocking Port Scans

Before You Begin
For background information, read "Understanding Port Scanning" on page 766.

You can use either J-Web or the CLI configuration editor to block port scans originating in a particular security zone.

This topic covers:

- J-Web Configuration on page 768
- CLI Configuration on page 768
- Related Topics on page 768

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Edit**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **5000-port-scan**.
6. Next to Tcp, click Configure.
7. Next to Port scan, select the check box and click Configure.
8. In the Threshold box, type **5000** and click OK.
9. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option 5000-port-scan tcp port-scan threshold
    5000
```

NOTE: The value unit is in microseconds. The default value is 5000 microseconds.

Related Topics

- Attack Detection and Prevention on page 761
- Reconnaissance Deterrence Overview on page 763

- Understanding Port Scanning on page 766

Understanding Network Reconnaissance Using IP Options

The Internet Protocol standard RFC 791, *Internet Protocol*, specifies a set of options to provide special routing controls, diagnostic tools, and security.

Before You Begin

For background information, read:

- Attack Detection and Prevention on page 761
- Reconnaissance Deterrence Overview on page 763

RFC 791 states that these options are " unnecessary for the most common communications" and, in reality, they rarely appear in IP packet headers. These options appear after the destination address in an IP packet header, as shown in Figure 128 on page 769. When they do appear, they are frequently being put to some illegitimate use.

Figure 128: Routing Options

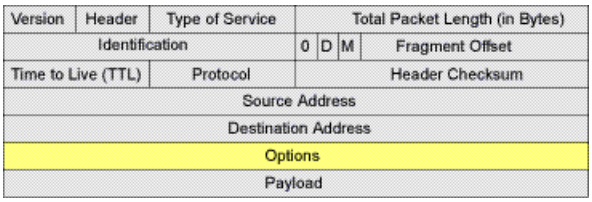

This topic covers:

- Uses for IP Packet Header Options on page 769
- Screen Options for Detecting IP Options Used For Reconnaissance on page 771
- Related Topics on page 771

Uses for IP Packet Header Options

Table 100 on page 769 lists the IP options and their accompanying attributes.

Table 100: IP Options and Attributes

Type	Class	Number	Length	Intended Use	Nefarious Use
End of Options	0*	0	0	Indicates the end of one or more IP options.	None.
No Options	0	1	0	Indicates there are no IP options in the header.	None.

Table 100: IP Options and Attributes *(continued)*

Type	Class	Number	Length	Intended Use	Nefarious Use
Security	0	2	11 bits	Provides a way for hosts to send security, TCC (closed user group) parameters, and Handling Restriction Codes compatible with Department of Defense (DoD) requirements. (This option, as specified in RFC 791, *Internet Protocol,* and RFC 1038, *Revised IP Security Option,* is obsolete.)	Unknown. However, because it is obsolete, its presence in an IP header is suspect.
Loose Source Route	0	3	Varies	Specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.	Evasion. The attacker can use the specified routes to hide the true source of a packet or to gain access to a protected network. (See "Blocking Packets with Either a Loose or Strict Source Route Option Set" on page 789.)
Record Route	0	7	Varies	Records the IP addresses of the network devices along the path that the IP packet travels. The destination machine can then extract and process the route information. (Due to the size limitation of 40 bytes for both the option and storage space, this can only record up to 9 IP addresses.)	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet passed.
Stream ID	0	8	4 bits	(Obsolete) Provided a way for the 16-bit SATNET stream identifier to be carried through networks that did not support the stream concept.	Unknown. However, because it is obsolete, its presence in an IP header is suspect.
Strict Source Route	0	9	Varies	Specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field.	Evasion. An attacker can use the specified routes to hide the true source of a packet or to gain access to a protected network. (See "Blocking Packets with Either a Loose or Strict Source Route Option Set" on page 789.)
Timestamp	2**	4		Records the time (in Universal Time***) when each network device receives the packet during its trip from the point of origin to its destination. The network devices are identified by IP number. This option develops a list of IP addresses of the devices along the path of the packet and the duration of transmission between each one.	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet passed.

Table 100: IP Options and Attributes *(continued)*

Type	Class	Number	Length	Intended Use	Nefarious Use
* The class of options identified as " 0" was designed to provide extra packet or network control.					
** The class of options identified as " 2" was designed for diagnostics, debugging, and measurement.					
*** The timestamp uses the number of milliseconds since midnight Universal Time (UT). UT is also known as Greenwich Mean Time (GMT), which is the basis for the international time standard.					

Screen Options for Detecting IP Options Used For Reconnaissance

The following screen options detect IP options that an attacker can use for reconnaissance or for some unknown but suspect purpose:

- Record Route:JUNOS Software detects packets where the IP option is 7 (Record Route) and records the event in the screen counters list for the ingress interface.
- Timestamp:JUNOS Software detects packets where the IP option list includes option 4 (Internet Timestamp) and records the event in the screen counters list for the ingress interface.
- Security: JUNOS Software detects packets where the IP option is 2 (Security) and records the event in the screen counters list for the ingress interface.
- Stream ID:JUNOS Software detects packets where the IP option is 8 (Stream ID) and records the event in the screen counters list for the ingress interface.

If a packet with any of the previous IP options is received, JUNOS Software flags this as a network reconnaissance attack and records the event for the ingress interface.

Related Topics

- Detecting Packets That Use IP Options for Reconnaissance on page 771

Detecting Packets That Use IP Options for Reconnaissance

Attackers can use the record route, timestamp, security, and stream ID IP options for reconnaissance or for some unknown but suspect purpose. To detect packets with these options set, you can use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 771
- CLI Configuration on page 773

J-Web Configuration

To configure screen and assign an Internet Protocol (IP) to it:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Edit**.
4. Next to Ids option, click **Add new entry**.
5. In the Name box, type **ip-record-route**.
6. Next to Ip, click **Configure**.
7. Next to Record route option, select the check box and click **OK**.
8. Next to Ids option, click **Add new entry**.
9. In the Name box, type **ip-timestamp-opt**.
10. Next to Ip, click **Edit**.
11. Next to Timestamp option, select the check box and click **OK**.
12. Next to Ids option, click **Add new entry**.
13. In the Name box, type **ip-security-opt**.
14. Next to Ip, click **Edit**.
15. Next to Security option, select the check box and click OK.
16. In the Name box, type **ip-stream-opt**.
17. Next to Ip, click **Edit**.
18. Next to Stream option, select the check box and click **OK**.
19. To save and commit the configuration, click **Commit**.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click **Add new entry**.
5. In the Name box, type **zone**.
6. In the Screen box, type **ip-record-route-opt** and click OK.
7. Next to Security zone, click **Add new entry**.
8. In the Name box, type **zone**.
9. In the Screen box, type **ip-timestamp-opt** and click **OK**.
10. Next to Security zone, click **Add new entry**.
11. In the Name box, type **zone**.
12. In the Screen box, type **ip-security-opt** and click **OK**.
13. Next to Security zone, click **Add new entry**.
14. In the Name box, type **zone**.

15. In the Screen box, type **ip-stream-opt** and click **OK**.
16. To save and commit the configuration, click **Commit**.

CLI Configuration

```
user@host# set security screen ids-option ip-record-route ip record-route-option
user@host# set security screen ids-option ip-timestamp-opt ip timestamp-option
user@host# set security screen ids-option ip-security-opt ip security-option
user@host# set security screen ids-option ip-stream-opt ip stream-option
user@host# set security zones security-zone zone screen ip-record-route-opt
user@host# set security zones security-zone zone screen ip-timestamp-opt
user@host# set security zones security-zone zone screen ip-security-opt
user@host# set security zones security-zone zone screen ip-stream-opt
```

Understanding Operating System Probes

Before launching an exploit, attackers might try to probe the targeted host to learn its operating system (OS). With that knowledge, they can better decide which attack to launch and which vulnerabilities to exploit. JUNOS Software can block reconnaissance probes commonly used to gather information about OS types.

Before You Begin

For background information, read:

- Attack Detection and Prevention on page 761
- Reconnaissance Deterrence Overview on page 763
- Understanding Network Reconnaissance Using IP Options on page 769

This topic covers:

- TCP Headers with SYN and FIN Flags Set on page 773
- TCP Headers With FIN Flag and Without ACK Flag on page 774
- TCP Header Without Flags Set on page 775
- Related Topics on page 775

TCP Headers with SYN and FIN Flags Set

Both the SYN and FIN control flags are not normally set in the same TCP segment header. The SYN flag synchronizes sequence numbers to initiate a TCP connection. The FIN flag indicates the end of data transmission to finish a TCP connection. Their purposes are mutually exclusive. A TCP header with the SYN and FIN flags set is anomalous TCP behavior, causing various responses from the recipient, depending on the OS. See Figure 129 on page 774.

Figure 129: TCP Header with SYN and FIN Flags Set

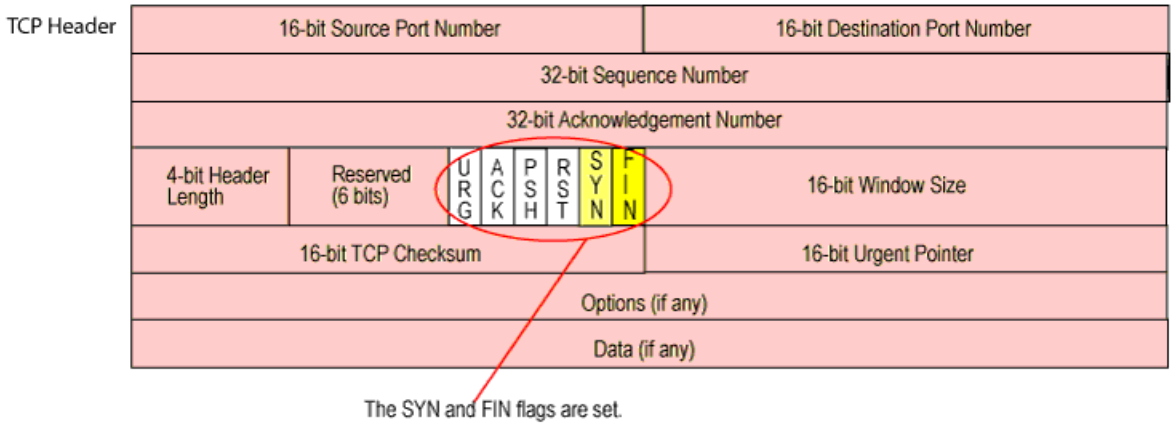

An attacker can send a segment with both flags set to see what kind of system reply is returned and thereby determine what kind of OS is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks.

When you enable this screen option, JUNOS Software checks if the SYN and FIN flags are set in TCP headers. If it discovers such a header, it drops the packet.

TCP Headers With FIN Flag and Without ACK Flag

Figure 130 on page 775 shows TCP segments with the FIN control flag set (to signal the conclusion of a session and terminate the connection). Normally, TCP segments with the FIN flag set also have the ACK flag set (to acknowledge the previous packet received). Because a TCP header with the FIN flag set but not the ACK flag is anomalous TCP behavior, there is no uniform response to this. The OS might respond by sending a TCP segment with the RST flag set. Another might completely ignore it. The victim's response can provide the attacker with a clue as to its OS.(Other purposes for sending a TCP segment with the FIN flag set are to evade detection while performing address and port scans and to evade defenses on guard for a SYN flood by performing a FIN flood instead. For information about FIN scans, see “FIN Scan” on page 780.)

NOTE: Vendors have interpreted RFC 793, *Transmission Control Protocol*, variously when designing their TCP/IP implementations. When a TCP segment arrives with the FIN flag set but not the ACK flag, some implementations send RST segments. Some drop the packet without sending an RST.

Figure 130: TCP Header with FIN Flag Set

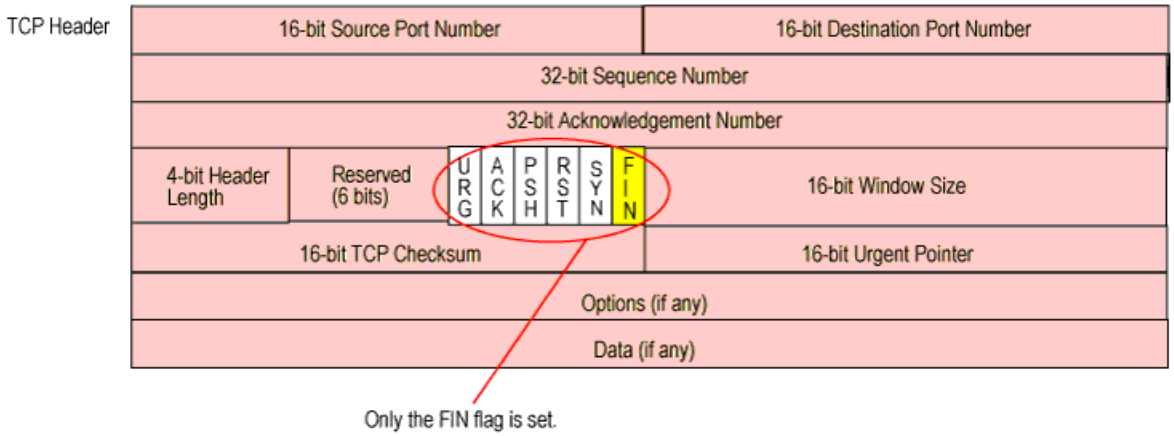

When you enable this screen option, JUNOS Software checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.

TCP Header Without Flags Set

A normal TCP segment header has at least one flag control set. A TCP segment with no control flags set is an anomalous event. Because different operating systems respond differently to such anomalies, the response (or lack of response) from the targeted device can provide a clue as to the type of OS it is running. See Figure 131 on page 775.

Figure 131: TCP Header with No Flags Set

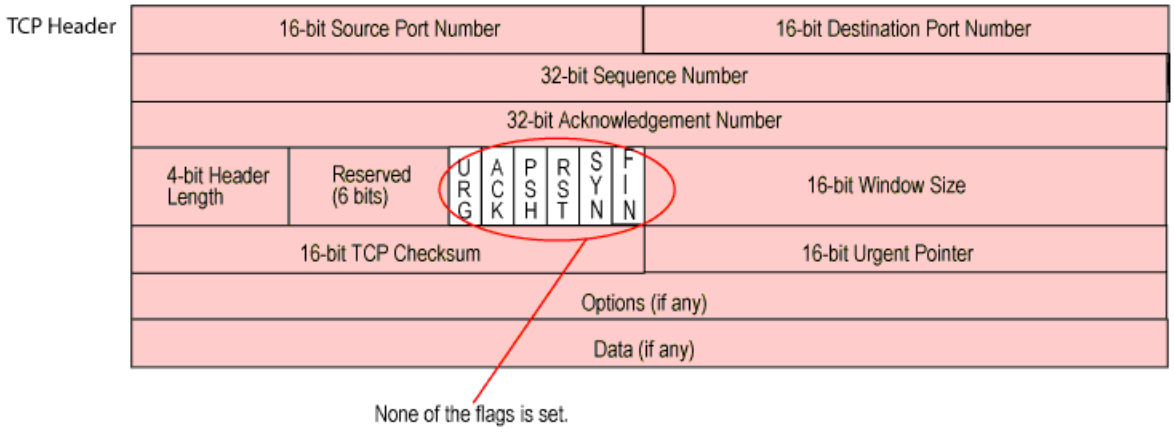

When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.

Related Topics

- Blocking Packets with SYN and FIN Flags Set on page 776

- Blocking Packets with FIN Flag/No ACK Flag Set on page 777
- Blocking Packets with SYN and FIN Flags Set on page 776

Blocking Packets with SYN and FIN Flags Set

A TCP header with the SYN and FIN flags set is anomalous TCP behavior, causing various responses from the recipient, depending on the OS. Blocking packets with SYN and FIN flags helps prevent OS system probes.

Before You Begin
For background information, read "Understanding Operating System Probes" on page 773.

You can use either J-Web or the CLI configuration editor to block packets with both the SYN and FIN flags set.

This topic covers:

- J-Web Configuration on page 776
- CLI Configuration on page 777
- Related Topics on page 777

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click **Add new entry**.
5. In the Name box, type **syn-fin**.
6. Next to Tcp, click **Configure**.
7. Next to Syn fin, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.

6. In the Screen box, type **syn-fin** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option syn-fin tcp syn-fin
user@host# set security zones security-zone zone screen syn-fin
```

Related Topics

- Blocking Packets with FIN Flag/No ACK Flag Set on page 777
- Blocking Packets with No Flags Set on page 778

Blocking Packets with FIN Flag/No ACK Flag Set

A TCP header with the FIN flag set but not the ACK flag is anomalous TCP behavior, causing various responses from the recipient, depending on the OS. Blocking packets with the FIN flag and without the ACK flag helps prevent OS system probes.

Before You Begin
For background information, read "Understanding Operating System Probes" on page 773.

You can use either J-Web or the CLI configuration editor to block packets with the FIN flag set but not the ACK flag.

This topic covers:

- J-Web Configuration on page 777
- CLI Configuration on page 778
- Related Topics on page 778

J-Web Configuration

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **screen**.
6. Next to Tcp, click Configure.
7. Next to Fin no ack, select the check box and click OK.
8. To save and commit the configuration, click Commit.

CLI Configuration

user@host# **set security screen ids-option <screen> tcp fin-no-ack**

Related Topics

- Blocking Packets with SYN and FIN Flags Set on page 776
- Blocking Packets with SYN and FIN Flags Set on page 776

Blocking Packets with No Flags Set

A TCP segment with no control flags set is an anomalous event, causing various responses from the recipient, depending on the OS. Blocking packets with no flags set helps prevent OS system probes. When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.

Before You Begin
For background information, read "Understanding Operating System Probes" on page 773.

You can use either J-Web or the CLI configuration editor to block packets with no flags set.

This topic covers:

- J-Web Configuration on page 778
- CLI Configuration on page 779
- Related Topics on page 779

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **tcp-no-flag**.
6. Next to Tcp, click Configure.
7. Next to tcp no flag, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **tcp-no-flag** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option tcp-no-flag tcp tcp-no-flag
user@host# set security zones security-zone zone screen tcp-no-flag
```

Related Topics

- Blocking Packets with SYN and FIN Flags Set on page 776
- Blocking Packets with FIN Flag/No ACK Flag Set on page 777

Understanding Attacker Evasion Techniques

Whether gathering information or launching an attack, it is generally expected that the attacker avoids detection. Although some IP address and port scans are blatant and easily detectable, more wily attackers use a variety of means to conceal their activity. Techniques such as using FIN scans instead of SYN scans—which attackers know most firewalls and intrusion detection programs detect—indicate an evolution of reconnaissance and exploit techniques to evade detection and successfully accomplish their tasks.

Before You Begin

For background information, read:

- Attack Detection and Prevention on page 761
- Reconnaissance Deterrence Overview on page 763
- Understanding Operating System Probes on page 773

This topic covers:

- FIN Scan on page 780
- Non-SYN Flags on page 780
- IP Spoofing on page 782

- IP Source Route Options on page 782
- Related Topics on page 784

FIN Scan

A FIN scan sends TCP segments with the FIN flag set in an attempt to provoke a response (a TCP segment with the RST flag set) and thereby discover an active host or an active port on a host. Attackers might use this approach rather than perform an address sweep with ICMP echo requests or an address scan with SYN segments because they know that many firewalls typically guard against the latter two approaches—but not necessarily against FIN segments. The use of TCP segments with the FIN flag set might evade detection and thereby help the attackers succeed in their reconnaissance efforts.

Non-SYN Flags

By default, JUNOS Software checks for SYN flags in the first packet of a session and rejects any TCP segments with non-SYN flags attempting to initiate a session. You can leave this packet flow as is or change it to so that JUNOS Software does not enforce SYN flag checking before creating a session. Figure 132 on page 780 illustrates packet flow sequences when SYN flag checking is enabled and when it is disabled.

Figure 132: SYN Flag Checking

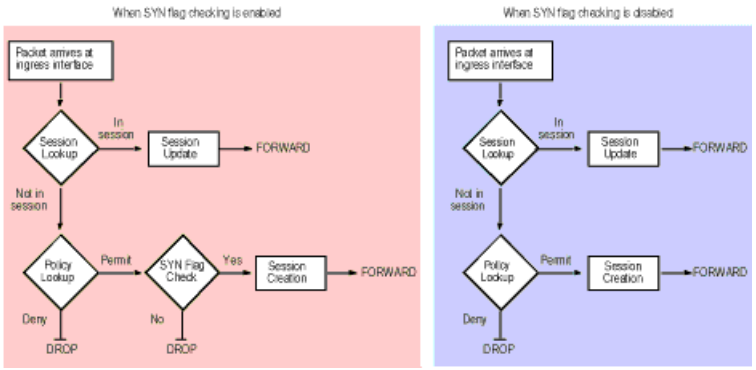

When JUNOS Software with SYN flag checking enabled receives a non-SYN TCP segment that does not belong to an existing session, it drops the packet and sends the source host to a TCP RST—unless the code bit of the initial non-SYN TCP packet is also RST. In that case, JUNOS Software simply drops the packet.

Not checking for the SYN flag in the first packets offers the following advantages:

- NSRP with Asymmetric Routing: In an active/active NSRP configuration in a dynamic routing environment, a host might send the initial TCP segment with the SYN flag set to one Juniper Networks device (Device-A) but the SYN/ACK might be routed to the other device in the cluster (Device-B). If this asymmetric routing occurs after Device-A has synchronized its session with Device-B, all is well. On the other hand, if the SYN/ACK response reaches Device-B before Device-A synchronizes the session and SYN checking is enabled, Device-B rejects the SYN/ACK, and the session cannot be established. With SYN checking disabled,

Device-B accepts the SYN/ACK response—even though there is no existing session to which it belongs—and creates a new session table entry for it.

- Uninterrupted Sessions: If you reset the device or even change a component in the core section of a policy and SYN checking is enabled, all existing sessions or those sessions to which the policy change applies are disrupted and must be restarted. Disabling SYN checking avoids such disruptions to network traffic flows.

NOTE: A solution to this scenario is to install the Juniper Networks device with SYN checking disabled initially. Then, after a few hours—when established sessions are running through the device—enable SYN checking. The core section in a policy contains the following main components: source and destination zones, source and destination addresses, one or more services, and an action.

However, the previous advantages exact the following security sacrifices:

- Reconnaissance Holes: When an initial TCP segment with a non-SYN flag—such as ACK, URG, RST, FIN—arrives at a closed port, many operating systems (Windows, for example) respond with a TCP segment that has the RST flag set. If the port is open, then the recipient does not generate any response.

  By analyzing these responses or lack thereof, an intelligence gatherer can perform reconnaissance on the protected network and also on the JUNOS Software policy set. If a TCP segment is sent with a non-SYN flag set and the policy permits it through, the destination host receiving such a segment might drop it and respond with a TCP segment that has the RST flag set. Such a response informs the perpetrator of the presence of an active host at a specific address and that the targeted port number is closed. The intelligence gatherer also learns that the firewall policy permits access to that port number on that host.

  By enabling SYN flag checking, JUNOS Software drops TCP segments without a SYN flag if they do not belong to an existing session. It does not return a TCP RST segment. Consequently, the scanner gets no replies regardless of the policy set or whether the port is open or closed on the targeted host.

- Session Table Floods: If SYN checking is disabled, an attacker can bypass the JUNOS Software SYN flood protection feature by flooding a protected network with a barrage of TCP segments that have non-SYN flags set. Although the targeted hosts drop the packets—and possibly send TCP RST segments in reply—such a flood can fill up the session table of the Juniper Networks device. When the session table is full, the device cannot process new sessions for legitimate traffic.

  By enabling SYN checking and SYN flood protection, you can thwart this kind of attack. Checking that the SYN flag is set on the initial packet in a session forces all new sessions to begin with a TCP segment that has the SYN flag set. SYN flood protection then limits the number of TCP SYN segments per second so that the session table does not become overwhelmed.

If you do not need SYN checking disabled, Juniper Networks strongly recommends that it be enabled (its default state for an initial installation of JUNOS Software). You

can enable it with the set flow tcp-syn-check command. With SYN checking enabled, the device rejects TCP segments with non-SYN flags set unless they belong to an established session.

IP Spoofing

One method of attempting to gain access to a restricted area of the network is to insert a bogus source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing. The mechanism to detect IP spoofing relies on route table entries. For example, if a packet with source IP address 10.1.1.6 arrives at ge-0/0/1, but JUNOS Software has a route to 10.1.1.0/24 through ge-0/0/0, a check for IP spoofing discovers that this address arrived at an invalid interface as defined in the route table. A valid packet from 10.1.1.6 can only arrive via ge-0/0/0, not ge-0/0/1. Therefore, JUNOS Software concludes that the packet has a spoofed source IP address and discards it.

IP Source Route Options

Source routing was designed to allow users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as " hops") along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis. If, for example, the transmission of a packet to a particular destination meets with irregular success, you might first use either the record route or timestamp IP option to discover the addresses of devices along the path or paths that the packet takes. You can then use either the loose or strict source route option to direct traffic along a specific path, using the addresses you learned from the results that the record route or timestamp options produced. By changing device addresses to alter the path and sending several packets along different paths, you can note changes that either improve or lessen the success rate. Through analysis and the process of elimination, you might be able to deduce where the trouble lies. See Figure 133 on page 783.

Figure 133: IP Source Routing

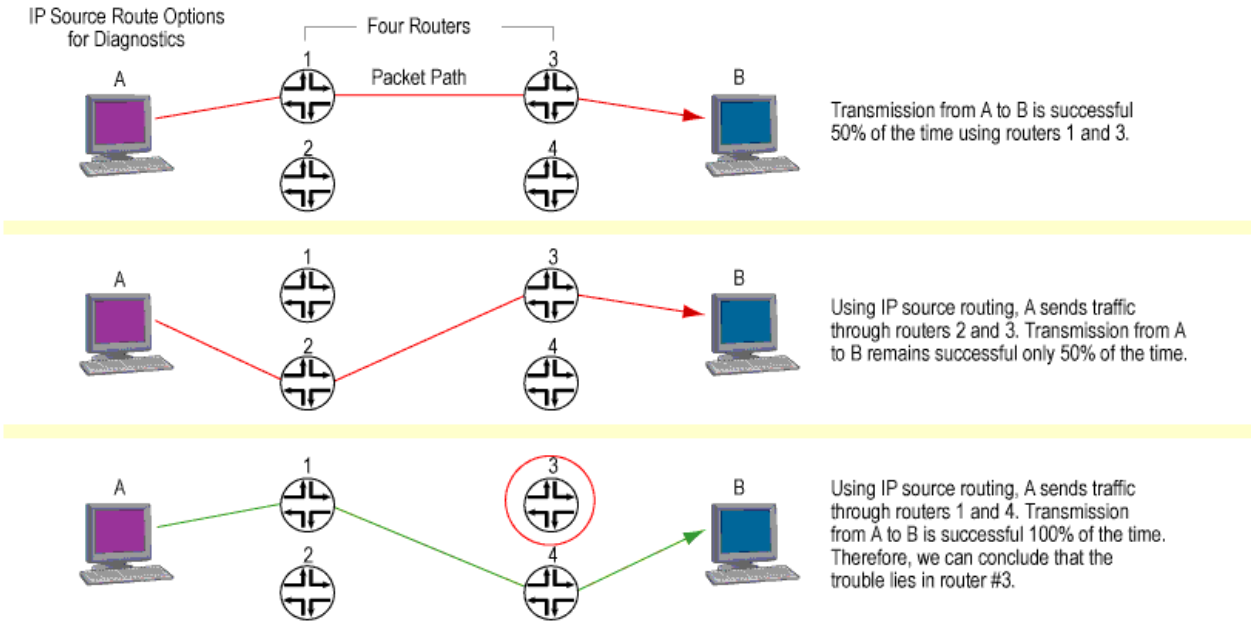

Although the uses of IP source route options were originally benign, attackers have learned to put them to more devious uses. They can use IP source route options to hide their true address and access restricted areas of a network by specifying a different path. For an example showing how an attacker can put both deceptions to use, consider the following scenario as illustrated in Figure 134 on page 783.

Figure 134: Loose IP Source Route Option for Deception

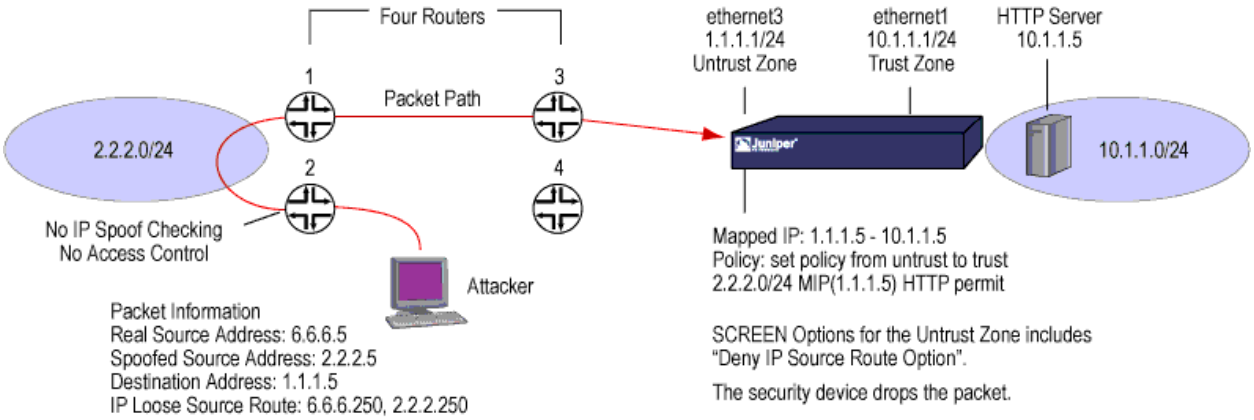

JUNOS Software only allows traffic 2.2.2.0/24 if it comes through ethernet1, an interface bound to zone_external. Devices 3 and 4 enforce access controls but devices 1 and 2 do not. Furthermore, device 2 does not check for IP spoofing. The attacker spoofs the source address, and by using the loose source route option, directs the packet through device 2 to the 2.2.2.0/24 network and from there out device 1. Device 1 forwards it to device 3, which forwards it to the Juniper Networks device. Because the packet came from the 2.2.2.0/24 subnet and has a source address from

that subnet, it seems to be valid. However, one remnant of the earlier chicanery remains: the loose source route option. In this example, you have enabled the deny IP source route screen option for the zone_external. When the packet arrives at ethernet3, the device rejects it.

You can enable the Juniper networks device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface. The screen options are as follows:

- Deny IP Source Route Option: Enable this option to block all IP traffic that employs the loose or strict source route option. Source route options can allow an attacker to enter a network with a false IP address.
- Detect IP Loose Source Route Option: The Juniper Networks device detects packets where the IP option is 3 (Loose Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.
- Detect IP Strict Source Route Option: The Juniper Networks device detects packets where the IP option is 9 (Strict Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field.

Related Topics

- Thwarting a FIN Scan on page 784
- Setting TCP SYN Checking on page 785
- Blocking IP Spoofing on page 787
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790

Thwarting a FIN Scan

A FIN scan sends TCP segments with the FIN flag set in an attempt to provoke a response (a TCP segment with the RST flag set) and thereby discover an active host or an active port on a host. The use of TCP segments with the FIN flag set might evade detection and thereby help the attacker succeed in his or her reconnaissance efforts.

Before You Begin
For background information, read “Understanding Attacker Evasion Techniques” on page 779.

To thwart FIN scans, use the JUNOS CLI configuration editor to take either or both of the following actions.

- Enable the screen option that specifically blocks TCP segments with the FIN flag set but not the ACK flag, which is anomalous for a TCP segment. Enter the following commands:

  ```
  user@host# set security screen fin-no-ack tcp fin-no-ack
  user@host# set security zones security-zone name screen fin-no-ack
  ```

  where *name* is the name of the zone to which you want to apply this screen option

- Change the packet processing behavior to reject all non-SYN packets that do not belong to an existing session. The SYN check flag is set as the default.

NOTE: Changing the packet flow to check that the SYN flag is set for packets that do not belong to existing sessions also thwarts other types of non-SYN scans, such as a null scan (when no TCP flags are set).

Related Topics

- Setting TCP SYN Checking on page 785
- Blocking IP Spoofing on page 787
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790

Setting TCP SYN Checking

With SYN checking enabled, the Juniper Networks device rejects TCP segments with non-SYN flags set unless they belong to an established session. Enabling SYN checking can help prevent attacker reconnaissance and session table floods.

Before You Begin

For background information, read "Understanding Attacker Evasion Techniques" on page 779.

You can use either J-Web or the CLI configuration editor to disable SYN checking. TCP SYN checking is on by default.

This topic covers:

- J-Web Configuration on page 786
- CLI Configuration on page 786
- Related Topics on page 786

J-Web Configuration

To disable SYN checking:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click **Configure** or **Edit**.
3. Next to flow, click Configure.
4. Next to Tcp session, click Configure.
5. Next to No syn check, select the check box and click OK.
6. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security flow tcp-session no-syn-check
```

Related Topics

- Thwarting a FIN Scan on page 784
- Blocking IP Spoofing on page 787
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790

Setting Strict Syn Checking

With strict SYN checking enabled, the Juniper Networks device enables the strict three-way handshake check for the TCP session. It enhances security by dropping data packets before the three-way handshake is done. By default, strict-syn-check is disabled.

Before You Begin
For background information, read “Understanding Attacker Evasion Techniques” on page 779.

You can use either J-Web or the CLI configuration editor to configure strict SYN checking. TCP strict SYN checking is disabled by default.

NOTE: strict-syn-check cannot be enabled if no-syn-check or no-syn-check-in-tunnel is enabled.

This topic covers:

- J-Web Configuration on page 787
- CLI Configuration on page 787
- Related Topics on page 787

J-Web Configuration

To enable strict SYN checking:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click **Configure** or **Edit**.
3. Next to flow, click Configure.
4. Next to Tcp session, click Configure.
5. Next to Strict syn check, select the check box and click OK.
6. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security flow tcp-session strict-syn-check
```

Related Topics

- Setting TCP SYN Checking on page 785
- Thwarting a FIN Scan on page 784
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790

Blocking IP Spoofing

One method of attempting to gain access to a restricted area of the network is to insert a bogus source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing.

Before You Begin
For background information, read "Understanding Attacker Evasion Techniques" on page 779.

You can use either J-Web or the CLI configuration editor to block *IP spoofing*.

This topic covers:

- J-Web Configuration on page 788
- CLI Configuration on page 788
- Related Topics on page 788

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **Ip-spoofing**.
6. Next to Ip, click Configure.
7. Next to Spoofing, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **ip-spoofing** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option ip-spoofing ip spoofing
user@host# set security zones security-zone zone screen ip-spoofing
```

Related Topics

- Thwarting a FIN Scan on page 784
- Setting TCP SYN Checking on page 785

- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790

Blocking Packets with Either a Loose or Strict Source Route Option Set

Checking for SYN flags can prevent attackers from using IP source route options to hide their true address and access restricted areas of a network by specifying a different path. TCP SYN checking is on by default.

Before You Begin

For background information, read "Understanding Attacker Evasion Techniques" on page 779.

You can use either J-Web or the CLI configuration editor to block packets with either a loose or strict source route option set. The specified security zone is the one from which the packets originated.

This topic covers:

- J-Web Configuration on page 789
- CLI Configuration on page 790
- Related Topics on page 790

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **ip-filter-src**.
6. Next to Ip, click Configure.
7. Next to Source route option , select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Zones, click **Edit**.
4. Next to Security zone, click Add new entry.

5. In the Name box, type **zone**.
6. In the Screen box, type **ip-filter-src** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option ip-filter-src ip source-route-option
user@host# set security zones security-zone zone screen ip-filter-src
```

Related Topics

- Thwarting a FIN Scan on page 784
- Setting TCP SYN Checking on page 785
- Blocking IP Spoofing on page 787
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 790

Detecting Packets with Either a Loose or Strict Source Route Option Set

Checking for SYN flags can also prevent attackers from using IP source route options to hide their true address and access restricted areas of a network by specifying a different path. TCP SYN checking is on by default.

Before You Begin
For background information, read "Understanding Attacker Evasion Techniques" on page 779.

You can use either J-Web or the CLI configuration editor to detect and record, but not block, packets with a loose or strict source route option set.

This topic covers:

- J-Web Configuration on page 790
- CLI Configuration on page 791
- Related Topics on page 791

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Screen, click **Edit**.
4. Next to Ids option, click Add new entry.

5. In the Name box, type **ip-loose-src-route**.
6. Next to Ip, click Configure.
7. Next to loose source route option, select the check box and click OK.
8. Next to Ip, click Edit.
9. Next to ip strict src route, select the check box and click OK.
10. To save and commit the configuration, click Commit.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Zones, click **Edit**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **ip-strict-src-route** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option ip-loose-src-route ip
  loose-source-route-option
user@host# set security screen ids-option ip-strict-src-route ip
  strict-source-route-option
user@host# set security zones security-zone zone screen ip-strict-src-route
```

Related Topics

- Thwarting a FIN Scan on page 784
- Setting TCP SYN Checking on page 785
- Blocking IP Spoofing on page 787
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 789

Suspicious Packet Attributes Overview

Attackers can craft packets to perform reconnaissance or launch denial-of-service (DoS) attacks. Sometimes it is unclear what the intent of a crafted packet is, but the very fact that it is crafted suggests that its being put to some kind of insidious use.

Before You Begin
For background information, read "Attack Detection and Prevention" on page 761.

All of the screen options listed in “Related Topics” on page 792 block suspicious packets that might contain hidden threats.

Related Topics

- Understanding ICMP Fragment Protection on page 792
- Understanding Large ICMP Packet Protection on page 794
- Understanding Bad IP Option Protection on page 797
- Understanding Unknown Protocol Protection on page 799
- Understanding IP Packet Fragment Protection on page 802
- Understanding SYN Fragment Protection on page 804

Understanding ICMP Fragment Protection

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.

Before You Begin
For background information, read “Suspicious Packet Attributes Overview” on page 791.

When you enable the ICMP fragment protection screen option, JUNOS Software blocks any ICMP packet that has the More Fragments flag set or that has an offset value indicated in the offset field. See Figure 135 on page 793.

Figure 135: Blocking ICMP Fragments

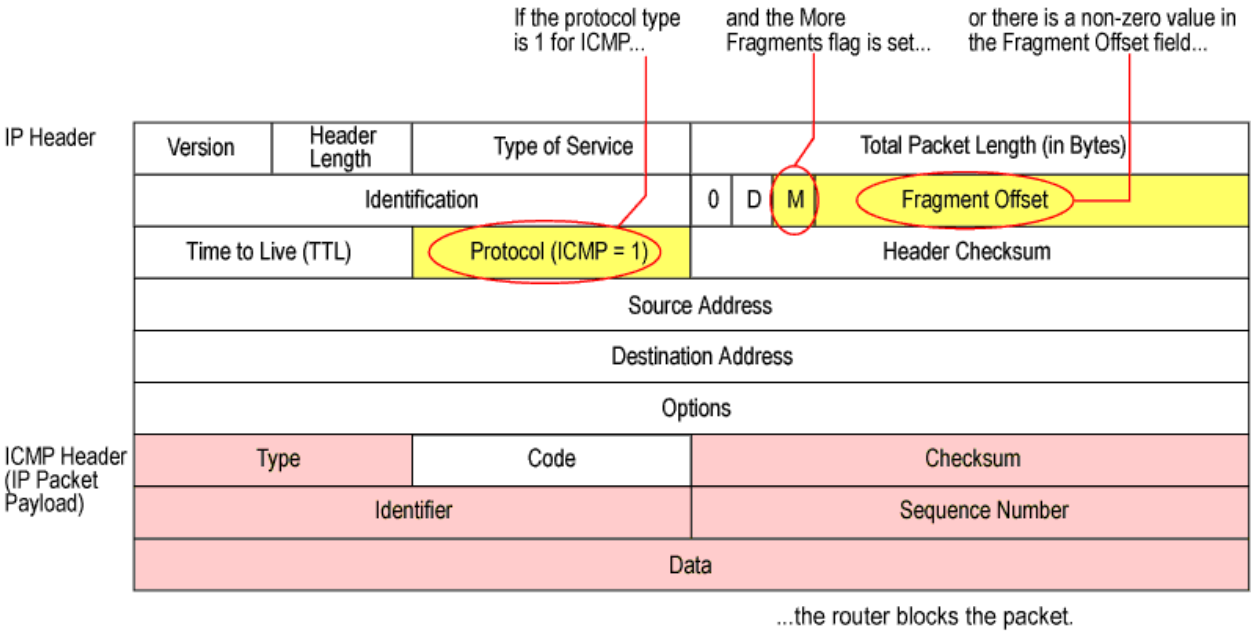

Related Topics

- Blocking Fragmented ICMP Packets on page 793
- Understanding Large ICMP Packet Protection on page 794

Blocking Fragmented ICMP Packets

Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.

Before You Begin
For background information, read “Understanding ICMP Fragment Protection” on page 792.

You can use either J-Web or the CLI configuration editor to block fragmented ICMP packets. The specified security zone is the one from which the fragments originated.

This topic covers:

- J-Web Configuration on page 794
- CLI Configuration on page 794
- Related Topics on page 794

J-Web Configuration

To configure screens:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Screen, click **Edit**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **icmp-fragment**.
6. Next to Icmp, click Configure.
7. Next to Fragment, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **CLI Tools > Point and Click CLI**.
2. Next to Security, click **Edit**.
3. Next to Zones, click **Edit**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **icmp-fragment** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option icmp-fragment icmp fragment
user@host# set security zones security-zone zone screen icmp-fragment
```

Related Topics

- Blocking Large ICMP Packets on page 796
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798
- Dropping Packets Using an Unknown Protocol on page 800
- Dropping Fragmented IP Packets on page 803
- Dropping IP Packets Containing SYN Fragments on page 805

Understanding Large ICMP Packet Protection

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate

reason for large ICMP packets. If an ICMP packet is unusually large, something is wrong.

Before You Begin

For background information, read “Suspicious Packet Attributes Overview” on page 791.

For example, the Loki program uses ICMP as a channel for transmitting covert messages. The presence of large ICMP packets might expose a compromised machine acting as a Loki agent. It also might indicate some other kind of questionable activity. See Figure 136 on page 795.

Figure 136: Blocking Large ICMP Packets

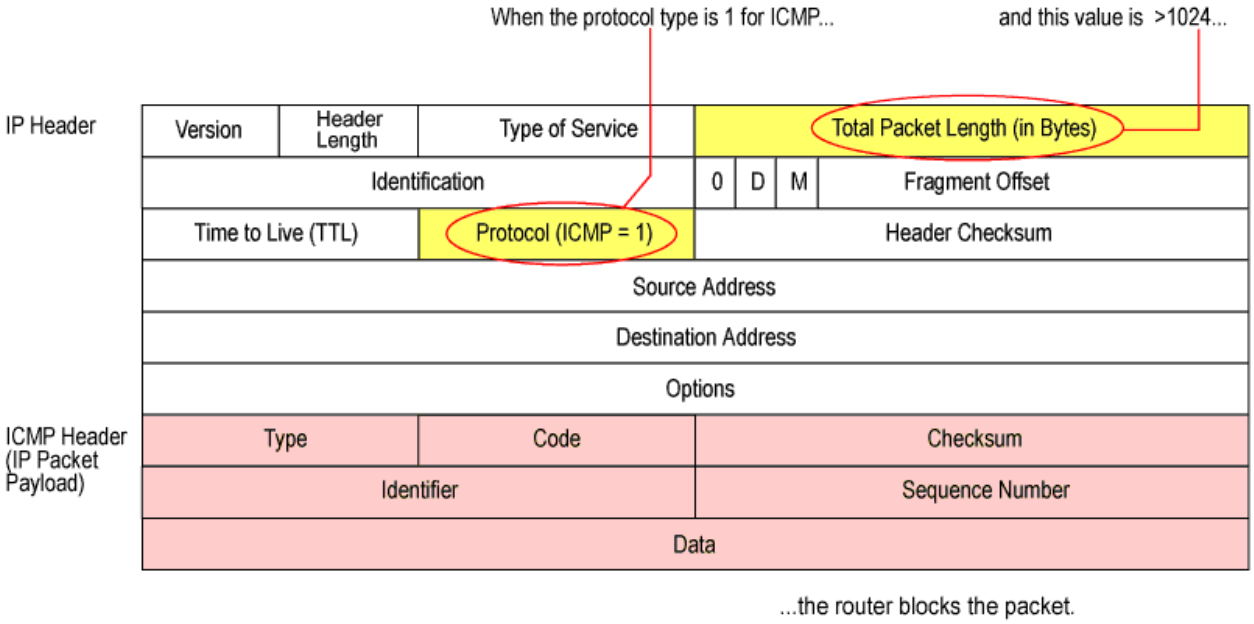

When you enable the Large Size ICMP Packet Protection screen option, JUNOS Software drops ICMP packets with a length greater than 1024 bytes.

Related Topics

- Blocking Large ICMP Packets on page 796

Blocking Large ICMP Packets

Because ICMP packets contain very short messages, there is no legitimate reason for large ICMP packets. If an ICMP packet is unusually large, something is wrong.

Before You Begin
For background information, read "Understanding Large ICMP Packet Protection" on page 794.

You can use either J-Web or the CLI configuration editor to block large ICMP packets. The specified security zone is the one from which the ICMP packets originated.

This topic covers:

- J-Web Configuration on page 796
- CLI Configuration on page 797
- Related Topics on page 797

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Edit**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **icmp-large**.
6. Next to Icmp, click Configure.
7. Next to Large, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.

6. In the Screen box, type **icmp-large** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option icmp-large icmp large
user@host# set security zones security-zone zone screen icmp-large
```

Related Topics

- Blocking Fragmented ICMP Packets on page 793
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798
- Dropping Packets Using an Unknown Protocol on page 800
- Dropping Fragmented IP Packets on page 803
- Dropping IP Packets Containing SYN Fragments on page 805

Understanding Bad IP Option Protection

The Internet Protocol standard RFC 791, *Internet Protocol*, specifies a set of eight options that provide special routing controls, diagnostic tools, and security. Although the original, intended uses for these options served worthy ends, people have figured out ways to twist these options to accomplish less commendable objectives. (For a summary of the exploits that attackers can initiate from IP options, see "Understanding Network Reconnaissance Using IP Options" on page 769.)

Before You Begin
For background information, read "Suspicious Packet Attributes Overview" on page 791.

Either intentionally or accidentally, attackers sometimes configure IP options incorrectly, producing either incomplete or malformed fields. Regardless of the intentions of the person who crafted the packet, the incorrect formatting is anomalous and potentially harmful to the intended recipient. See Figure 137 on page 798.

Figure 137: Incorrectly Formatted IP Options

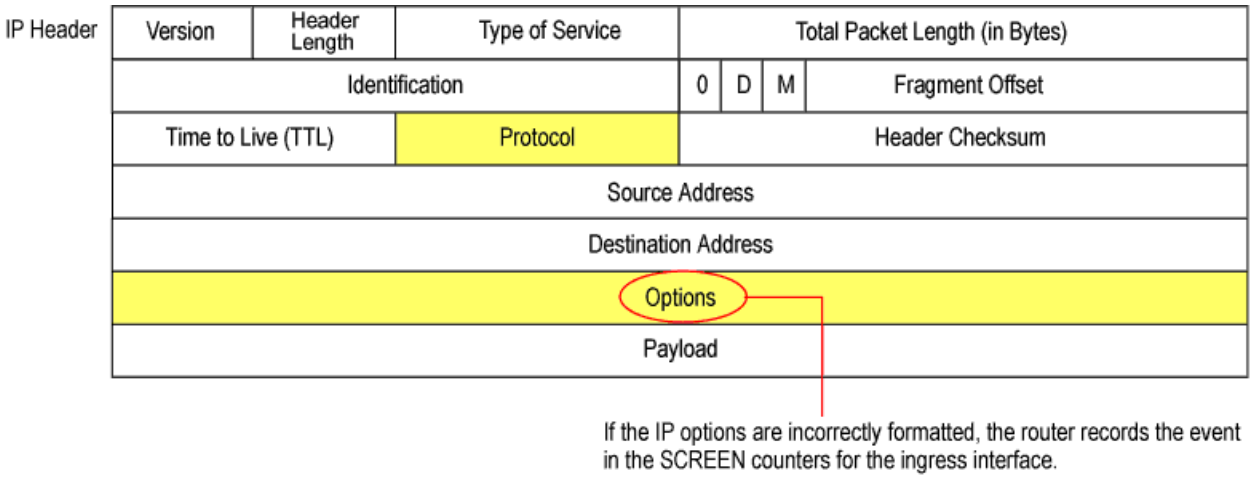

When you enable the bad IP option protection screen option, JUNOS Software blocks packets when any IP option in the IP packet header is incorrectly formatted. Additionally, JUNOS Software records the event in the event log.

Related Topics

- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798

Detecting and Blocking IP Packets with Incorrectly Formatted Options

Attackers sometimes configure IP options incorrectly, producing either incomplete or malformed fields. The incorrect formatting is anomalous and potentially harmful to the intended recipient.

Before You Begin
For background information, read “Understanding Bad IP Option Protection” on page 797.

You can use either J-Web or the CLI configuration editor to detect and block IP packets with incorrectly formatted IP options. The specified security zone is the one from which the packets originated.

This topic covers:

- J-Web Configuration on page 799
- CLI Configuration on page 799
- Related Topics on page 799

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **zone**.
6. Next to Ip, click Configure.
7. Next to Bad option, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **ip-bad-option** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option ip-bad-option ip bad-option
user@host# set security zones security-zone zone screen ip-bad-option
```

Related Topics

- Blocking Fragmented ICMP Packets on page 793
- Blocking Large ICMP Packets on page 796
- Dropping Packets Using an Unknown Protocol on page 800
- Dropping Fragmented IP Packets on page 803
- Dropping IP Packets Containing SYN Fragments on page 805

Understanding Unknown Protocol Protection

Based on RFC 1700, these protocol types with ID numbers of 137 or greater are reserved and undefined at this time. Precisely because these protocols are undefined,

there is no way to know in advance if a particular unknown protocol is benign or malicious.

Before You Begin
For background information, read “Suspicious Packet Attributes Overview” on page 791.

Unless your network makes use of a nonstandard protocol with an ID number of 137 or greater, a cautious stance is to block such unknown elements from entering your protected network. See Figure 138 on page 800.

Figure 138: Unknown Protocols

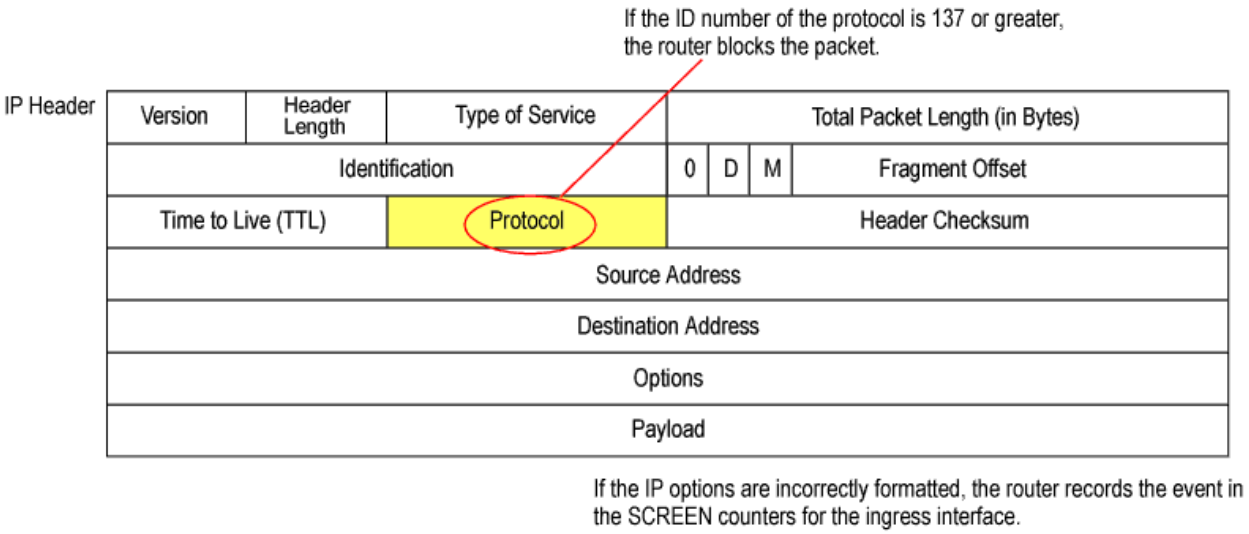

When you enable the unknown protocol protection screen option, JUNOS Software drops packets when the protocol field contains a protocol ID number of 137 or greater by default.

Related Topics

- Dropping Packets Using an Unknown Protocol on page 800

Dropping Packets Using an Unknown Protocol

Protocol types with ID numbers of 137 or greater are reserved and undefined at this time. Therefore, there is no way to know in advance if a particular unknown protocol is benign or malicious.

Before You Begin
For background information, read “Understanding Unknown Protocol Protection” on page 799.

You can use either J-Web or the CLI configuration editor to drop packets that use an unknown protocol. The specified security zone is the one from which the packets originates.

This topic covers:

- J-Web Configuration on page 801
- CLI Configuration on page 801
- Related Topics on page 801

J-Web Configuration

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **unknown-protocol** and click OK.
7. To save and commit the configuration, click Commit.

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **unknown-protocol**.
6. Next to Ip, click Configure.
7. Next to unknown-protocol, select the check box and click OK.
8. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security zones security-zone zone screen unknown-protocol
user@host# set security screen ids-option unknown-protocol ip unknown-protocol
```

Related Topics

- Blocking Fragmented ICMP Packets on page 793
- Blocking Large ICMP Packets on page 796
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798

- Dropping Fragmented IP Packets on page 803
- Dropping IP Packets Containing SYN Fragments on page 805

Understanding IP Packet Fragment Protection

As packets traverse different networks, it is sometimes necessary to break a packet into smaller pieces (fragments) based upon the maximum transmission unit (MTU) of each network. IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the victim receives these packets, the results can range from processing the packets incorrectly to crashing the entire system. See Figure 139 on page 802.

Before You Begin

For background information, read "Suspicious Packet Attributes Overview" on page 791.

Figure 139: IP Packet Fragments

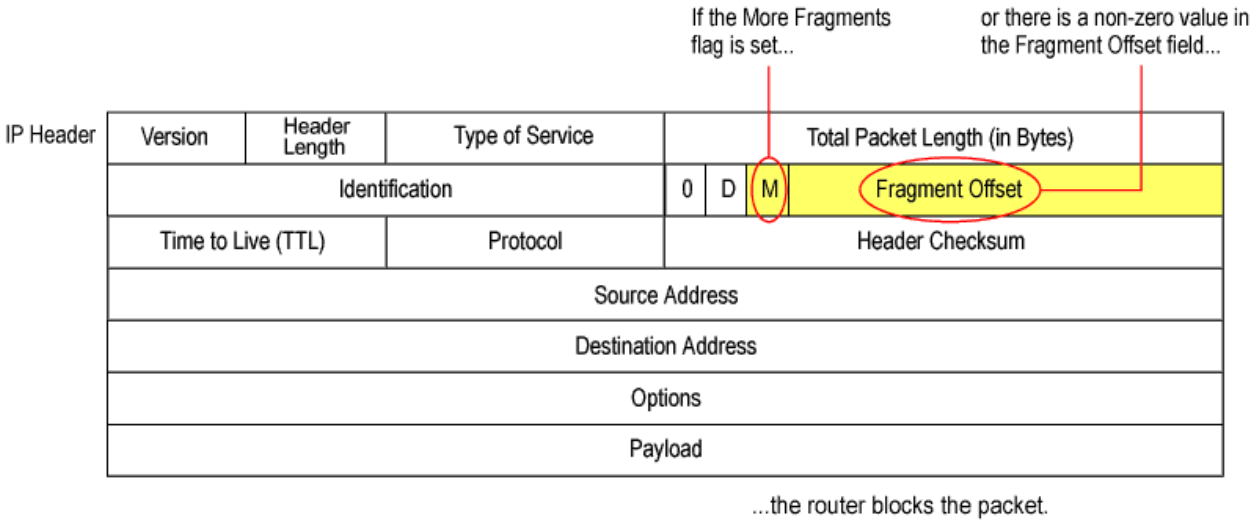

When you enable JUNOS Software to deny IP fragments on a security zone, it blocks all IP packet fragments that it receives at interfaces bound to that zone.

Related Topics

- “Dropping Fragmented IP Packets” on page 803

Dropping Fragmented IP Packets

IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations.

Before You Begin
For background information, read “Understanding IP Packet Fragment Protection” on page 802.

You can use either J-Web or the CLI configuration editor to drop fragmented IP packets. The specified security zone is the one from which the fragments originated.

This topic covers:

- J-Web Configuration on page 803
- CLI Configuration on page 804
- Related Topics on page 804

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **block-frag**.
6. Next to Ip, click Configure.
7. Next to Block-frag, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.

6. In the Screen box, type **block-frag** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option block-frag ip block-frag
user@host# set security zones security-zone zone screen block-frag
```

Related Topics

- Blocking Fragmented ICMP Packets on page 793
- Blocking Large ICMP Packets on page 796
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798
- Dropping Packets Using an Unknown Protocol on page 800
- Dropping IP Packets Containing SYN Fragments on page 805

Understanding SYN Fragment Protection

The IP encapsulates a TCP SYN segment in the IP packet that initiates a TCP connection. Because the purpose of this packet is to initiate a connection and invoke a SYN/ACK segment in response, the SYN segment typically does not contain any data. Because the IP packet is small, there is no legitimate reason for it to be fragmented.

Before You Begin
For background information, read “Suspicious Packet Attributes Overview” on page 791.

A fragmented SYN packet is anomalous, and as such it is suspect. To be cautious, block such unknown elements from entering your protected network. See Figure 140 on page 805.

Figure 140: SYN Fragments

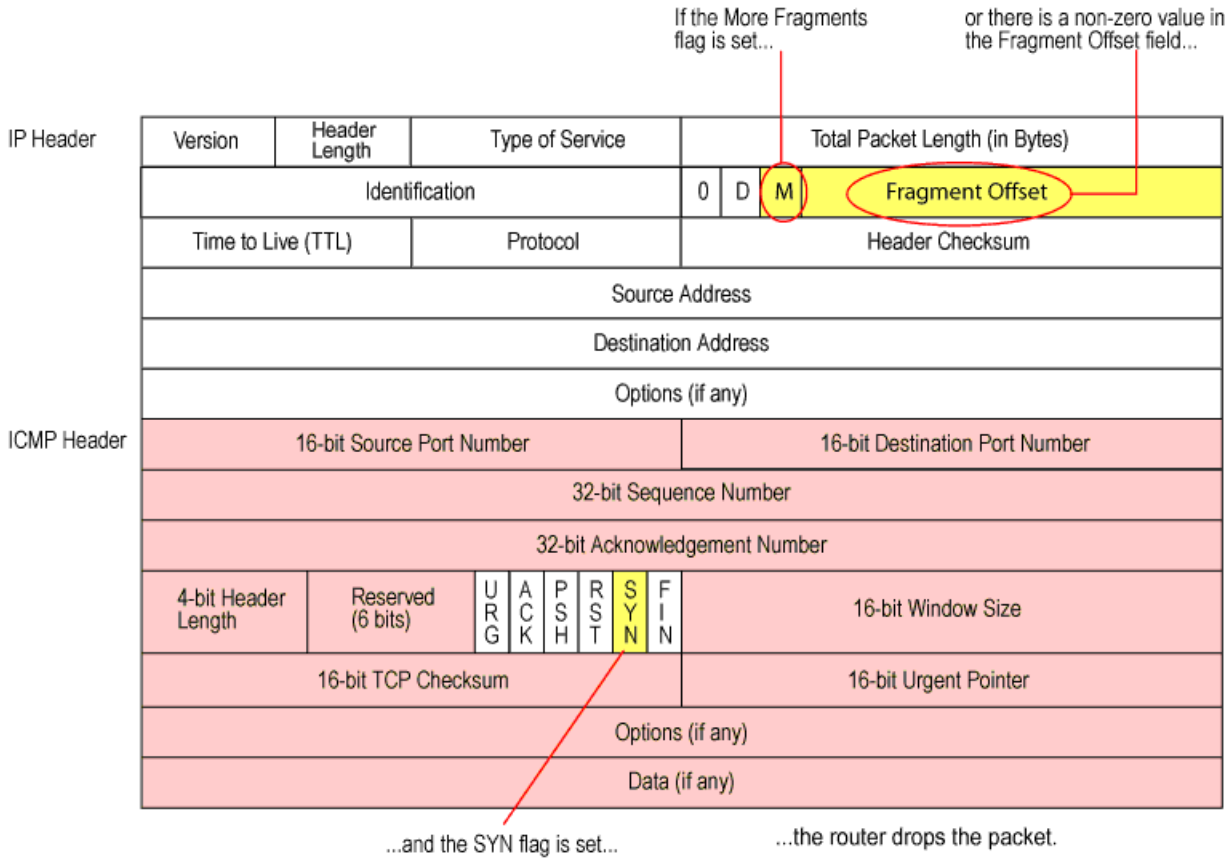

When you enable the SYN fragment detection screen option, JUNOS Software detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is set in the TCP header. JUNOS Software records the event in the screen counters list for the ingress interface.

Related Topics

- Dropping IP Packets Containing SYN Fragments on page 805

Dropping IP Packets Containing SYN Fragments

A fragmented SYN packet is anomalous, and as such it is suspect. To be cautious, block such unknown elements from entering your protected network.

Before You Begin
For background information, read “Understanding SYN Fragment Protection” on page 804.

You can use either J-Web or the CLI configuration editor to drop IP packets containing SYN fragments. The specified security zone is the one from which the packets originated.

This topic covers:

- J-Web Configuration on page 806
- CLI Configuration on page 806
- Related Topics on page 806

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **syn-frag**.
6. Next to tcp, click Configure.
7. Next to syn frag, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **syn-frag** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option syn-frag tcp syn-frag
user@host# set security zones security-zone zone screen syn-frag
```

Related Topics

- Blocking Fragmented ICMP Packets on page 793
- Blocking Large ICMP Packets on page 796
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 798

- Dropping Packets Using an Unknown Protocol on page 800
- Dropping Fragmented IP Packets on page 803

Denial-of-Service Attack Overview

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed. The target can be the Juniper Networks firewall, the network resources to which the firewall controls access, or the specific hardware platform or operating system of an individual host.

Before You Begin
For background information, read "Attack Detection and Prevention" on page 761.

If a DoS attack originates from multiple source addresses, it is known as a distributed denial-of-service (DDoS) attack. Typically, the source address of a DoS attack is spoofed. The source addresses in a DDoS attack might be spoofed, or the actual addresses of compromised hosts might be used as " zombie agents" to launch the attack.

The device can defend itself and the resources it protects from DoS and DDoS attacks.

Related Topics

- Firewall DoS Attacks Overview on page 815
- Network DoS Attacks Overview on page 823
- OS-Specific DoS Attacks Overview on page 847

Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview

The intent of an application-level distributed denial-of-service (AppDDoS) attack is to overwhelm the targeted server, such as a DNS or HTTP server. This is done by making a tremendous amount of application requests from malicious bot clients that often times use spoofed IP addresses. With all of these requests coming into the victim, it can not keep up and process legitimate traffic.

AppDDoS attacks are different than traditional L3/L4 DDoS attacks, such as a SYN flood. From a L3 and L4 perspective, the attack may appear as legitimate transactions. Traditional L3/L4 DDoS solutions can only rate limit these attacks and begin the application transactions, instead of denying those from attackers.

Related Topics

- Understanding Application-level DDoS Rulebase on page 504

- Enabling Protection Against Application-level Distributed Denial-of-Service (AppDDoS) Attack on page 812
- Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview on page 808

Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview

- Understanding Application-level DDoS Module on page 808
- Understanding Application-ddos Definition on page 809
- Understanding AppDDoS Rule on page 810
- Understanding AppDDoS IP-action on page 811
- Understanding Session Action on page 811
- Related Topics on page 811

Understanding Application-level DDoS Module

The AppDDoS IDP module uses application-level metrics to differentiate between normal and malicious application requests. It will then identify the offending source addresses and can drop or deny these requests. Based on user configured application thresholds, when the client application transactions exceeds the defined thresholds, session and ip-actions will be applied on traffic from the offending client address. This feature will protect servers against DNS and HTTP application DDoS attacks.

In order to identify malicious bot clients, you create a policy with rulebase-ddos to monitor specific traffic, define AppDDoS application metrics and thresholds to monitor that traffic. When the threshold are exceeded, your defined action is taken on the client to protect the application server.

IDP performs multi-stage analysis from connection monitoring, protocol analysis, to bot client classification, and maintains the state for each protected server. You can configure connection-rate-threshold in the AppDDoS application definition to monitor connection rate. When connection rate are exceeded, IDP transitions to protocol analysis for deep content inspection and maintains statistical data on application transactions. The AppDDoS attacks can be classified into either heavy hitters or random hitters. Heavy hitters perform identical application transactions in a fast and repeated fashion. For example, querying a non-existing domain name repeatedly. Random hitters perform random application transactions, for example, querying a random domain name, one per each request. You can configure context value-hit-rate-threshold to detect heavy hitters and context hit-rate-threshold to detect random hitters. If either of the above context thresholds are exceeded, IDP transitions to the bot client classification stage where it tracks the application transactions on a per client basis based on user configured time-binding thresholds. A benign client will not perform identical and repeated transactions; whereas a malicious bot clients will. Once time-binding thresholds are exceeded, identified bot clients will be blocked with the configured ip-action and session actions.

You can also configure a list of regular expressions under exclude-context-values to exempt certain context values from being considered for AppDDoS processing. This

is helpful for requests for well known resources that may often hit context thresholds, for example, a DNS query for domain name google.com.

Protocol analysis stage uses a default interval of 60 seconds for context hit-rate-threshold and value-hit-rate-threshold. For example, if you configure 10,000 as the value-hit-rate threshold, the context value would be monitored against 10,000 hits limit in a 60 seconds interval.

IDP also uses hysteresis for state transitions to avoid thrashing between the states. A default of 20% lower limit will be used from the configured connection and context thresholds for falling behind in state. For example. if you configure a context value-hit-rate-threshold of 10,000, IDP transitions from protocol analysis to bot client classification after 10000 hits in 60 seconds for identical context values, and falls behind in state only when such hits are smaller than 8000 in 60 seconds.

It is recommended to configure time-binding thresholds in the application-ddos definition, since it is critical to enumerate benign client versus a malicious bot client. However, if you chose to not define time-binding thresholds, IDP does not do bot client classification. In this case, if application transactions exceed context thresholds, the configured ip-action and session actions will be performed. Note that without bot client classification, benign clients might get denied when making a request to the protected server.

IDP maintains AppDDoS state for the current policy only. For more information, see the policy aging reference in “Multiple IDP Detector Support” on page 561. Traffic from sessions using older policy will not be inspected for AppDDoS. If a new policy is loaded, AppDDoS state for each protected server will be relearned.

Understanding Application-ddos Definition

You can only configure one application-ddos definition for each protected server. However you can use the same application-ddos definition in two or more rules with specific destination-address and/or to-zone to protect two or more servers with same desired applicaiton-ddos thresholds. Table Table 101 on page 810 shows the parameters that can be set for application-ddos.

NOTE: Application-level denial-of-service (AppDDoS) detection will not work if two rules with different AppDDoS applications hit traffic going to a single destination application server. When setting up AppDDoS rules, make sure you do not configure rulebase-ddos rules that have two different application-ddos objects while the traffic destined to one application server can hit more than one rule. Essentially, for each protected application server, you have to configure AppDDoS rules so traffic destined for one protected server only hits one AppDDoS rule.

AppDDoS rules are terminal, which means that once traffic is processed by one rule, it will not be processed by other rules.

For more details on the following commands, see the *JUNOS Software CLI Reference* guide.

Table 101: Application DDoS Parameters

Parameter	Description
service ***service-name***	The application layer service to be monitored, such as DNS or HTTP.
connection-rate threshold	The connections-per-second threshold at which to start monitoring the application context values.
context ***context-name***	Name of the application context that the IDP protocol decoder generates while parsing the application protocol from traffic data.
hit-rate-threshold	Number of context hits in tick interval (60 seconds by default) to start bot client classification, if time binding parameters are configured. If time binding parameters are not configured, the configured policy actions are taken.
value-hit-rate-threshold	Number of context value hits in tick interval to start bot client classification, if time binding parameters are configured. If time binding parameters are not configured, the configured policy actions are taken.
max-context-values	The top *n* context values that should be monitored, reported, or both.
time-binding-period	The time-binding period to determine if a client should be classified as a malicious bot client or not. This setting is used in conjunction with time-binding count to detect an attack if a client request for the same context value exceeds time-binding-count times in time-binding-period seconds.
time-binding-count	The number of context or context value hits from each client over the time binding period to determine if it should be considered a malicious bot client.

Understanding AppDDoS Rule

You configure one or more AppDDoS rules to define the traffic that should be monitored in order to protect your servers. Table Table 102 on page 810 shows the parameters that can be set for application-ddos.

Table 102: AppDDoS Rule Parameters

Parameter	Description
from-zone	Match source zone.
source-address	Match source address.
to-zone	Match destination zone.
destination-address	Match destination address.
application	Choose default to select the application service from the application-ddos definition.
application-ddos	Specify the DDoS application.

Understanding AppDDoS IP-action

You configure ip-action to drop future sessions from identified bot client addresses for a specified time. Table Table 103 on page 811 shows the parameters that can be set for ip-action.

Table 103: AppDDoS IP-action Parameters

Parameter	Description
ip-block	Block future connections of any session that matches the IP action.
ip-close	Closes future connections of any client address that matches ip-action by sending an RST packet to the client.
ip-notify	Does not take any action against matching future connections, but do log the event.

Understanding Session Action

Configure the action to perform on the identified bot client.

Table Table 104 on page 811 shows the parameters that can be set for action.

Table 104: AppDDoS Action Parameters

Parameter	Description
close-server	Closes the connection and sends an RST packet to the server but not to the client.
drop-connection	Drops all packets associated with the connection, preventing traffic for the connection from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
drop-packet	Drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source-IP address.
no-action	No action is taken. Use this action when you want to only generate logs for some traffic.

Related Topics

- Enabling Protection Against Application-level Distributed Denial-of-Service (AppDDoS) Attack on page 812
- Understanding Application-level DDoS Rulebase on page 504
- Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview on page 807

Enabling Protection Against Application-level Distributed Denial-of-Service (AppDDoS) Attack

This example will describe how you can use the AppDDoS module to protect a DNS server.

When setting up AppDDoS protection for a DNS server, you will first observe the average load of DNS requests. If the DNS server is expected to handle a normal load of 1000 requests per second, choose 20% in excess of the normal load (1200 requests per second) as the connection-rate-threshold. This is essentially 60,000 transactions in 60 seconds, so choose 20% in excess of this load as context hit-rate-threshold (72,000). You can choose context value-hit-rate-threshold based on the maximum load of requests for the same domain name being queried for non exempt context dns-type-cname values. For example, if it is impractical for DNS to receive queries for domain xyz.com in excess of 2000 times in 60 seconds, context value-hit-rate-threshold should be set to 20% more than that value, which would be 2400 times in 60 seconds. For monitoring and reporting, you will optionally set max-context-values to 100, so at the maximum, the most active 100 DNS query requests will be monitored and reported. If a client is in this range, it is mostly likely a malicious bot client. Once bot clients are identified, you can configure ip-action as ip-block with timeout as 600 seconds (the bot client gets access denied for 1 hour) and session action is set as drop-packet.

In the above example, IDP starts deep protocol analysis when the number of connections per second exceeds 1200, and will start bot client classification either if total number of queries for context dns-type-name exceeds 72000 or requests for the same query value exceeds 2400.

NOTE: When an AppDDoS attack occurs on the application server, it will have much higher transaction rates than it does under normal or even peak load. With this in mind, it is best to set higher thresholds than the normal peak of the application server so it does not trigger unnecessary client classification processing. This will improve the over all performance of the Juniper device since the AppDDoS module will not start client classification until the server has actually reached abnormal transaction rates.

Before You Begin

For background information, read “Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview” on page 807, and “Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview” on page 808.

To enable protection against AppDDoS for a DNS server, use the CLI configuration editor.

For detailed information about the following commands, see the *JUNOS Software CLI Reference.*

This topic covers:

■ CLI Configuration on page 813

CLI Configuration

1. Access the IDP Security configuration hierarchy.

```
[edit]
user@host# edit security idp
```

2. Configure application-ddos to define the type of traffic, the protocol context that will be monitored, and thresholds that will be used to trigger an action. In this example, DNS traffic and the protocol context dns-type-name will be monitored.

```
[edit security idp]
user@host# set application-ddos dns-server-1
[edit security idp]
user@host# set application-ddos dns-server service dns
[edit security idp]
user@host# set application-ddos dns-server-1 connection-rate-threshold 1200
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  hit-rate-threshold 72000
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  value-hit-rate-threshold 2400
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  max-context-values 100
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  time-binding-count 10
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  time-binding-period 30
```

NOTE: You can continue to set other protocol contexts and thresholds that you would like to monitor.

3. (Optional) Set contexts values that will be exempt from monitoring.

```
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  exclude-context-values .*google.com
[edit security idp]
user@host# set application-ddos dns-server-1 context dns-type-name
  exclude-context-values .*yahoo.com
```

4. Set IDP policy rule for rulebase-ddos to define the source and destination of traffic that will be monitored.

NOTE: You can only define one DDoS application per AppDDoS rule. Create additional rules to monitor multiple DDoS applications.

Each AppDDoS rule is a terminal rule, meaning that only one matching rule is considered for incoming traffic matching.

```
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   match source-address any
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   match to-zone any
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   match destination-address any
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   match application default
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   match application-ddos dns-server-1
[edit security idp]
```

5. Define the action to be taken when AppDDoS attack traffic is detected.

```
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   then action drop-packet
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   then ip-action ip-block
[edit security idp]
user@host# set idp-policy AppDDoS-policy-1 rulebase-ddos rule AppDDoS-rule1
   then ip-action timeout 600
```

Related Topics

- Application-level Distributed Denial-of-Service (AppDDoS) Protection Overview on page 808
- Understanding Application-level DDoS Rulebase on page 504
- Application-level Distributed Denial-of-Service (AppDDoS) Attack Overview on page 807

Firewall DoS Attacks Overview

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed.

Before You Begin

For background information, read "Denial-of-Service Attack Overview" on page 807.

If attackers discover the presence of the Juniper Networks firewall, they might launch a DoS attack against it instead of the network behind it. A successful DoS attack against a firewall amounts to a successful DoS attack against the protected network in that it thwarts attempts of legitimate traffic to traverse the firewall.

An attacker might use session table floods and SYN-ACK-ACK proxy floods to fill up the session table of JUNOS Software and thereby produce a DoS.

Related Topics

- Understanding Session Table Flood Attacks on page 815
- Understanding SYN-ACK-ACK Proxy Flood Attacks on page 821
- Network DoS Attacks Overview on page 823
- OS-Specific DoS Attacks Overview on page 847

Understanding Session Table Flood Attacks

A successful DoS attack overwhelms its victim with such a massive barrage of false simulated traffic that it becomes unable to process legitimate connection requests. DoS attacks can take many forms—SYN flood, SYN-ACK-ACK flood, UDP flood, ICMP flood, and so on—but they all seek the same objective, which is to fill up their victim's session table.

Before You Begin

For background information, read "Firewall DoS Attacks Overview" on page 815.

When the session table is full, that host cannot create any new sessions and begins rejecting new connection requests. The source-based session limits screen option and the destination-based session limit screen option help mitigate such attacks.

This topic covers:

- Source-Based Session Limits on page 816
- Destination-Based Session Limits on page 817
- Related Topics on page 817

Source-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. One benefit of setting a source-based session limit is that it can stem an attack such as the Nimda virus (which is actually both a virus and a worm) that infects a server and then begins generating massive amounts of traffic from that server. Because all the virus-generated traffic originates from the same IP address, a source-based session limit ensures that the firewall can curb such excessive amounts of traffic. See Figure 141 on page 816.

Figure 141: Limiting Sessions Based on Source IP Address

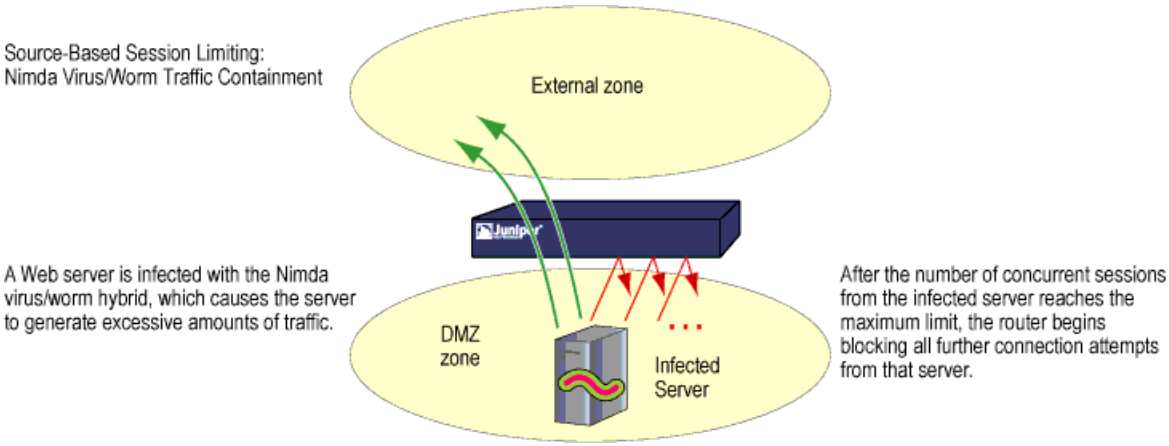

Another benefit of source-based session limiting is that it can mitigate attempts to fill up the firewall's session table, if all the connection attempts originate from the same source IP address.

Determining what constitutes an acceptable number of connection requests requires a period of observation and analysis to establish a baseline for typical traffic flows. You also need to consider the maximum number of concurrent sessions required to fill up the session table of the particular Juniper Networks platform you are using. To see the maximum number of sessions that your session table supports, use the CLI command get session, and then look at the first line in the output, which lists the number of current (allocated) sessions, the maximum number of sessions, and the number of failed session allocations:

```
alloc 420/max 128000, alloc failed 0
```

The default maximum for source-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

Destination-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. A wily attacker can launch a distributed denial-of-service (DDos) attack. In a DDoS attack, the malicious traffic can come from hundreds of hosts, known as " zombie agents," that are surreptitiously under the control of an attacker. In addition to the SYN, UDP, and ICMP flood detection and prevention screen options, setting a destination-based session limit can ensure that JUNOS Software allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. See Figure 142 on page 817.

Figure 142: Distributed DOS Attack

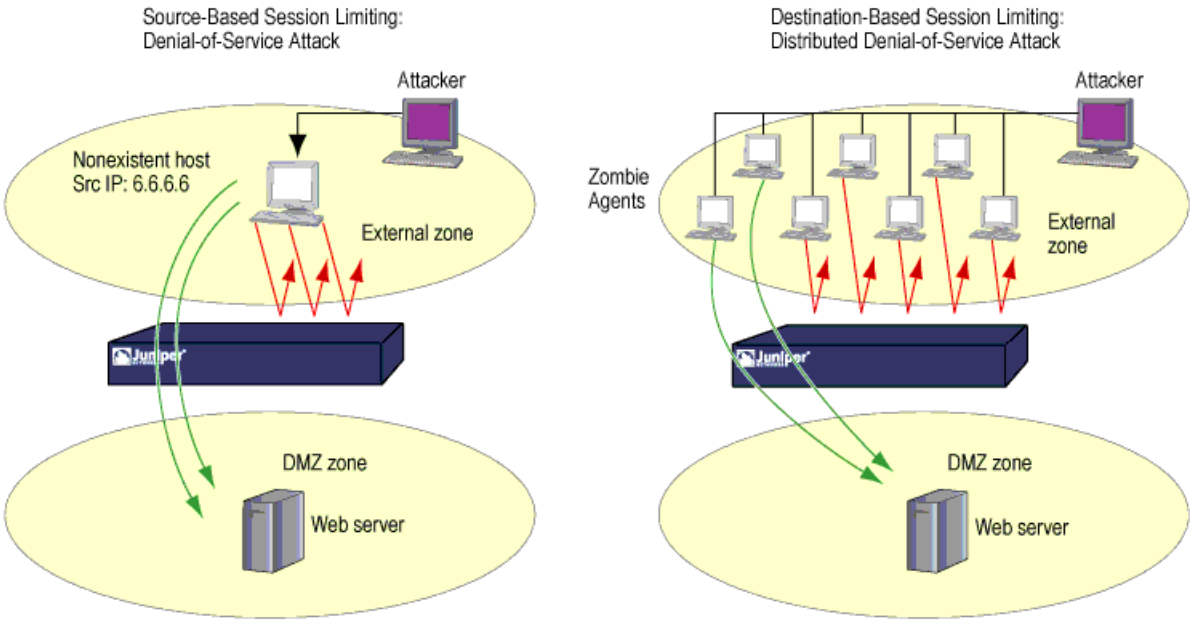

The default maximum for destination-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

Related Topics

- Setting Source-Based Session Limits on page 818
- Setting Destination-Based Session Limits on page 819

Setting Source-Based Session Limits

A source-based session limit can stem an attack that infects a server and then begins generating massive amounts of traffic from that server.

Before You Begin

For background information, read "Understanding Session Table Flood Attacks" on page 815.

In this example, you want to limit the amount of sessions that any one server in the DMZ and zone_a zones can initiate. Because the DMZ zone only contains Web servers, none of which should initiate traffic, you set the source-session limit at the lowest possible value: 1 session. On the other hand, the zone_a zone contains personal computers, servers, printers, and so on, many of which do initiate traffic. For the zone_a zone, you set the source-session limit maximum to 80 concurrent sessions.

You can use either J-Web or the CLI configuration editor to set the source-session limit. In this example you are setting the source-session limit maximum to 80 concurrent sessions.

This topic covers:

- J-Web Configuration on page 818
- CLI Configuration on page 819
- Related Topics on page 819

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **1-limit-session**.
6. Next to Limit session, select the check box and click Configure.
7. In the Source ip based box, type **1** and click OK.
8. To configure another Source-ip-based 100 and 80, repeat Step 4 through Step 7 and click OK.
9. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **dmz**.
6. In the Screen box, type **100-limit-session** and click OK.
7. Next to Security zone, click Add new entry.
8. In the Name box, type **zone_a**.
9. In the Screen box, type **100-limit-session** and click OK.
10. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option 1-limit-session limit-session source-ip-based
  1
user@host# set security screen ids-option 100-limit-session limit-session
  source-ip-based 100
user@host# set security screen ids-option 80-limit-session limit-session
  source-ip-based 80
user@host# set security zones security-zone dmz screen 100-limit-session
user@host# set security zones security-zone zone_a screen 100-limit-session
```

Related Topics

- Setting Destination-Based Session Limits on page 819
- Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack on page 821

Setting Destination-Based Session Limits

In addition to the SYN, UDP, and ICMP flood detection and prevention screen options, setting a destination-based session limit can ensure that JUNOS Software allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host.

Before You Begin
For background information, read "Understanding Session Table Flood Attacks" on page 815.

In this example, you want to limit the amount of traffic to a Web server at 1.2.2.5. The server is in the DMZ zone. After observing the traffic flow from the external zone to this server for a month, you have determined that the average number of concurrent sessions it receives is 2000. Based on this information, you decide to set the new session limit at 4000 concurrent sessions. Although your observations show

that traffic spikes sometimes exceed that limit, you opt for firewall security over occasional server inaccessibility.

To set the destination-session limit, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 820
- CLI Configuration on page 820
- Related Topics on page 821

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **4000-limit-session**.
6. Next to the Limit session, select the check box and click Configure.
7. In the Destination ip based box, type **4000** and click Commit.
8. To configure another Destination-ip-based 100, repeat Step d through Step g and click Commit.
9. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **external_zone**.
6. In the Screen box, type **100-limit-session** and click Commit.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option 4000-limit-session limit-session
  destination-ip-based 4000
user@host# set security screen ids-option 100-limit-session limit-session
  destination-ip-based 100
user@host# set security zones security-zone external_zone screen 100-limit-session
```

Related Topics

- Setting Source-Based Session Limits on page 818
- Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack on page 821

Understanding SYN-ACK-ACK Proxy Flood Attacks

When an authentication user initiates a Telnet or FTP connection, the user sends a SYN segment to the Telnet or FTP server. JUNOS Software intercepts the SYN segment, creates an entry in its session table, and proxies a SYN-ACK segment to the user. The user then replies with an ACK segment. At this point, the initial three-way handshake is complete. JUNOS Software sends a login prompt to the user. If the user, with malicious intent, does not log in, but instead continues initiating SYN-ACK-ACK sessions, the firewall session table can fill up to the point where the device begins rejecting legitimate connection requests.

Before You Begin
For background information, read "Firewall DoS Attacks Overview" on page 815.

To prevent such an attack, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, JUNOS Software rejects further connection requests from that IP address. By default, the threshold is 512 connections from any single IP address. You can change this threshold (to any number between 1 and 250,000) to better suit the requirements of your network environment.

Related Topics

- Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack on page 821

Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack

Malicious users can fill up the firewall session table to the point where the device begins rejecting legitimate connection requests by continuously initiating SYN-ACK-ACK sessions.

Before You Begin
For background information, read "Understanding SYN-ACK-ACK Proxy Flood Attacks" on page 821.

To enable protection against a SYN-ACK-ACK proxy flood, use either J-Web or the CLI configuration editor. The specified zone is where the attack originated.

NOTE: The value unit is connections per source address. The default value is 512 connections from any single address.

This topic covers:

- J-Web Configuration on page 822
- CLI Configuration on page 822
- Related Topics on page 822

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **1000-syn-ack-ack-proxy**.
6. Next to Tcp, click Configure.
7. Next to Syn ack ack proxy box, select the check box and click Configure.
8. In the Threshold box, type **1000** and click OK.
9. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **1000-syn-ack-ack-proxy** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option 1000-syn-ack-ack-proxy tcp
  syn-ack-ack-proxy threshold 1000
user@host# set security zones security-zone zone screen 1000-syn-ack-ack-proxy
```

Related Topics

- Setting Source-Based Session Limits on page 818

- Setting Destination-Based Session Limits on page 819

Network DoS Attacks Overview

A denial-of-service (DoS) attack directed against one or more network resources floods the target with an overwhelming number of SYN, ICMP, or UDP packets, or with an overwhelming number of SYN fragments.

Before You Begin

For background information, read "Denial-of-Service Attack Overview" on page 807.

Depending on the attackers' purpose and the extent and success of previous intelligence gathering efforts, the attackers might single out a specific host, such as a device or server; they might aim at random hosts across the targeted network. Either approach has the potential of upsetting service to a single host or to the entire network, depending on how critical the role of the victim is to the rest of the network.

Related Topics

- Understanding SYN Flood Attacks on page 823
- Understanding SYN Cookie Protection on page 836
- Understanding ICMP Flood Attacks on page 839
- Understanding UDP Flood Attacks on page 842
- Understanding Land Attacks on page 844

Understanding SYN Flood Attacks

A SYN flood occurs when a host becomes so overwhelmed by SYN segments initiating incomplete connection requests that it can no longer process legitimate connection requests.

Before You Begin

For background information, read "Network DoS Attacks Overview" on page 823.

Two hosts establish a TCP connection with a triple exchange of packets known as a *three-way handshake*: A sends a SYN segment to B; B responds with a SYN/ACK segment; and A responds with an ACK segment. A SYN flood attack inundates a site with SYN segments containing forged (spoofed) IP source addresses with nonexistent or unreachable addresses. B responds with SYN/ACK segments to these addresses and then waits for responding ACK segments. Because the SYN/ACK segments are sent to nonexistent or unreachable IP addresses, they never elicit responses and eventually time out. See Figure 143 on page 824.

Figure 143: SYN Flood Attack

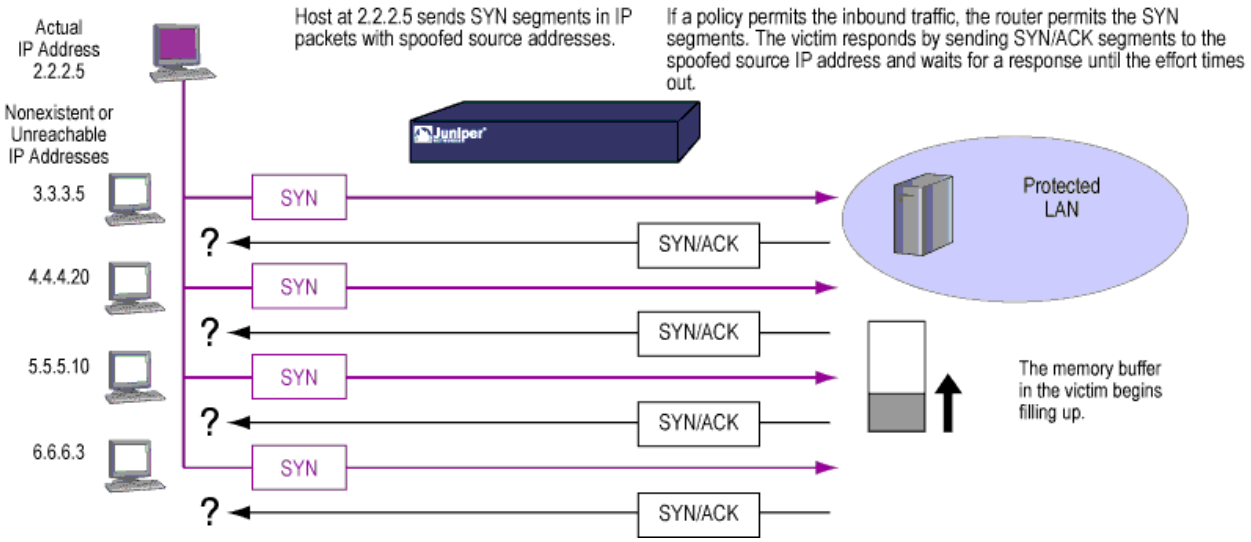

By flooding a host with incomplete TCP connections, the attacker eventually fills the memory buffer of the victim. Once this buffer is full, the host can no longer process new TCP connection requests. The flood might even damage the victim's operating system. Either way, the attack disables the victim and its normal operations.

This topic covers:

- SYN Flood Protection on page 824
- SYN Flood Options on page 826
- Related Topics on page 828

SYN Flood Protection

JUNOS Software can impose a limit on the number of SYN segments permitted to pass through the firewall per second. You can base the attack threshold on the destination address and port, the destination address only, or the source address only. When the number of SYN segments per second exceeds one of these thresholds, JUNOS Software starts proxying incoming SYN segments, replying with SYN/ACK segments and storing the incomplete connection requests in a connection queue. The incomplete connection requests remain in the queue until the connection is completed or the request times out. In Figure 144 on page 825, the SYN attack threshold has passed, and JUNOS Software has started proxying SYN segments.

Figure 144: Proxying SYN Segments

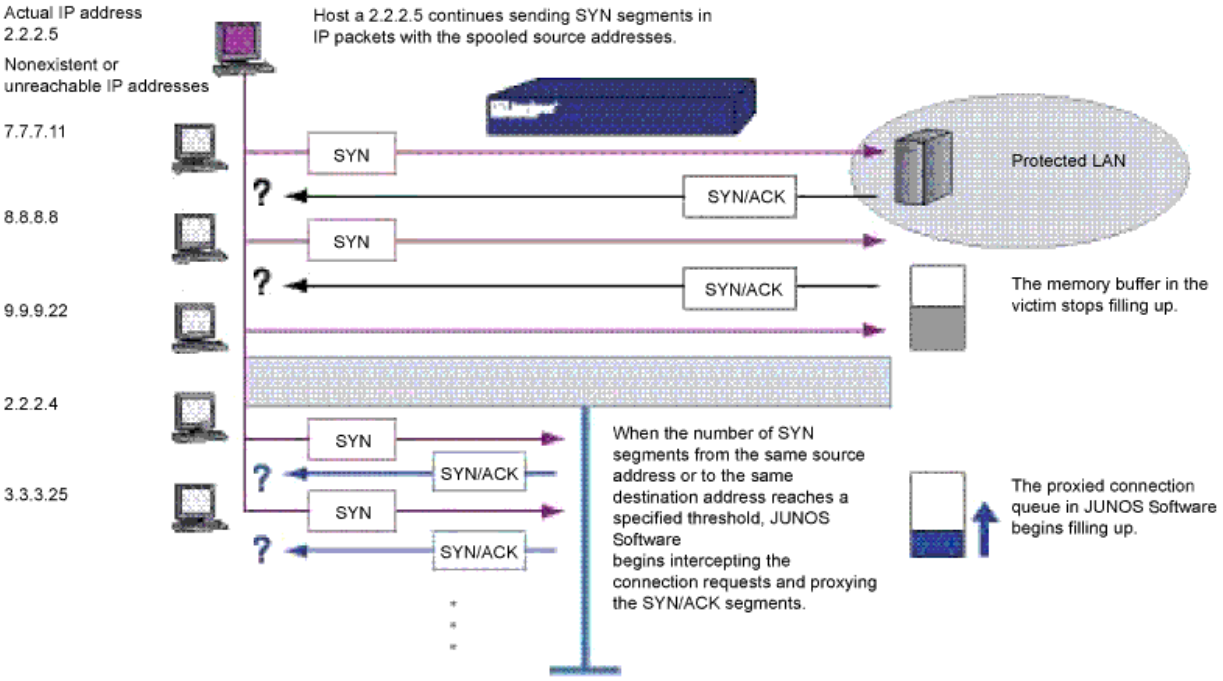

In Figure 145 on page 826, the proxied connection queue has completely filled up, and JUNOS Software is rejecting new incoming SYN segments. This action shields hosts on the protected network from the bombardment of incomplete three-way handshakes.

Figure 145: Rejecting New SYN Segments

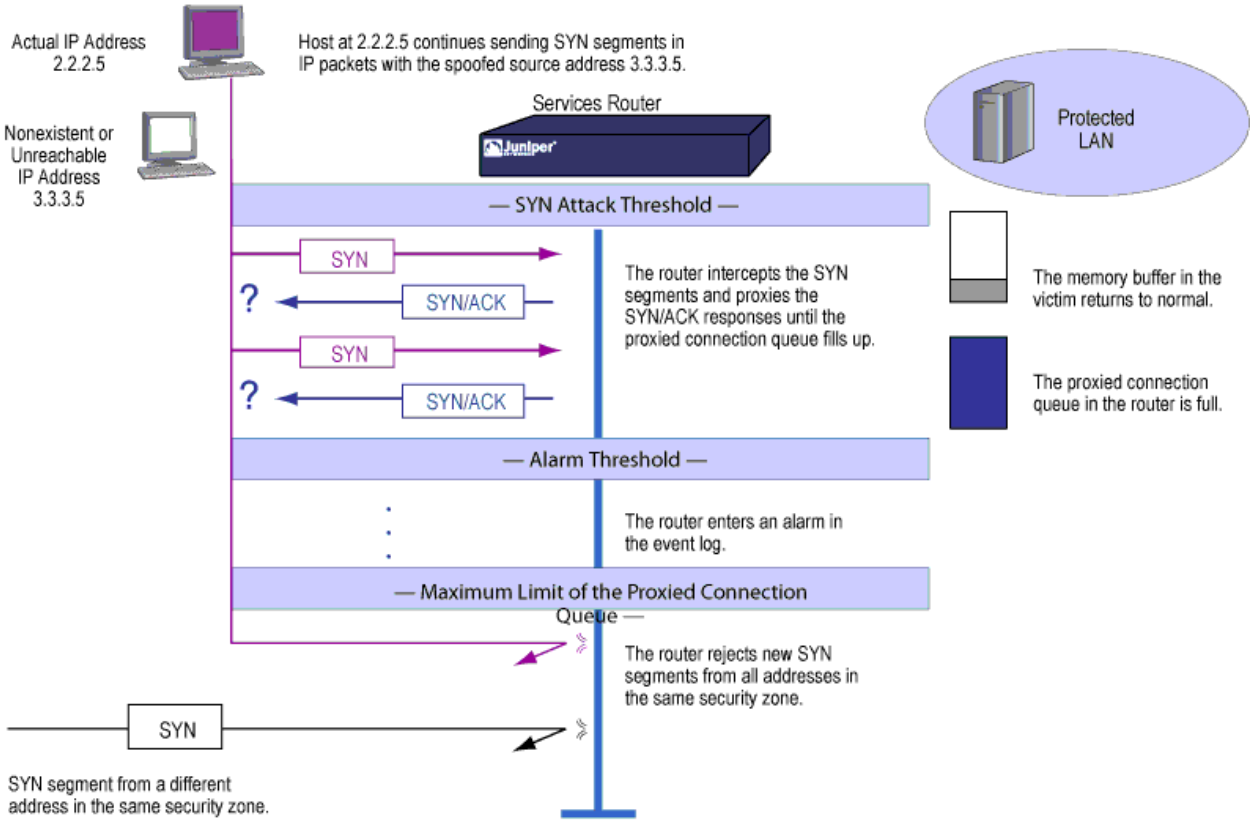

The device starts receiving new SYN packets when the proxy queue drops below the maximum limit.

NOTE: The procedure of proxying incomplete SYN connections above a set threshold pertains only to traffic permitted by existing policies. Any traffic for which a policy does not exist is automatically dropped.

SYN Flood Options

You can set the following parameters for proxying uncompleted TCP connection requests:

- Attack Threshold: This option allows you to set the number of SYN segments (that is, TCP segments with the SYN flag set) to the same destination address and port number per second required to activate the SYN proxying mechanism. Although you can set the threshold to any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, if it is an e-business site that normally gets 20,000 SYN segments per second, you might want to set the threshold to 30,000 per second. If a smaller site normally gets 20 SYN segments per second, you might consider setting the threshold to 40.

```
user@host# set security screen zone-syn-flood tcp syn-flood attack-threshold
  number
user@host# set security zones security-zone zone screen zone-syn-flood
```

- Alarm Threshold: This option allows you to set the number of proxied, half-complete TCP connection requests per second after which JUNOS Software enters an alarm in the event log. The value you set for an alarm threshold triggers an alarm when the number of proxied, half-completed connection requests to the same destination address and port number per second exceeds that value. For example, if you set the SYN attack threshold at 2000 SYN segments per second and the alarm at 1000, then a total of 3001 SYN segments to the same destination address and port number per second is required to trigger an alarm entry in the log. More precisely:

  1. The firewall passes the first 2000 SYN segments per second that meet policy requirements.
  2. The firewall proxies the next 1000 SYN segments in the same second.
  3. The 1001st proxied connection request (or 3001st connection request in that second) triggers the alarm.

```
user@host# set security screen zone-syn-flood tcp syn-flood alarm-threshold
<number>

user@host# set security zones security-zone zone screen zone-syn-flood
```

  For each SYN segment to the same destination address and port number in excess of the alarm threshold, the attack detection module generates a message. At the end of the second, the logging module compresses all similar messages into a single log entry that indicates how many SYN segments to the same destination address and port number arrived after exceeding the alarm threshold. If the attack persists beyond the first second, the event log enters an alarm every second until the attack stops.

- Source Threshold: This option allows you to specify the number of SYN segments received per second from a single source IP address—regardless of the destination IP address and port number—before JUNOS Software begins dropping connection requests from that source.

```
user@host# set security screen zone-syn-flood tcp syn-flood
source-threshold <number>

user@host# set security zones security-zone zone screen zone-syn-flood
```

Tracking a SYN flood by source address uses different detection parameters from tracking a SYN flood by destination address and destination port number. When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection mechanism and the source-based SYN flood tracking mechanism in effect.

- Destination Threshold: This option allows you to specify the number of SYN segments received per second for a single destination IP address before JUNOS Software begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based on destination IP address only—regardless of the destination port number.

```
user@host# set security screen zone-syn-flood tcp syn-flood
destination-threshold <number>

user@host# set security zones security-zone zone screen zone-syn-flood
```

  When you set a SYN attack threshold and a destination threshold, you put both the basic SYN flood protection mechanism and the destination-based SYN flood tracking mechanism in effect.

  Tracking a SYN flood by destination address uses different detection parameters from tracking a SYN flood by destination address and destination port number. Consider the following case where JUNOS Software has policies permitting FTP requests (port 21) and HTTP requests (port 80) to the same server. If the SYN flood attack threshold is 1000 packets per second (pps) and an attacker sends 999 FTP packets and 999 HTTP packets per second, neither set of packets (where a set is defined as having the same destination address and port number) activates the SYN proxying mechanism. The basic SYN flood attack mechanism tracks destination address and port number, and neither set exceeds the attack threshold of 1000 pps. However, if the destination threshold is 1000 pps, JUNOS Software treats both FTP and HTTP packets with the same destination address as members of a single set and rejects the 1001st packet—FTP or HTTP—to that destination.

- Timeout: This option allows you to set the maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 0–50 seconds. You might try decreasing the timeout value to a shorter length until you begin to see any dropped connections during normal traffic conditions. Twenty seconds is a very conservative timeout for a three-way handshake ACK response.

```
user@host# set security screen zone-syn-flood tcp syn-flood timeout
<number>

user@host# set security zones security-zone zone screen zone-syn-flood
```

Related Topics

- Understanding SYN Cookie Protection on page 836
- Enabling SYN Cookie Protection on page 837
- Enabling SYN Flood Protection on page 835

Example: SYN Flood Protection

In this example, you protect four Web servers in the DMZ zone from SYN flood attacks originating in the external zone by enabling the SYN flood protection SCREEN option for the external zone.

Before You Begin
For background information, read “Understanding SYN Flood Attacks” on page 823.

NOTE: We recommend that you augment the SYN flood protection that JUNOS Software provides with device-level SYN flood protection on each of the Web servers. In this example, the Web servers are running UNIX, which also provides some SYN flood defenses, such as adjusting the length of the connection request queue and changing the timeout period for incomplete connection requests.

Figure 146: Device-Level SYN Flood Protection

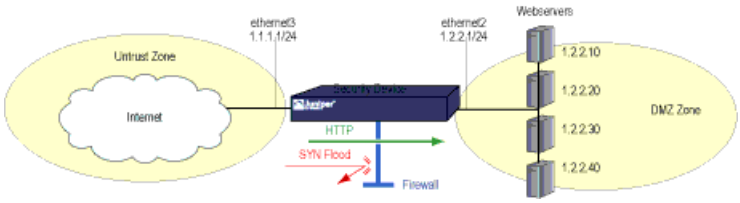

To configure the SYN flood protection parameters with appropriate values for your network, you must first establish a baseline of typical traffic flows. For one week, you run a sniffer on ethernet3—the interface bound to zone_external—to monitor the number of new TCP connection requests arriving every second for the four Web servers in the DMZ zone. Your analysis of the data accumulated from one week of monitoring produces the following statistics:

- Average number of new connection requests per server: 250 per second
- Average peak number of new connection requests per server: 500 per second

NOTE: A sniffer is a network-analyzing device that captures packets on the network segment to which you attach it. Most sniffers allow you to define filters to collect only the type of traffic that interests you. Later, you can view and evaluate the accumulated information. In this example, you want the sniffer to collect all TCP packets with the SYN flag set arriving at ethernet3 and destined for one of the four Web servers in the DMZ. You might want to continue running the sniffer at regular intervals to see if there are traffic patterns based on the time of day, days of the week, the time of month, or the season of the year. For example, in some organizations, traffic might increase dramatically during a critical event. Significant changes probably warrant adjusting the various thresholds.

Based on this information, you set the following SYN flood protection parameters for zone_external, as shown in Table 105 on page 830.

Table 105: SYN Flood Protection Parameters

Parameter	Value	Reason for Each Value
Attack Threshold	625 packets per second (pps)	This is 25% higher than the average peak number of new connection requests per second per server, which is unusual for this network environment. When the number of SYN packets per second for any one of the four Web servers exceeds this number, the device begins proxying new connection requests to that server. (In other words, beginning with the 626th SYN packet to the same destination address and port number in one second, the device begins proxying connection requests to that address and port number.)
Alarm Threshold	250 pps	When the device proxies 251 new connection requests in one second, it makes an alarm entry in the event log. By setting the alarm threshold somewhat higher than the attack threshold, you can avoid alarm entries for traffic spikes that only slightly exceed the attack threshold.
Source Threshold	25 pps	When you set a source threshold, the device tracks the source IP address of SYN packets, regardless of the destination address and port number. (Note that this source-based tracking is separate from the tracking of SYN packets based on destination address and destination port number that constitutes the basic SYN flood protection mechanism.) In the one week of monitoring activity, you observed that no more than 1/25 of new connection requests for all servers came from any one source within a one-second interval. Therefore, connection requests exceeding this threshold are unusual and provide sufficient cause for the device to execute its proxying mechanism. (25 pps is 1/25 of the attack threshold, which is 625 pps.) If the device tracks 25 SYN packets from the same source IP address, beginning with the 26th packet, it rejects all further SYN packets from that source for the remainder of that second and the next second as well.

Table 105: SYN Flood Protection Parameters *(continued)*

Parameter	Value	Reason for Each Value
Destination Threshold	0 pps	When you set a destination threshold, the device runs a separate tracking of only the destination IP address, regardless of the destination port number. Because the four Web servers only receive HTTP traffic (destination port 80)—no traffic to any other destination port number reaches them—setting a separate destination threshold offers no additional advantage.
Timeout	20 seconds	The default value of 20 seconds is a reasonable length of time to hold incomplete connection requests.
Parameter	Value	Reason for Each Value
Attack Threshold	625 packets per second (pps)	This is 25% higher than the average peak number of new connection requests per second per server, which is unusual for this network environment. When the number of SYN packets per second for any one of the four Web servers exceeds this number, the device begins proxying new connection requests to that server. (In other words, beginning with the 626th SYN packet to the same destination address and port number in one second, the device begins proxying connection requests to that address and port number.)

* Half-completed connection requests are incomplete three-way handshakes. A three-way handshake is the initial phase of a TCP connection. It consists of a TCP segment with the SYN flag set, a response with the SYN and ACK flags set, and a response to that with the ACK flag set.

You can use either J-Web or the CLI configuration editor to configure SYN flood protection parameters.

This topic covers:

- J-Web Configuration on page 831
- CLI Configuration on page 834
- Related Topics on page 835

J-Web Configuration

To set interfaces:

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Interfaces, click **Configure** or **Edit**.
3. Next to Interface, click Add new entry.
4. In the Interface name box, type **ge-0/0/0**.
5. Next to Unit, click Add new entry.
6. Next to Interface unit number, type **0**.
7. Next to Inet, select the check box and click Configure.
8. Next to Address, click Add new entry.

9. Next to Source, type **1.2.2.1/24** and click OK.
10. To configure another interface, **fe-1/0/0**, and address, **1.1.1.1/24**, repeat Step b through i and click OK.
11. To save and commit the configuration, click Commit.

To configure a zone and assign an interface:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone_dmz**.
6. Next to Interfaces, click Add new entry.
7. In the Interface unit box, type **ge-0/0/0.0** and click OK.
8. To configure another security zone, **zone_external** and assign an interface **fe-1/0/0.0**, repeat Step 4 through Step 7 and click OK.
9. To save and commit the configuration, click Commit.

To define an address:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone_dmz**.
6. Next to Address book, click Configure.
7. Next to Address, click Add new entry.
8. In the Address name box, type **ws1 1.2.2.10/32** and click OK.
9. To configure other address entries such as **ws2 1.2.2.20/32**, **ws3 1.2.2.30/32**, **ws4 1.2.2.40/32**, repeat Step 7 through Step 8 and click OK.
10. Next to Address set, click Add new entry.
11. In the Address set name box, type **web_servers**.
12. Next to Address, click Add new entry.
13. In the Address name box, type **ws1**.
14. To configure other address-set entries such as **ws2**, **ws3**,**ws4**, repeat Step l through Step m and click OK.
15. To save and commit the configuration, click Commit.

To configure a policy:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Policies, select the check box and click **Configure** or **Edit**.
4. Next to Policy, click Add new entry.
5. In the From zone name box, type **zone_external**.
6. In the To zone name box, type **zone_dmz** and click OK.
7. Under the From zone name column, click **private**.
8. Next to Policy, click Add new entry.
9. In the Policy name box, type **id_1**.
10. Select the Match check box.
11. Select the Then check box.
12. Next to Match, click Configure.
13. Next to Source address choice list, select **Source address**.
14. Next to Source address, click Add new entry.
15. From the Value keyword list, select **any** and click OK.
16. From the Destination address choice list, select **Destination address**.
17. Next to Destination address, click Add new entry.
18. From the Value keyword list, select **Enter Specific Value**.
19. In the Address box, type **web_servers** and click OK.
20. From the Application choice list, select Application.
21. Next to Application, click Add new entry.
22. In the Value keyword box, type **junos-http** and click OK.
23. Next to Then, click Configure.
24. Next to Action, select **permit** and click OK.
25. To save and commit the configuration, click Commit.

To configure screen options:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **zone_external-syn-flood**.
6. Next to Tcp, click Configure.
7. Next to Syn flood box, select the check box and click Configure.
8. In the Alarm threshold box, type **250**.

9. In the Attack threshold box, type **625**.
10. In the Source threshold box, type **25**.
11. In the Timeout box, type **20**.
12. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone_external**.
6. In the Screen box, type **zone_external-syn-flood** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

1. Set interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 1.2.2.1/24
user@host# set interfaces fe-1/0/0 unit 0 family inet address 1.1.1.1/24
user@host# set security zones security-zone zone_dmz interfaces ge-0/0/0.0
user@host# set security zones security-zone zone_external interfaces fe-1/0/0.0
```

2. Define addresses.

```
user@host# set security zones security-zone zone_dmz address-book address
  ws1 1.2.2.10/32
user@host# set security zones security-zone zone_dmz address-book address
  ws2 1.2.2.20/32
user@host# set security zones security-zone zone_dmz address-book address
  ws3 1.2.2.30/32
user@host# set security zones security-zone zone_dmz address-book address
  ws4 1.2.2.40/32
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws1
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws2
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws3
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws4
```

3. Configure policy.

```
user@host# set security policies from-zone zone_external to-zone zone_dmz
  policy id_1 match source-address any
user@host# set security policies from-zone zone_external to-zone zone_dmz
  policy id_1 match destination-address web_servers
```

```
user@host# set security policies from-zone zone_external to-zone zone_dmz
    policy id_1 match application junos-http
user@host# set security policies from-zone zone_external to-zone zone_dmz
    policy id_1 then permit
```

4. Configure SCREEN options.

```
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
    alarm-threshold 250
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
    attack-threshold 625
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
    source-threshold 25
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
    timeout 20
user@host# set security zones security-zone zone_external screen
    zone_external-syn-flood
```

Related Topics

- Enabling SYN Flood Protection on page 835

Enabling SYN Flood Protection

A SYN flood occurs when a host becomes so overwhelmed by SYN segments initiating incomplete connection requests that it can no longer process legitimate connection requests.

Before You Begin

For background information, read "Understanding SYN Flood Attacks" on page 823.

To enable the SYN flood protection screen option and define its parameters, the JUNOS CLI configuration editor. The specified zone is where a flood might originate.

```
user@host# set security screen zone-syn-flood tcp syn-flood timeout 20
user@host# set security zones security-zone zone screen zone-syn-flood
user@host# set zone zone screen syn-flood
```

Related Topics

- Enabling SYN Cookie Protection on page 837
- Enabling ICMP Flood Protection on page 840
- Enabling UDP Flood Protection on page 843
- Enabling Protection Against a Land Attack on page 845

Understanding SYN Cookie Protection

SYN Cookie is a stateless SYN proxy mechanism you can use in conjunction with the defenses against a SYN flood attack.

Before You Begin

For background information, read:

- Network DoS Attacks Overview on page 823
- Understanding SYN Flood Attacks on page 823

As with traditional SYN proxying, SYN Cookie is activated when the SYN flood attack threshold is exceeded. However, because SYN Cookie is stateless, it does not set up a session or policy and route lookups upon receipt of a SYN segment, and it maintains no connection request queues. This dramatically reduces CPU and memory usage and is the primary advantage of using SYN Cookie over the traditional SYN proxying mechanism.

When SYN Cookie is enabled on JUNOS Software and becomes the TCP-negotiating proxy for the destination server, it replies to each incoming SYN segment with a SYN/ACK containing an encrypted cookie as its Initial Sequence Number (ISN). The cookie is an MD5 hash of the original source address and port number, destination address and port number, and ISN from the original SYN packet. After sending the cookie, JUNOS Software drops the original SYN packet and deletes the calculated cookie from memory. If there is no response to the packet containing the cookie, the attack is noted as an active SYN attack and is effectively stopped.

If the initiating host responds with a TCP packet containing the cookie + 1 in the TCP ACK field, JUNOS Software extracts the cookie, subtracts 1 from the value, and recomputes the cookie to validate that it is a legitimate ACK. If it is legitimate, JUNOS Software starts the TCP proxy process by setting up a session and sending a SYN to the server containing the source information from the original SYN. When JUNOS Software receives a SYN/ACK from the server, it sends ACKs to the server and to the initiation host. At this point the connection is established and the host and server are able to communicate directly.

Figure 147 on page 837 shows how a connection is established between an initiating host and a server when SYN Cookie is active on JUNOS Software.

Figure 147: Establishing a Connection with SYN Cookie Active

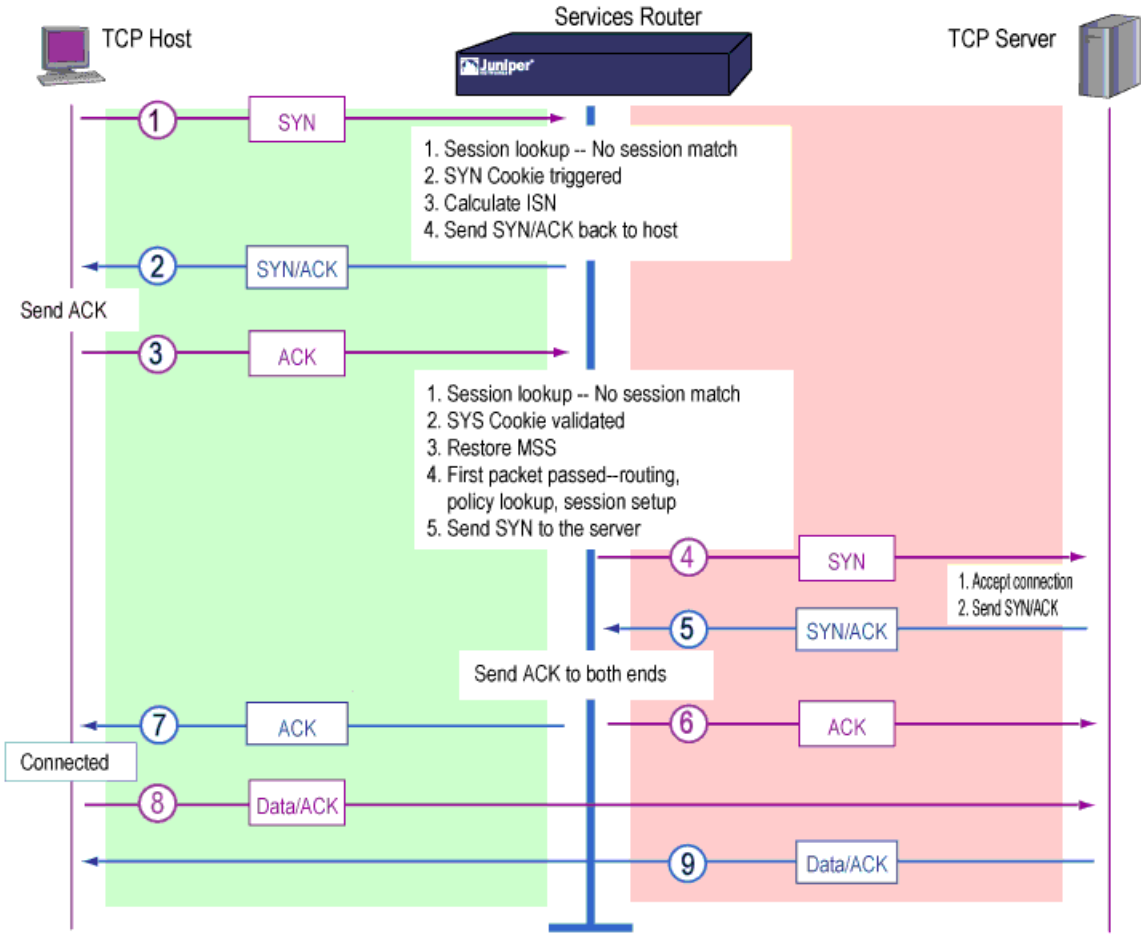

Related Topics

- Enabling SYN Cookie Protection on page 837

Enabling SYN Cookie Protection

SYN Cookie is a stateless SYN proxy mechanism you can use in conjunction with the defenses against a SYN flood attack.

Before You Begin
For background information, read “Understanding SYN Cookie Protection” on page 836.

You can use either J-Web or the CLI configuration editor to enable SYN Cookie, set the SYN flood attack threshold.

This topic covers:

- J-Web Configuration on page 838
- CLI Configuration on page 838
- Related Topics on page 839

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **external-syn-flood**.
6. Next to Tcp, click Configure.
7. Next to Syn flood box, select the check box and click Configure.
8. In the Timeout box, type **20** and click OK.
9. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **external**.
6. In the Screen box, type **external-syn-flood** and click OK.
7. To save and commit the configuration, click Commit.

To configure flow:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Flow, click **Configure**.
4. From the Syn flood protection mode choice list, select **syn-cookie** and click OK.
5. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option external-syn-flood tcp syn-flood timeout
  20
```

```
user@host# set security zones security-zone external screen external-syn-flood
user@host# set security flow syn-flood-protection-mode syn-cookie
```

NOTE: The SYN Cookie feature can only detect and protect against spoofed SYN-Flood attacks, thus minimizing the negative impact to hosts that are secured by JUNOS Software. If an attacker is using a legitimate IP source address, rather than a spoofed IP source, then the SYN-Cookie mechanism does not stop the attack.

Related Topics

- Enabling SYN Flood Protection on page 835
- Enabling ICMP Flood Protection on page 840
- Enabling UDP Flood Protection on page 843
- Enabling Protection Against a Land Attack on page 845

Understanding ICMP Flood Attacks

An ICMP flood typically occurs when ICMP echo requests overload its victim with so many requests that it expends all its resources responding until it can no longer process valid network traffic.

Before You Begin

For background information, read "Network DoS Attacks Overview" on page 823.

When enabling the ICMP flood protection feature, you can set a threshold that once exceeded invokes the ICMP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, JUNOS Software ignores further ICMP echo requests for the remainder of that second plus the next second as well. See Figure 148 on page 840.

NOTE: An ICMP flood can consist of any type of ICMP message. Therefore, JUNOS Software monitors all ICMP message types, not just echo requests.

Figure 148: ICMP Flooding

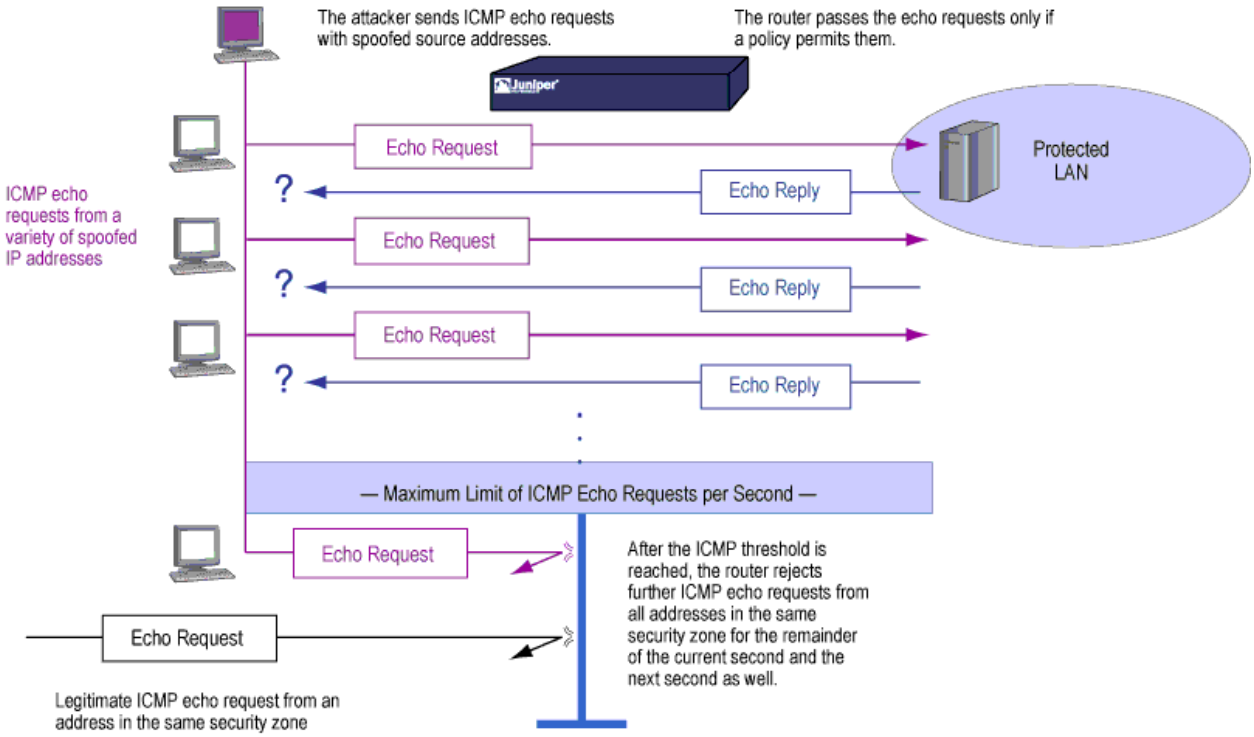

Related Topics

- Enabling ICMP Flood Protection on page 840

Enabling ICMP Flood Protection

An ICMP flood typically occurs when ICMP echo requests overload its victim with so many requests that it expends all its resources responding until it can no longer process valid network traffic.

Before You Begin
For background information, read “Understanding ICMP Flood Attacks” on page 839.

To enable ICMP flood protection, you can use either J-Web or the CLI configuration editor. The specified zone is where a flood might originate.

NOTE: The value unit is ICMP packets per second. The default value is 1000 packets per second.

This topic covers:

- J-Web Configuration on page 841
- CLI Configuration on page 841
- Related Topics on page 841

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **1000-Icmp-flood**.
6. Next to Icmp, click Configure.
7. Next to Flood, select the check box and click Configure.
8. In the Threshold box, type **1000** and click OK.
9. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **1000-Icmp-flood** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option 1000-icmp-flood icmp flood threshold
  1000
user@host# set security zones security-zone zone screen 1000-icmp-flood
```

Related Topics

- Enabling SYN Flood Protection on page 835
- Enabling SYN Cookie Protection on page 837
- Enabling UDP Flood Protection on page 843
- Enabling Protection Against a Land Attack on page 845

Understanding UDP Flood Attacks

Similar to the ICMP flood, UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the victim to the point that the victim can no longer handle valid connections.

Before You Begin

For background information, read:

- Network DoS Attacks Overview on page 823
- Understanding ICMP Flood Attacks on page 839

After enabling the UDP flood protection feature, you can set a threshold that, once exceeded, invokes the UDP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the number of UDP datagrams from one or more sources to a single destination exceeds this threshold, JUNOS Software ignores further UDP datagrams to that destination for the remainder of that second plus the next second as well. See Figure 149 on page 842.

Figure 149: UDP Flooding

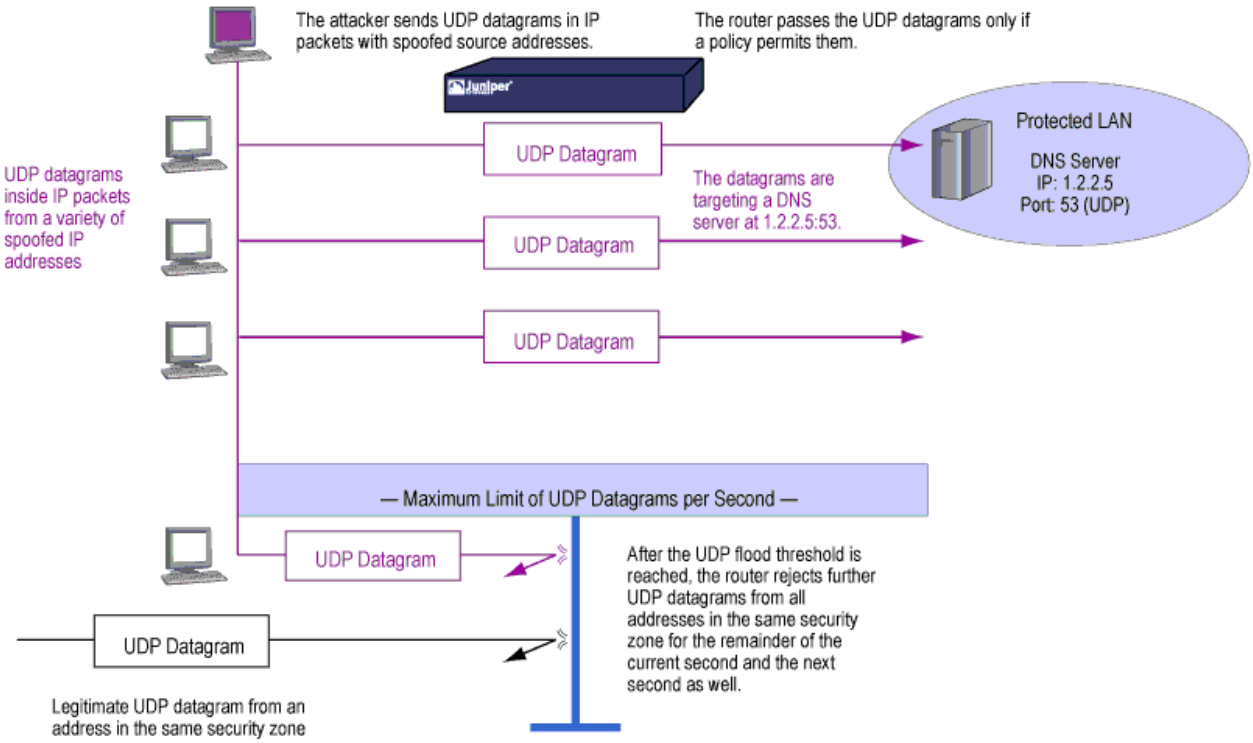

Related Topics

- Enabling Protection Against a Land Attack on page 845

Enabling UDP Flood Protection

UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the victim to the point that the victim can no longer handle valid connections.

Before You Begin
For background information, read "Understanding UDP Flood Attacks" on page 842.

You can use either J-Web or the CLI configuration editor to enable UDP flood protection. The specified zone is where a flood might originate.

NOTE: The value unit is UDP packets per second. The default value is 1000 packets per second.

This topic covers:

- J-Web Configuration on page 843
- CLI Configuration on page 844
- Related Topics on page 844

J-Web Configuration

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **external**.
6. In the Screen box, type **external-udp-flood** and click OK.
7. To save and commit the configuration, click Commit.

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.

3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **1000-udp-flood**.
6. Next to Udp, click Configure.
7. Next to Flood, select the check box and click Configure.
8. In the Threshold box, type **1000** and click OK.
9. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security zones security-zone external screen external-udp-flood
user@host# set security screen ids-option 1000-udp-flood udp flood threshold 1000
```

Related Topics

- Enabling SYN Flood Protection on page 835
- Enabling SYN Cookie Protection on page 837
- Enabling ICMP Flood Protection on page 840
- Enabling Protection Against a Land Attack on page 845

Understanding Land Attacks

Combining a SYN attack with IP spoofing, a land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.

Before You Begin

For background information, read:

- Network DoS Attacks Overview on page 823
- IP Spoofing on page 782

The receiving system responds by sending the SYN-ACK packet to itself, creating an empty connection that lasts until the idle timeout value is reached. Flooding a system with such empty connections can overwhelm the system, causing a denial of service. See Figure 150 on page 845.

Figure 150: Land Attack

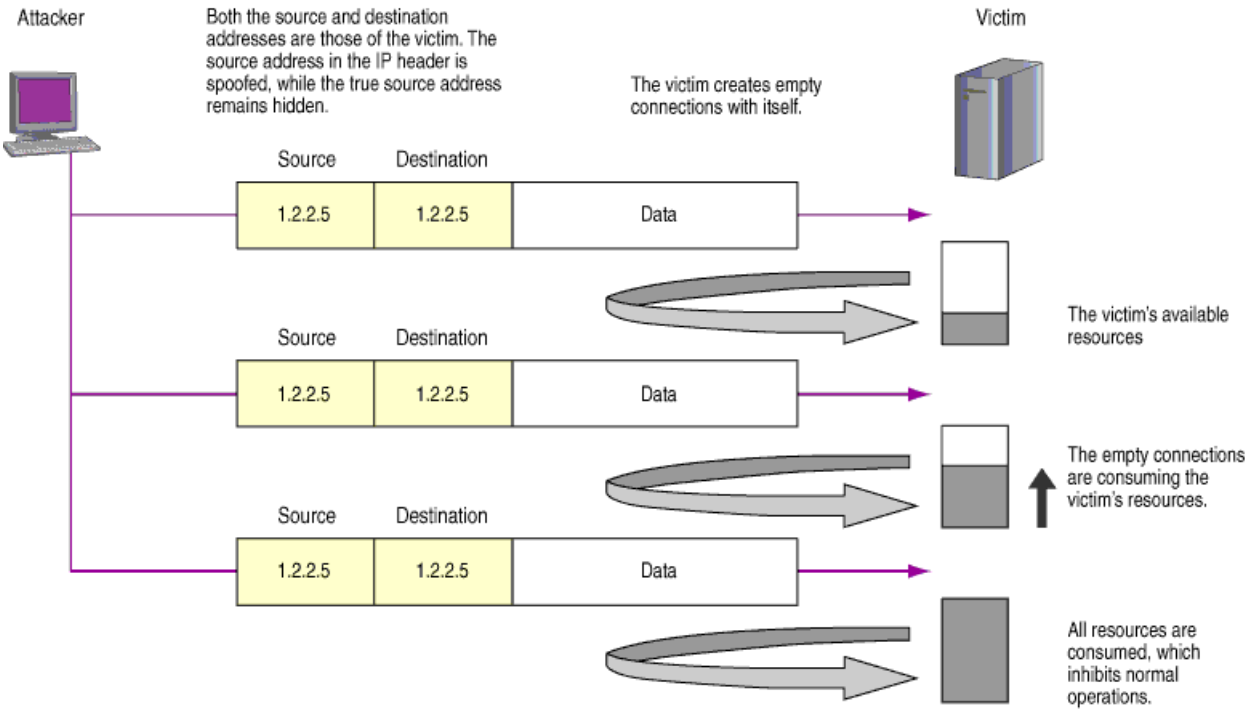

When you enable the screen option to block land attacks, JUNOS Software combines elements of the SYN flood defense and IP spoofing protection to detect and block any attempts of this nature.

Related Topics

- Enabling Protection Against a Land Attack on page 845

Enabling Protection Against a Land Attack

Combining a SYN attack with IP spoofing, a land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.

Before You Begin
For background information, read "Understanding Land Attacks" on page 844.

You can use either J-Web or the CLI configuration editor to enable protection against a land attack.

This topic covers:

- J-Web Configuration on page 846
- CLI Configuration on page 846
- Related Topics on page 846

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **land**.
6. Next to Tcp, click Configure.
7. Next to Land, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **land** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen land tcp land
user@host# set security zones security-zone zone screen land
```

Related Topics

- Enabling SYN Flood Protection on page 835
- Enabling SYN Cookie Protection on page 837
- Enabling ICMP Flood Protection on page 840
- Enabling UDP Flood Protection on page 843

OS-Specific DoS Attacks Overview

If an attacker not only identifies the IP address and responsive port numbers of an active host but also its operating system (OS), instead of resorting to brute-force attacks, the attacker can launch more elegant attacks that can produce one- or two-packet " kills."

Before You Begin

For background information, read "Denial-of-Service Attack Overview" on page 807.

OS-specific DoS attacks, including ping of death attacks, teardrop attacks, and WinNuke attacks, can cripple a system with minimum effort. If JUNOS Software is protecting hosts susceptible to these attacks, you can configure JUNOS Software to detect these attacks and block them before they reach their target.

Related Topics

- Understanding Ping of Death Attacks on page 847
- Understanding Teardrop Attacks on page 849
- Understanding WinNuke Attacks on page 852

Understanding Ping of Death Attacks

OS-specific DoS attacks such as ping of death attacks can cripple a system with minimum effort.

Before You Begin

For background information, read "OS-Specific DoS Attacks Overview" on page 847.

The maximum allowable IP packet size is 65,535 bytes, including the packet header, which is typically 20 bytes long. An ICMP echo request is an IP packet with a pseudo header, which is 8 bytes long. Therefore, the maximum allowable size of the data area of an ICMP echo request is 65,507 bytes (65,535 - 20 - 8 = 65,507).

However, many ping implementations allow the user to specify a packet size larger than 65,507 bytes. A grossly oversized ICMP packet can trigger a range of adverse system reactions such as denial of service (DoS), crashing, freezing, and rebooting.

When you enable the ping of death screen option, JUNOS Software detects and rejects such oversized and irregular packet sizes even when the attacker hides the total packet size by purposefully fragmenting it. See Figure 151 on page 848.

NOTE: For information about IP specifications, see RFC 791, *Internet Protocol*.For information about ICMP specifications, see RFC 792, *Internet Control Message Protocol*.For information about ping of death attacks, see http://www.insecure.org/sploits/ping-o-death.html.

Figure 151: Ping of Death

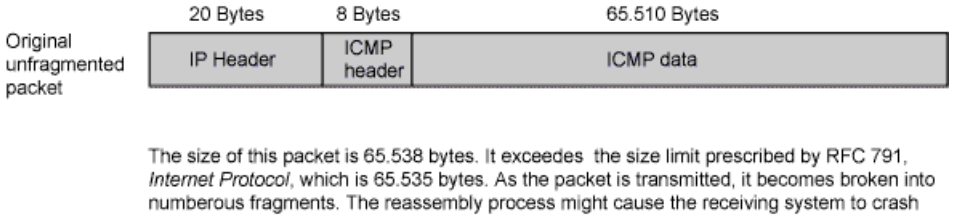

Related Topics

- Enabling Protection Against a Ping of Death Attack on page 848

Enabling Protection Against a Ping of Death Attack

A grossly oversized ICMP packet can trigger a range of adverse system reactions such as denial of service (DoS), crashing, freezing, and rebooting.

Before You Begin
For background information, read "Understanding Ping of Death Attacks" on page 847.

You can use either J-Web or the CLI configuration editor to enable protection against a ping of death attack. The specified zone is where the attack originates.

This topic covers:

- J-Web Configuration on page 848
- CLI Configuration on page 849
- Related Topics on page 849

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.

5. In the Name box, type **ping-death**.
6. Next to Icmp, click Configure.
7. Next to Ping death, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **ping-death** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option ping-death icmp ping-death
user@host# set security zones security-zone zone screen ping-death
```

Related Topics

- Enabling Protection Against a Teardrop Attack on page 851
- Enabling Protection Against a WinNuke Attack on page 853

Understanding Teardrop Attacks

OS-specific DoS attacks such as teardrop attacks can cripple a system with minimum effort.

Before You Begin
For background information, read "OS-Specific DoS Attacks Overview" on page 847.

Teardrop attacks exploit the reassembly of fragmented IP packets. In the IP header, one of the fields is the fragment offset field, which indicates the starting position, or offset, of the data contained in a fragmented packet relative to the data of the original unfragmented packet. See Figure 152 on page 850.

Figure 152: Teardrop Attacks

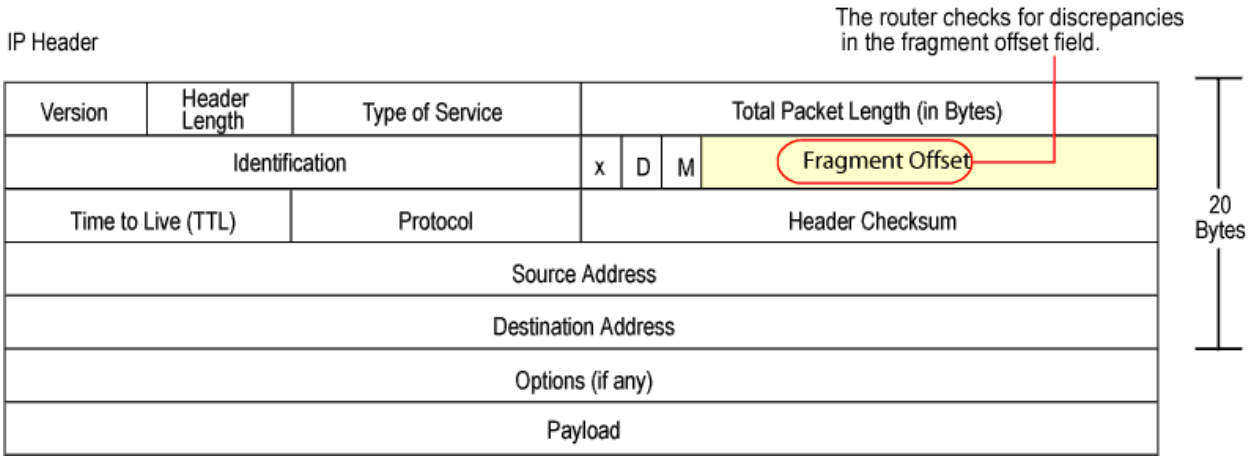

When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash, especially if it is running an older operating system that has this vulnerability. See Figure 153 on page 850.

Figure 153: Fragment Discrepancy

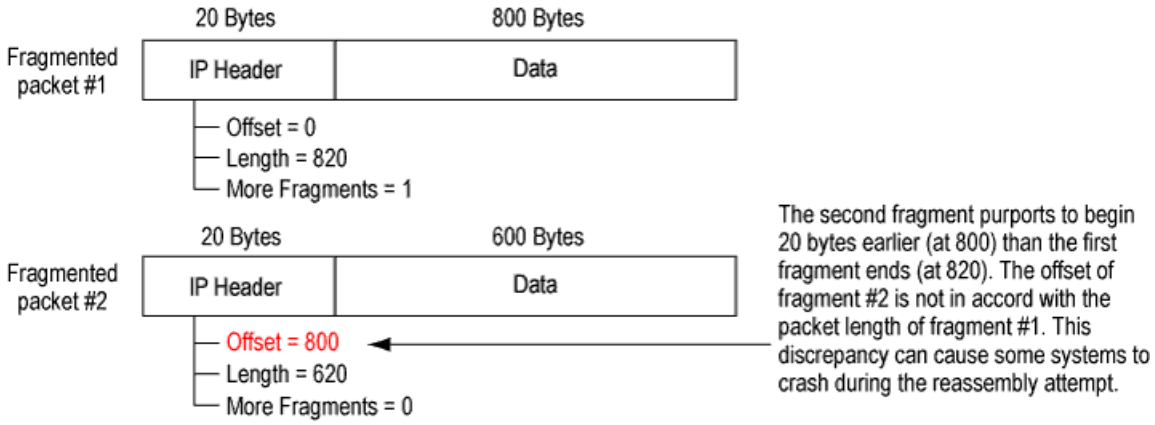

After you enable the teardrop attack screen option, whenever JUNOS Software detects this discrepancy in a fragmented packet, it drops it.

Related Topics

- Enabling Protection Against a Teardrop Attack on page 851

Enabling Protection Against a Teardrop Attack

Teardrop attacks exploit the reassembly of fragmented IP packets.

Before You Begin
For background information, read "Understanding Teardrop Attacks" on page 849.

You can use either J-Web or the CLI configuration editor to enable protection against a Teardrop attack. The specified zone is where the attack originates.

This topic covers:

- J-Web Configuration on page 851
- CLI Configuration on page 852
- Related Topics on page 852

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **tear-drop**.
6. Next to Ip, click Configure.
7. Next to Tear drop, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.

6. In the Screen box, type **tear-drop** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen ids-option tear-drop ip tear-drop
user@host# set security zones security-zone zone screen tear-drop
```

Related Topics

- Enabling Protection Against a Ping of Death Attack on page 848
- Enabling Protection Against a WinNuke Attack on page 853

Understanding WinNuke Attacks

OS-specific DoS attacks such as WinNuke attacks can cripple a system with minimum effort.

Before You Begin
For background information, read "OS-Specific DoS Attacks Overview" on page 847.

WinNuke is a DoS attack targeting any computer on the Internet running Windows. The attacker sends a TCP segment—usually to NetBIOS port 139 with the urgent (URG) flag set—to a host with an established connection (See Figure 154 on page 853). This introduces a NetBIOS fragment overlap, which causes many machines running Windows to crash. After rebooting the attacked machine, the following message appears, indicating that an attack has occurred:

```
An exception OE has occurred at 0028:[address] in VxD MSTCP(01) +
000041AE. This was called from 0028:[address] in VxD NDIS(01) +
00008660. It may be possible to continue normally.
Press any key to attempt to continue.
Press CTRL+ALT+DEL to restart your computer. You will lose any unsaved information
   in all applications.
Press any key to continue.
```

Figure 154: WinNuke Attack Indicators

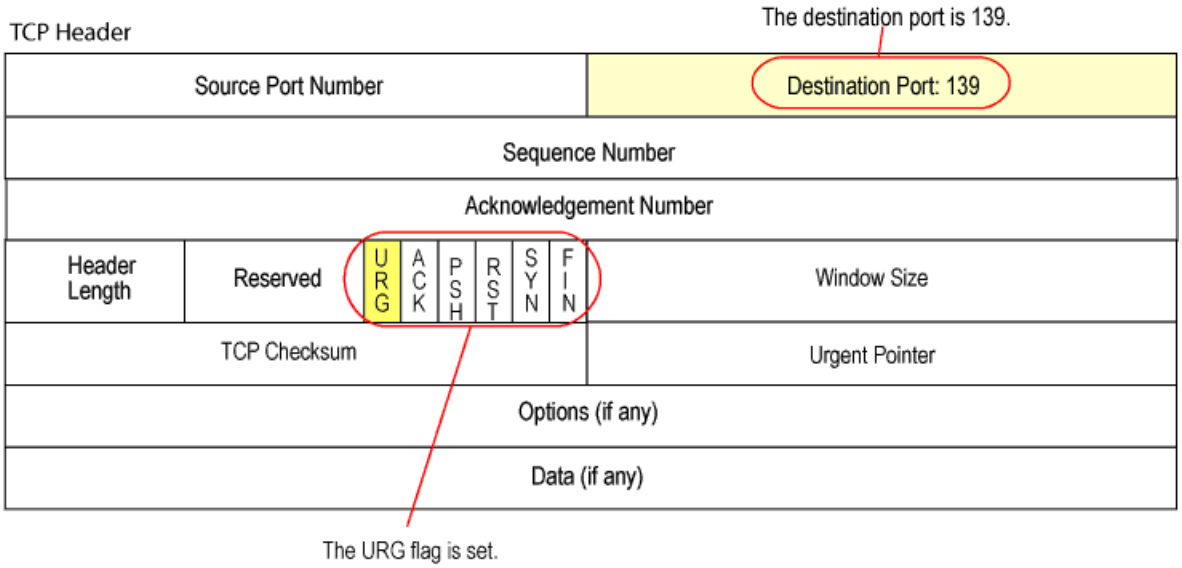

If you enable the WinNuke attack defense screen option, JUNOS Software scans any incoming Microsoft NetBIOS session service (port 139) packets. If JUNOS Software observes that the URG flag is set in one of those packets, it unsets the URG flag, clears the URG pointer, forwards the modified packet, and makes an entry in the event log noting that it has blocked an attempted WinNuke attack.

Related Topics

- Enabling Protection Against a WinNuke Attack on page 853

Enabling Protection Against a WinNuke Attack

WinNuke is a DoS attack targeting any computer on the Internet running Windows.

Before You Begin
For background information, read “Understanding WinNuke Attacks” on page 852.

You can use either J-Web or the CLI configuration editor to enable protection against a WinNuke attack. The specified zone is where the attack originates.

This topic covers:

- J-Web Configuration on page 854
- CLI Configuration on page 854
- Related Topics on page 854

J-Web Configuration

To configure screens:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Screen, click **Configure**.
4. Next to Ids option, click Add new entry.
5. In the Name box, type **winnuke**.
6. Next to Tcp, click Configure.
7. Next to Winnuke, select the check box and click OK.
8. To save and commit the configuration, click Commit.

To configure zones:

1. Select **Configure > CLI Tools > Point and Click CLI**.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Zones, click **Configure**.
4. Next to Security zone, click Add new entry.
5. In the Name box, type **zone**.
6. In the Screen box, type **winnuke** and click OK.
7. To save and commit the configuration, click Commit.

CLI Configuration

```
user@host# set security screen winnuke tcp winnuke
user@host# set security zones security-zone zone screen winnuke
```

Related Topics

- Enabling Protection Against a Ping of Death Attack on page 848
- Enabling Protection Against a Teardrop Attack on page 851

Configuring Firewall Screen Options—Quick Configuration

You can use J-Web Quick Configuration to quickly configure Firewall/NAT flow

Before You Begin

For background information, read

- "Configuring Host Inbound Traffic" on page 53
- "Reconnaissance Deterrence Overview" on page 763
- "Understanding Operating System Probes" on page 773
- "Understanding Attacker Evasion Techniques" on page 779

Figure 155 on page 855 shows the Quick Configuration for Firewall/NAT Screen page.

Figure 155: Quick Configuration Page for Firewall/NAT Screen

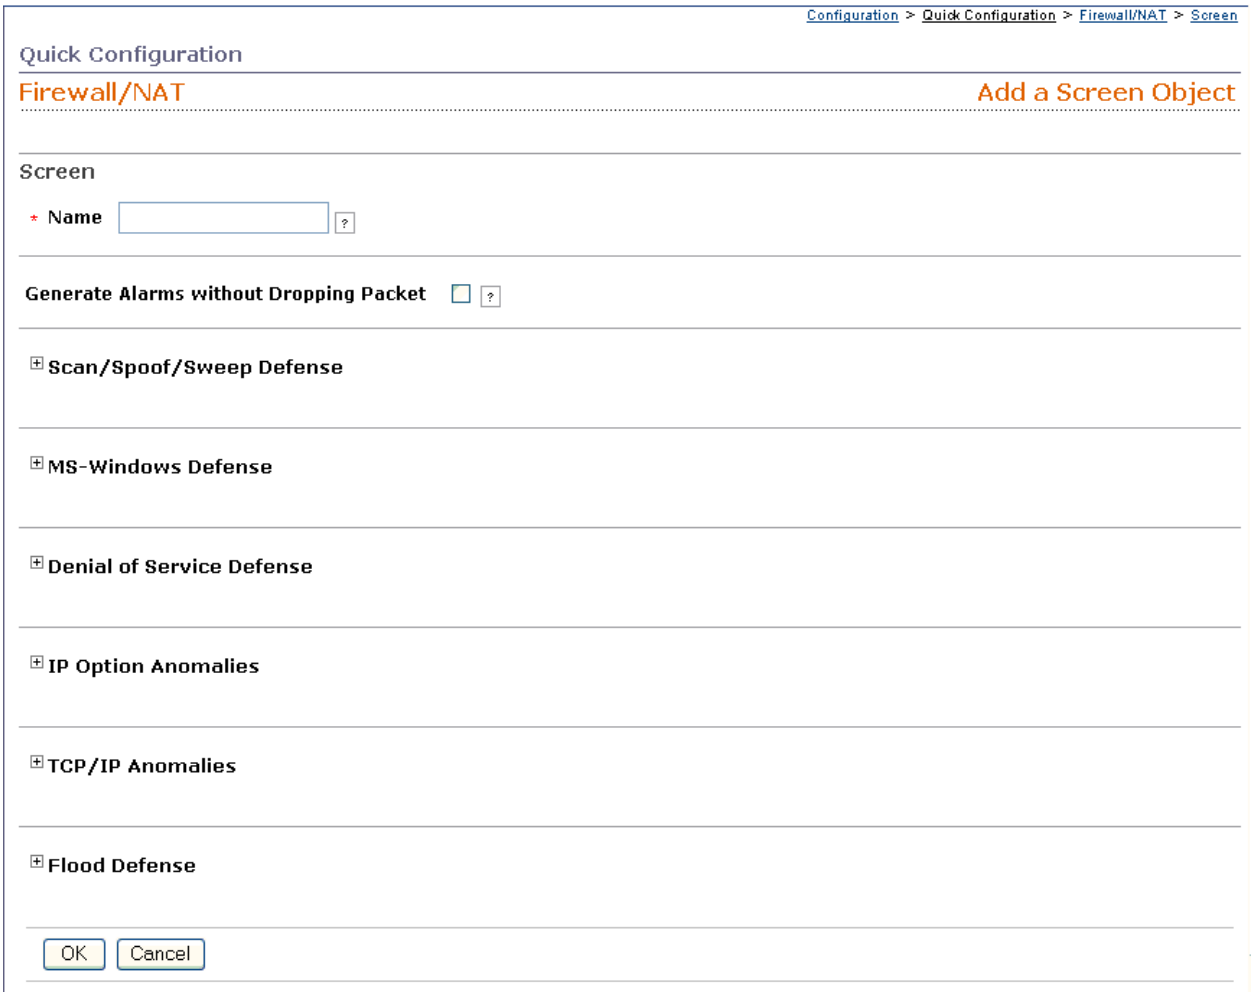

To configure screen options with Quick Configuration:

1. Select **Conigure > Security > Screens**.
2. Click Add to define screen objects; the screen objects page appears as shown in Figure 155 on page 855.
3. Fill in the screen options as shown in Table 106 on page 856.
4. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

Table 106: Firewall/NAT Screen Configuration Options

Field	Function	Action
Screen		
Name	Name of the screen object.	Specify a unique name for the screen object you are defining.
Generate Alarms without Dropping Packets	Generates alarms without dropping packets.	Select this check box to enable alarm generation but do not drop any packets.
Scan/Spoof/Sweep Defense		
IP Address Spoof	Enables IP address spoofing. IP spoofing is when a bogus source address is inserted in the packet header to make the packet appear to come from a trusted source.	Select this check box to enable IP address spoofing.
IP Address Sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.	Select this check box to enable IP address sweep. Configure a time interval (in microseconds). If a remote host sends ICMP traffic to 10 addresses within this interval, an address sweep attack is flagged and further ICMP packets from the remote host are rejected. Valid values are between 1000 and 1000000 microseconds. The default value is 5000 microseconds.
Port Scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.	Select this check box to enable port scanning. Configure a time interval (in microseconds). If a remote host scans 10 ports within this interval, a port scan attack is flagged and further packets from the remote host are rejected. Valid values are between 1000 and 1000000 microseconds. The default value is 5000 microseconds.
MS-Windows Defense		
WinNuke Attack Protection	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a DoS attack targeting any computer on the Internet running Windows.	Select this check box to enable WinNuke attack protection option.

Table 106: Firewall/NAT Screen Configuration Options *(continued)*

Field	Function	Action
Denial of Service Defense		
Land Attack Protection	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.	Select this check box to enable land attack protection option.
Teardrop Attack Protection	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.	Select this check box to enable teardrop protection option.
ICMP Fragment Protection	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.	Select this check box to enable ICMP fragment protection option.
Ping of Death Attack Protection	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).	Select this check box to enable ping of death attack protection option.
Large Size ICMP Packet Protection	Number of large ICMP packets.	Select this check box to enable large (size > 1024) ICMP packet protection option.
Block Fragment Traffic	Number of IP block fragments.	Select this check box to enable IP fragment blocking.
Source IP Based Session Limit	Limits sessions from the same source IP.	Select this check box to enable source IP based session limit. Configure the threshold between 1 and 50000 sessions. The default value is 128 sessions. **NOTE:** For SRX Series devices, the applicable range is 1 through 8000000 sessions per second.
Destination IP Based Session Limit	Limits sessions to the same destination IP.	Select this check box to enable destination IP based session limit. Configure the threshold between 1 and 50000 sessions. The default value is 128 sessions. **NOTE:** For SRX Series devices, the applicable range is 1 through 8000000 sessions per second.
SYN-ACK-ACK Proxy Protection	Number of TCP flags enabled with SYN-ACK-ACK. This is designed to prevent flooding with SYN-ACK-ACK sessions. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, JUNOS 8.5 software rejects further connection requests from that IP address.	Select this check box to enable the SYN-ACK-ACK proxy protection screen option. Configure the threshold value between 1 and 250000 unauthenticated connections. The default value is 512.
IP Option Anomalies		
Bad IP Option	Number of bad options counter.	Select this check box to enable IP with bad option IDs screen option.

Table 106: Firewall/NAT Screen Configuration Options *(continued)*

Field	Function	Action
Record Route Option	Records the IP addresses of the network devices along the path that the IP packet travels.	Select this check box to enable IP with record route option.
Timestamp Option	Records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.	Select this check box to enable IP with timestamp option.
Security Option	Provides a way for hosts to send security.	Select this check box to enable IP with security option.
Stream Option	Provides a way for the 16-bit SATNET stream identifier to be carried through networks that did not support the stream concept.	Select this check box to enable IP with stream option.
Loose Source Route Option	Specifies a partial route list for a packet to take on its journey from source to destination.	Select this check box to enable IP with loose source route option.
Strict Source Route Option	Specifies the complete route list for a packet to take on its journey from source to destination.	Select this check box to enable IP with strict source route option.
Source Route Option	Number of IP addresses of the devices set at the source that an IP transmission is allowed to take along the path on its way to its destination.	Select this check box to enable IP with source route option.
TCP/IP Anomalies		
SYN Fragment Protection	Number of TCP SYN fragments.	Select this check box to enable SYN Fragment option.
SYN and FIN Flags Set Protection	Number of TCP SYN and FIN flags. When you enable this option, JUNOS Software checks if the SYN and FIN flags are set in TCP headers. If it discovers such a header, it drops the packet.	Select this check box to enable SYN and FIN flags Set option.
FIN Flag without ACK Flag Set Protection	Number of TCP FIN flags without the acknowledge (ACK) flag. When you enable this option, JUNOS Software checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.	Select this check box to enable FIN flag without ACK option and FIN Flag Set option.
TCP Packet without Flag Set Protection	Number of TCP headers without flags set. A normal TCP segment header has at least one flag control set.	Select this check box to enable TCP Packet without Flag Set option.
Unknown Protocol Protection	Number of internet protocols (IP) that are unknown.	Select this check box to enable Unknown Protocol Protection option.
Flood Defense		

Table 106: Firewall/NAT Screen Configuration Options *(continued)*

Field	Function	Action
ICMP Flood Protection	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.	Select this check box to enable ICMP Flood Protection option. Configure threshold value for ICMP flood between 1 and 100000 ICMP packets per second (pps). The default value is 1000 pps. **NOTE:** For SRX Series devices, the applicable range is 1 through 4000000 ICMP Packets per second.
UDP Flood Protection	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.	Select this check box to enable UDP Flood Protection option. Configure threshold value for UDP flood between 1 and 100000 UDP packets with same destination address per second (pps). The default value is 1000 pps. **NOTE:** For SRXSeries devices, the applicable range is 1 through 4000000 UDP packets per second.

Table 106: Firewall/NAT Screen Configuration Options *(continued)*

Field	Function	Action
SYN Flood Protection	Attack Threshold—Number of SYN packets per second required to trigger the SYN proxy mechanism. Alarm Threshold—Define the number of half-complete proxy connections per second at which the device makes entries in the event alarm log. Source Threshold—Number of SYN segments received per second from a single source IP address (regardless of the destination IP address and port number) before the device begins dropping connection requests from that source. Destination Threshold—Number of SYN segments received per second for a single destination IP address before the device begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based only on destination IP address, regardless of the destination port number. Timeout—Maximum length of time before a half-completed connection is dropped from the queue. You can decrease the timeout value until you see any connections dropped during normal traffic conditions.	Attack Threshold—Configure a value between 1 and 100000 proxied requests per second. The default value is 200. **NOTE:** For SRX Series devices, the applicable range is 1 through 1000000 proxied requests per second. Alarm Threshold—Configure a value between 1 and 100000 segments received per second for SYN flood alarm. The default value is 512. **NOTE:** For SRX Series devices, the applicable range is 1 through 1000000 segments per second. Source Threshold—Configure a value for SYN flood from the same source between 4 and 100000 segments received per second. The default value is 4000. **NOTE:** For SRX Series devices, the applicable range is 4 through 1000000 segments per second. Destination Threshold—Configure a value for SYN flood to the same destination between 4 and 100000. The default value is 4000. **NOTE:** For SRX Series devices, the applicable range is 4 through 1000000 segments per second. Timeout—Configure a value for SYN attack protection between 1 and 50 seconds. The default value is 20 seconds.

Verifying Application Security Information Using Trace Options

The JUNOS Software trace function provides a tool for applications to write security debugging information to a file. The information that appears in this file is based on criteria you set. You can use this information to analyze security application issues.

The trace function operates in a distributed manner, with each thread writing to its own trace buffer. These trace buffers are then collected at one point, sorted, and written to trace files. Trace messages are delivered using the IPC (InterProcess Communications) protocol. A trace message has a lower priority than that of control

protocol packets such as BGP, OSPF, and IKE and therefore delivery is not considered to be as reliable.

Setting Security Trace Options

This topic covers:

- J-Web Configuration on page 861
- CLI Configuration on page 862
- Example: Show Security Traceoptions Output on page 863

J-Web Configuration

1. Select Configure>CLI Tools>Point and Click CLI.
2. Next to Security, click **Configure** or **Edit**.
3. Next to Traceoptions, click Configure.
4. Select the No remote trace checkbox to disable remote tracing.
5. In the Rate limit edit box, enter parameters to limit the incoming rate of trace messages. The edit field here accepts a numeric value between 0 and 4294967295.
6. Select the Use local files checkbox to write trace messages to a local file. The trace file is saved in the /var/log/directory.
7. In the Filename edit box, enter a name for the file to which you want trace information written. This name is a text string between 1 and 1024 character and it cannot include spaces, / or % characters. If you do not create a name, the default name for this file is “security.”
8. In the Files edit box, enter the maximum number of trace files that can accumulate. This is a numeric value between 2 and 1000 used to limit the amount of trace files created. The default value is 3.
9. In the Match edit box, enter matching criteria in the for information logged to the trace file. This is a regular expression for matching against. Wildcard (*) characters are accepted.
10. In the Size edit box, enter size parameters to limit the maximum size to which a trace file can grow. This is a numeric value from 10240 to 1073741824. Once the file reaches that size, it is compressed and renamed to <filename>0.gz and the next one is named <filename>1.gz, and so on.
11. Select either the Yes or the No Word readable checkbox. If you select Yes, you are allowing any user to read the trace file.
12. To turn traceoptions on and to perform more than one tracing operation, click the Add new entry link beside Flag to include multiple flag commands. You can include the following flags:
    - all — Trace everything
    - compilation — Trace compilation events

- configuration — Trace configuration events
- routing-socket — Trace routing socket events

13. Click OK to return to the configuration page.
14. In the Advanced section, you can enter groups to which to apply these trace option settings or you can enter groups to exclude from these settings. Click the Add new entry link beside the Apply groups or Apply groups except option to enter one or more groups for inclusion or exclusion.
15. When finished, click one of the following buttons:
    - OK — This applies your settings and returns you to the previous level in the configuration hierarchy.
    - Cancel — This clears the settings you have not yet applied and returns you to the previous level in the configuration hierarchy.
    - Refresh — This updates the display with any changes to the configuration made by other users.
    - Commit — This verifies your entries and applies them to the current configuration file running on the routing platform.
    - Discard — This removes settings applied to, or deletes existing statements or identifiers from, the configuration.

CLI Configuration

Use the following commands to set the described trace options:

- Set remote tracing as disabled.

```
user@host # set security traceoptions no-remote-trace
```

- Set the local writing of trace files. (The trace file is saved in the /var/log/ directory.)

```
user@host # set security traceoptions use-local-files
```

- Set file name. If you do not create a name, the default name for this file is "security."

```
user@host # set security traceoptions file <filename>
```

- Set the maximum number of trace files that can accumulate. (The default is 3.)

```
user@host # set security traceoptions file files 3
```

- Set matching criteria for logging data to the trace file. This criteria is a regular expression for matching against. Wildcard (*) characters are accepted.

```
user@host # set security traceoptions file match *thread
```

- Set world-readable or not.

```
user@host # set security traceoptions file world-readable
user@host # set security traceoptions file no-world-readable
```

- Set maximum trace file size. Once the file reaches that size, it is compressed and renamed to <filename>0.gz and the next one is named <filename>1.gz, and so on.

```
user@host # set security traceoptions file size 10240
```

- To turn traceoptions on and to perform more than one tracing operation, set the following flags:

```
user@host # set security traceoptions flag all
user@host # set security traceoptions flag compilation
user@host # set security traceoptions flag configuration
user@host # set security traceoptions flag routing-socket
```

- Set trace options to apply to entered groups or to exclude entered groups. You can enter groups to which to apply trace option settings or you can enter groups to exclude from trace option settings.

```
user@host # set security traceoptions apply-groups <value>
user@host # set security traceoptions apply-groups-except <value>
```

Example: Show Security Traceoptions Output

The following CLI command is used as follows:

```
user@host # show security traceoptions file usp_trace
user@host # show security traceoptions flag all
user@host # show security traceoptions rate-limit 888
```

The output is as follows:

```
Apr 11 16:06:42
  21:13:15.750395:CID-906489336:FPC-01:PIC-01:THREAD_ID-01:PFE:now update
  0x3607edf8df8in 0x3607e8d0
Apr 11 16:06:42
  21:13:15.874058:CID-1529687608:FPC-01:PIC-01:THREAD_ID-01:CTRL:Enter
  Function[util_ssam_handler]
Apr 11 16:06:42
  21:13:15.874485:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default1: Rate limit
  changed to 888
Apr 11 16:06:42
  21:13:15.874538:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default1: Destination
  ID set to 1
Apr 11 16:06:42
  21:13:15.874651:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default2: Rate limit
  changed to 888
Apr 11 16:06:42
  21:13:15.874832:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default2: Destination
  ID set to 1
Apr 11 16:06:42
  21:13:15.874942:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default3: Rate limit
  changed to 888
```

```
Apr 11 16:06:42
  21:13:15.874997:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default3: Destination
  ID set to 1
```

Verifying Application Security Flow Information

For flow trace options, you can define a packet filter using combinations of destination-port, destination-prefix, interface, protocol, source-port, and source-prefix. If the security flow trace flag for a certain module is set, the packet matching the specific packet filter triggers flow tracing and writes debugging information to the trace file.

The following example displays the options you can set by using :

- For the packet filter named filter1, set the destination port “imap” for matching.

  ```
  user@host # set security flow traceoptions packet-filter filter1 destination-port
    imap
  ```

- For the packet filter named filter1, set the destination IPv4 prefix address 1.2.3.4.

  ```
  user@host # set security flow traceoptions packet-filter filter1 destination-prefix
    1.2.3.4
  ```

- For the packet filter named filter1, set the logical interface as fxp0.

  ```
  user@host # set security flow traceoptions packet-filter filter1 interface fxp0
  ```

- For the packet filter named filter1, set the IP protocol for matching as TCP.

  ```
  user@host # set security flow traceoptions packet-filter filter1 protocol tcp
  ```

- For the packet filter named filter1, set the source port for matching as http.

  ```
  user@host # set security flow traceoptions packet-filter filter1 source-port http
  ```

- For the packet filter named filter1, set the source IPv4 prefix address as 5.6.7.8.

  ```
  user@host # set security flow traceoptions packet-filter filter1 source-prefix 5.6.7.8
  ```

Part 10
Chassis Cluster

- Chassis Cluster on page 867

Chapter 28

Chassis Cluster

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- Understanding Chassis Cluster on page 868
- Understanding Chassis Cluster Formation on page 869
- Understanding Redundancy Groups on page 870
- Understanding Monitoring on page 877
- Understanding Redundant Ethernet Interfaces on page 879
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885
- Understanding Failover on page 888
- Hardware Setup for SRX Series Chassis Clusters on page 890
- Hardware Setup for J Series Chassis Clusters on page 894
- What Happens When You Enable Chassis Cluster on page 895
- Creating an SRX Series Chassis Cluster—Overview on page 904
- Creating a J Series Chassis Cluster—Overview on page 906
- Configuring the Control Ports on page 909
- Setting the Node ID and Cluster ID on page 910
- Configuring the Management Interface on page 911
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912
- Configuring Redundant Ethernet Interfaces—Quick Configuration on page 915
- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920
- Configuring Redundancy Groups on page 921
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Control Link Recovery on page 924
- Configuring Interface Monitoring on page 925

- Configuring a Dampening Time Between Back-to-Back Redundancy Group Failovers on page 926
- Initiating a Manual Redundancy Group Failover on page 927
- Configuring Conditional Route Advertising on page 929
- Verifying the Chassis Cluster Configuration on page 931
- SNMP Failover Traps on page 938
- Upgrading Devices in a Chassis Cluster on page 939
- Disabling Chassis Cluster on page 942
- Chassis Cluster Configuration Scenarios on page 942

Understanding Chassis Cluster

Chassis clustering provides network node redundancy by grouping a pair of the same kind of supported SRX Series devices or J Series devices into a cluster. The devices must be running JUNOS Software. The control ports on the respective nodes are connected to form a control plane that synchronizes configuration and kernel state to facilitate the high availability of interfaces and services. Similarly, the data plane on the respective nodes is connected over the fabric ports to form a unified data plane. The fabric link allows for the management of cross-node flow processing and for the management of session redundancy.

The control plane software operates in active/backup mode. When configured as a chassis cluster, the two nodes back up each other, with one node acting as the primary device and the other as the secondary device, ensuring stateful failover of processes and services in the event of system or hardware failure. If the primary device fails, the secondary device takes over processing of traffic.

The data plane software operates in active/active mode. In a chassis cluster, session information is updated as traffic traverses either device and this information is transmitted between the nodes over the fabric link to guarantee that established sessions are not dropped when a failover occurs. In active/active mode, it is possible for traffic to ingress the cluster on one node and egress from the other node.

NOTE: For SRX3000 and SRX5000 line chassis clusters, the fabric interfaces can also be connected over 10-Gigabit Ethernet links.

Chassis cluster functionality includes:

- Resilient system architecture, with a single active control plane for the entire cluster and multiple Packet Forwarding Engines. This architecture presents a single device view of the cluster.
- Synchronization of configuration and dynamic runtime states between nodes within a cluster.
- Monitoring of physical interfaces, and failover if the failure parameters cross a configured threshold.

The different states that a cluster can be in at any given instant are as follows: hold, primary, secondary-hold, secondary, ineligible, and disabled. A state transition can be triggered due to any event, such as interface monitoring, SPU monitoring, failures, and manual failovers.

Related Topics

- Understanding Chassis Cluster Formation on page 869

Understanding Chassis Cluster Formation

To form a chassis cluster, a pair of the same kind of supported SRX Series devices or J Series devices are combined to act as a single system that enforces the same overall security. For SRX5600 and SRX5800 chassis clusters, the placement and type of Services Processing Cards (SPCs) must match in the two clusters. For SRX3400 and SRX3600 chassis clusters, the placement and type of SPCs, Input/Output Cards (IOCs), and Network Processing Cards (NPCs) must match in the two devices.

For J Series chassis clusters, although the devices must be the same kind, they can contain different Physical Interface Modules (PIMs).

When a device joins a cluster, it becomes a node of that cluster. With the exception of unique node settings and management IP addresses, nodes in a cluster share the same configuration.

You can deploy up to 15 chassis clusters in a Layer 2 domain. Clusters and nodes are identified in the following way:

- A cluster is identified by a *cluster ID* (cluster-id) specified as a number from 1 through 15.
- A cluster node is identified by a *node ID* (node) specified as a number from 0 to 1.

Before You Begin
Read “Understanding Chassis Cluster” on page 868

Related Topics

- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885
- Setting the Node ID and Cluster ID on page 910
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912

- What Happens When You Enable Chassis Cluster on page 895
- Understanding Failover on page 888

Understanding Redundancy Groups

Chassis clustering provides high availability of interfaces and services through redundancy groups and primacy within groups.

Before You Begin
Read “Understanding Chassis Cluster Formation” on page 869

This topic includes:

- About Redundancy Groups on page 870
- Redundancy Group 0: Routing Engines on page 871
- Redundancy Groups 1 Through 128 on page 871
- Redundancy Group Interface Monitoring on page 874
- Redundancy Group IP Address Monitoring on page 875
- Related Topics on page 877

About Redundancy Groups

A redundancy group is an abstract construct that includes and manages a collection of objects. A redundancy group contains objects on both nodes. A redundancy group is primary on one node and backup on the other at any time. When a redundancy group is said to be primary on a node, its objects on that node are active.

Redundancy groups are independent units of failover. Each redundancy group fails over from one node to the other independent of other redundancy groups. When a redundancy group fails over, all its objects fail over together.

Three things determine the primacy of a redundancy group: the priority configured for the node, the node ID (in case of tied priorities), and the order in which the node comes up. If a lower priority node comes up first, then it will assume the primacy for a redundancy group (and will stay as primary if preempt is not enabled).

A chassis cluster can include many redundancy groups, some of which might be primary on one node and some of which might be primary on the other. Alternatively, all redundancy groups can be primary on a single node. One redundancy group's primacy does not affect another redundancy group's primacy. You can create up to 128 redundancy groups.

NOTE: The maximum number of redundancy groups is equal to the number of redundant Ethernet interfaces that you configure.

You can configure redundancy groups to suit your deployment. You configure a redundancy group to be primary on one node and backup on the other node. You specify the node on which the group is primary by setting priorities for both nodes within a redundancy group configuration. The node with the higher priority takes precedence, and the redundancy group's objects on it are active.

If a redundancy group is configured so that both nodes have the same priority, the node with the lowest node ID number always takes precedence, and the redundancy group is primary on it. In a two-node cluster, node 0 always takes precedence in a priority tie.

Redundancy Group 0: Routing Engines

When you initialize a device in chassis cluster mode, the system creates a redundancy group referred to in this document as redundancy group 0. Redundancy group 0 manages the primacy and failover between the Routing Engines on each node of the cluster. As is the case for all redundancy groups, redundancy group 0 can be primary on only one node at a time. The node on which redundancy group 0 is primary determines which Routing Engine is active in the cluster. A node is considered the primary node of the cluster if its Routing Engine is the active one.

The redundancy group 0 configuration specifies the priority for each node. Below is how redundancy group 0 primacy is determined. Note that the three seconds value is the interval if the default heartbeat-threshold and heartbeat-interval values are used.

- The node that comes up first (at least three seconds prior to the other node) is the primary node.
- If both nodes come up at the same time (or within three seconds of each other):
  - The node with the higher configured priority is the primary node.
  - If there is a tie (either because the same value was configured or because default settings were used), the node with the lower node ID (node 0) is the primary node.

The above priority scheme applies to redundancy groups *x* (redundancy groups numbered 1 through 128) as well, provided preempt is not configured.

You cannot enable preemption for redundancy group 0. If you want to change the primary node for redundancy group 0, you must do a manual failover.

CAUTION: Be cautious and judicious in your use of redundancy group 0 manual failovers. A redundancy group 0 failover implies a Routing Engine failover, in which case all processes running on the primary node are killed and then spawned on the new master Routing Engine. This failover could result in loss of state, such as routing state, and degrade performance by introducing system churn.

Redundancy Groups 1 Through 128

You can configure one or more redundancy groups numbered 1 through 128, referred to in this chapter as redundancy group *x*. The maximum number of redundancy

groups is equal to the number of redundant Ethernet interfaces that you configure (for more information, see Table 109 on page 880). Each redundancy group *x* acts as an independent unit of failover and is primary on only one node at a time.

Each redundancy group *x* contains one or more redundant Ethernet interfaces. A redundant Ethernet interface is a pseudointerface that contains a pair of physical Gigabit Ethernet interfaces or a pair of Fast Ethernet interfaces. If a redundancy group is active on node 0, then the child links of all the associated redundant Ethernet interfaces on node 0 are active. If the redundancy group fails over to node 1, then the child links of all redundant Ethernet interfaces on node 1 become active.

On SRX Series and J Series chassis clusters, you can configure multiple redundancy groups to load-share traffic across the cluster. For example, you can configure some redundancy groups *x* to be primary on one node and some redundancy groups *x* to be primary on the other node. You can also configure a redundancy group *x* in a one-to-one relationship with a single redundant Ethernet interface to control which interface traffic flows through.

The traffic for a redundancy group is processed on the node where the redundancy group is active. Because more than one redundancy group can be configured, it is possible that the traffic from some redundancy groups will be processed on one node while the traffic for other redundancy groups is processed on the other node (depending on where the redundancy group is active). Multiple redundancy groups make it possible for traffic to arrive over an interface of one redundancy group and egress over an interface that belongs to another redundancy group. In this situation, the ingress and egress interfaces might not be active on the same node. When this happens, the traffic is forwarded over the fabric link to the appropriate node.

When you configure a redundancy group *x*, you must specify a priority for each node to determine the node on which the redundancy group *x* is primary. The node with the higher priority is selected as primary. The primacy of a redundancy group *x* can fail over from one node to the other. When a redundancy group *x* fails over to the other node, its redundant Ethernet interfaces on that node are active and their interfaces are passing traffic.

Table 107 on page 872 gives an example of redundancy group *x* in an SRX Series chassis cluster and indicates the node on which the group is primary. It shows the redundant Ethernet interfaces and their interfaces configured for redundancy group *x*.

NOTE: SRX210 devices have both Gigabit Ethernet ports and Fast Ethernet ports.

Table 107: Redundancy Groups Example for an SRX Series Chassis Cluster (SRX3000 and SRX5000 Lines)

Group	Primary	Priority	Objects	Interface	Interface
Redundancy group 0	Node 0	Node 0: 254	Routing Engine on node 0	—	—
		Node 1: 2	Routing Engine on node 1	—	—

Table 107: Redundancy Groups Example for an SRX Series Chassis Cluster (SRX3000 and SRX5000 Lines) *(continued)*

Group	Primary	Priority	Objects	Interface	Interface
Redundancy group 1	Node 0	Node 0: 254	Redundant Ethernet interface 0	ge-1/0/0	ge-23/0/0
		Node 1: 2	Redundant Ethernet interface 1	ge-1/3/0	ge-23/3/0
Redundancy group 2	Node 1	Node 0: 2	Redundant Ethernet interface 2	ge-2/0/0	ge-24/0/0
		Node 1: 254	Redundant Ethernet interface 3	ge-2/3/0	ge-24/3/0
Redundancy group 3	Node 0	Node 0: 254	Redundant Ethernet interface 4	ge-3/0/0	ge-25/0/0
		Node 1: 2	Redundant Ethernet interface 5	ge-3/3/0	ge-25/3/0

As the example for an SRX Series chassis cluster in Table 107 on page 872 shows:

- The Routing Engine on node 0 is active because redundancy group 0 is primary on node 0. (The Routing Engine on node 1 is passive, serving as backup.)
- Redundancy group 1 is primary on node 0. Interfaces ge-1/0/0 and ge-1/3/0 belonging to redundant Ethernet interface 0 and redundant Ethernet interface 1 are active and handling traffic.
- Redundancy group 2 is primary on node 1. Interfaces ge-24/0/0 and ge-24/3/0 belonging to redundant Ethernet interface 2 and redundant Ethernet interface 3 are active and handling traffic.
- Redundancy group 3 is primary on node 0. Interfaces ge-3/0/0 and ge-3/3/0 belonging to redundant Ethernet interface 4 and redundant Ethernet interface 5 are active and handling traffic.

Table 108 on page 874 gives an example of redundancy groups *x* in a J Series chassis cluster and indicates the node on which each group is primary. It shows the redundant Ethernet interfaces and their interfaces configured for each redundancy group *x*.

Table 108: Redundancy Groups Example for a J Series Chassis Cluster

Group	Primary	Priority	Objects	Interface	Interface
Redundancy group 0	Node 1	Node 0: 50	Routing Engine on node 0	—	—
		Node 1: 100	Routing Engine on node 1	—	—
Redundancy group 1	Node 1	Node 0: 50	Redundant Ethernet interface 0	fe-1/0/0	fe-8/0/0
		Node 1: 100	Redundant Ethernet interface 1	fe-1/0/1	fe-8/0/1
Redundancy group 2	Node 1	Node 0: 50	Redundant Ethernet interface 2	ge-2/0/0	ge-9/0/0
		Node 1: 100	Redundant Ethernet interface 3	ge-2/0/1	ge-9/0/1
Redundancy group 3	Node 0	Node 0: 100	Redundant Ethernet interface 4	ge-3/0/0	ge-10/0/0
		Node 1: 50	Redundant Ethernet interface 5	ge-3/0/1	ge-10/0/1

As the example for a J Series chassis cluster in Table 108 on page 874 shows:

- The Routing Engine on node 1 is active because redundancy group 0 is primary on node 1. (The Routing Engine on node 0 is passive, serving as backup.)
- Redundancy group 1 is primary on node 1. Interfaces fe-8/0/0 and fe-8/0/1 belonging to redundant Ethernet interface 0 and redundant Ethernet interface 1 are active and handling traffic.
- Redundancy group 2 is primary on node 1. Interfaces ge-9/0/0 and ge-9/0/1 belonging to redundant Ethernet interface 2 and redundant Ethernet interface 3 are active and handling traffic.
- Redundancy group 3 is primary on node 0. Interfaces ge-3/0/0 and ge-3/0/1 belonging to redundant Ethernet interface 4 and redundant Ethernet interface 5 are active and handling traffic.

Redundancy Group Interface Monitoring

For a redundancy group to automatically fail over to another node, its interfaces must be monitored. When you configure a redundancy group, you can specify a set of interfaces that the redundancy group is to monitor for status (or “health”) to

determine whether the interface is up or down. A monitored interface can be a child interface of any of its redundant Ethernet interfaces. When you configure an interface for a redundancy group to monitor, you give it a weight.

Every redundancy group has a threshold tolerance value initially set to 255. When an interface monitored by a redundancy group becomes unavailable, its weight is subtracted from the redundancy group's threshold. When a redundancy group's threshold reaches 0, it fails over to the other node. For example, if redundancy group 1 was primary on node 0, on the threshold-crossing event, redundancy group 1 becomes primary on node 1. In this case, all the child interfaces of redundancy group 1's redundant Ethernet interfaces begin handling traffic.

A redundancy group failover occurs because the cumulative weight of the redundancy group's monitored interfaces has brought its threshold value to 0. When the monitored interfaces of a redundancy group on both nodes reach their thresholds at the same time, the redundancy group is primary on the node with the lower node ID, in this case node 0.

NOTE: If you want to dampen the failovers occurring because of interface monitoring failures, use the hold-down-interval statement.

Redundancy Group IP Address Monitoring

NOTE: IP address monitoring is not supported on the backup device in a chassis cluster.

Redundancy group IP address monitoring checks end-to-end connectivity and allows a redundancy group to fail over because of a failure in the ability of a redundant Ethernet (reth) interface to reach a configured IP address. Redundancy groups on the master device in a cluster can be configured to monitor specific IP addresses to determine whether or not an upstream device in the network is reachable. The redundancy group can be configured such that if the monitored IP address becomes unreachable, the redundancy group will fail over to its backup to maintain service. The primary difference between this monitoring feature and interface monitoring is that IP address monitoring allows for failover when the interface is still up but the network device it is connected to is not reachable for some reason. It may be possible under those circumstances for the other node in the cluster to route traffic around the problem.

NOTE: If you want to dampen the failovers occurring because of IP address monitoring failures, use the hold-down-interval statement.

IP address monitoring configuration requires that you set not only the address to monitor and its failover value but also an IP address monitoring threshold value. Only after the IP address monitoring threshold is reached due to monitored address reachability failure will the IP address monitoring value be deducted from the redundant group’s failover threshold. Thus, multiple addresses may not only be

monitored simultaneously, but also monitored in such a way as to weight their importance to maintaining traffic flow. Also, the threshold value of an IP address that is unreachable and then becomes reachable again will be restored to the monitoring threshold. This will not, however, cause a failback unless the preempt option has been enabled.

When configured, the IP address monitoring failover value is considered along with interface monitoring—if set—and built-in failover monitoring, including SPU monitoring, cold-sync monitoring, and NPC monitoring (on supported platforms). The primary IP addresses that should be monitored are router gateway addresses to ensure that valid traffic coming into the services router can be forwarded to the appropriate network router.

The interval to check the reachability of a monitored IP address is once per second. After failing to reach the configured IP address for five consecutive attempts, the IP is determined to be unreachable and the failover value is deducted from the redundancy group's priority threshold. If the recalculated threshold is not 0, the IP address is marked unreachable only on the primary node and is still marked as reachable on the backup node. If the redundancy group threshold reaches 0 and there are unreachable IP addresses, the redundancy groups will continuously fail over and fail back between the nodes until an unreachable IP address becomes reachable or unreachable monitored IP addresses are removed from monitoring.

To prevent the failure of a monitored IP address from causing failover and failback behavior every six seconds, a 180 second (3 minute) pause timer has been included in the IP address monitoring feature. (Note that there is one pause timer per redundancy group and not one per monitored IP address.) When a monitored IP address failure causes a redundancy group to fail over, a 180 second countdown is started on the new secondary node. During that countdown, the failed monitored IP address failover value is re-added to the redundancy group’s threshold, effectively blocking it from causing a failback. If, after three minutes, the IP address is still unreachable, its value will be deducted from the redundancy group’s threshold. If the failure of multiple monitored IP addresses results in a failover, the behavior is the same; the 180 second pause timer is started to prevent an immediate failback.

The maximum number of monitoring IPs that can be configured per cluster is 64 for the SRX5000 line and 32 for the SRX3000 line. Monitoring can be accomplished only if the IP address is reachable on a redundant Ethernet (reth) interface, and IP addresses cannot be monitored over a tunnel. The feature also cannot be used on a chassis cluster running in transparent mode.

Every redundancy group *x* has a threshold tolerance value initially set to 255. When an IP address monitored by a redundancy group *x* becomes unavailable, its weight is subtracted from the redundancy group *x*'s threshold. When a redundancy group *x*'s threshold reaches 0, it fails over to the other node. For example, if redundancy group 1 was primary on node 0, on the threshold-crossing event, redundancy group 1 becomes primary on node 1. In this case, all the child interfaces of redundancy group 1's redundant Ethernet interfaces begin handling traffic.

A redundancy group *x* failover occurs because the cumulative weight of the redundancy group *x*'s monitored IP addresses has brought its threshold value to 0. When the monitored IP addresses of a redundancy group *x* on both nodes reach their thresholds at the same time, the redundancy group *x* is primary on the node with the lower node ID, in this case node 0.

Related Topics

- Understanding Redundant Ethernet Interfaces on page 879
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912
- Configuring Redundant Ethernet Interfaces on page 923
- Initiating a Manual Redundancy Group Failover on page 927

Understanding Monitoring

There are various types of objects to monitor as you work with devices configured as chassis clusters, including global-level objects and objects that are specific to redundancy groups. This section describes the monitoring of global-level objects. For information about monitoring objects specific to redundancy groups, see "Understanding Redundancy Groups" on page 870.

Before You Begin

For background information, read:

- Understanding Chassis Cluster Formation on page 869
- Understanding Redundancy Groups on page 870

The SRX3000 and SRX5000 lines have one or more Services Processing Units (SPUs) that run on a Services Processing Card (SPC). All flow-based services run on the SPU. Other SRX Series devices and all J Series devices have a flow-based forwarding process, *flowd*, which forwards packets through the device.

This topic includes:

- SPU Monitoring on page 877
- Flowd Monitoring on page 878
- Cold-Sync Monitoring on page 878
- Related Topics on page 879

SPU Monitoring

SPU monitoring tracks the health of the SPUs and of the central point (CP). The chassis manager on each SPC monitors the SPUs and the central point, and also maintains the heartbeat with the routing engine (RE) chassisd. In this hierarchical monitoring system, chassisd is the center for hardware failure detection. SPU monitoring is enabled by default.

NOTE: SPU monitoring is supported on the SRX3000 and SRX5000 lines.

Persistent SPU and central point failure on a node is deemed a catastrophic Packet Processing Engine (PFE) failure. In this case, the node's PFE will be disabled in the cluster by reducing the priorities of redundancy groups *x* to 0.

- A central point failure triggers failover to the secondary node. The failed node's PFE, which includes all SPCs and all I/O cards (IOCs), will automatically be restarted. If the secondary central point has failed as well, the cluster will be unable to come up because there will be no primary device. Only the data plane (redundancy group *x*) will be failed over.
- A single, failed SPU causes failover of redundancy group *x* to the secondary node. All IOCs and SPCs on the failed node will be restarted and redundancy group *x* will be failed over to the secondary node. Failover to the secondary node is automatic without the need for user intervention. When the failed (former) primary node has its failing component restored, failback will be determined by the preempt configuration for the redundancy group *x*. The interval for dead SPU detection is 30 seconds.

Flowd Monitoring

Flowd monitoring tracks the health of the flowd process. Flowd monitoring is enabled by default.

NOTE: Flowd monitoring is supported on SRX210 and SRX100 devices. It is not supported on J Series devices.

Persistent flowd failure on a node is deemed a catastrophic Packet Processing Engine (PFE) failure. In this case, the node's PFE will be disabled in the cluster by reducing the priorities of redundancy groups *x* to 0.

A failed flowd process causes failover of redundancy group *x* to the secondary node. Failover to the secondary node is automatic without the need for user intervention. When the failed (former) primary node has its failing component restored, failback will be determined by the preempt configuration for the redundancy group *x*.

Cold-Sync Monitoring

The process of synchronizing the data plane runtime objects (RTOs) on the startup of the SPUs or flowd is called *cold sync*. When all the RTOs are synchronized, the cold-sync process is complete and, the SPU or flowd on the node is ready to take over for the primary node, if needed. The process of monitoring the cold-sync state of all the SPUs or flowd on a node is called *cold-sync monitoring*. Keep in mind that when preempt is enabled, cold-sync monitoring will prevent the node from taking over the mastership until the cold-sync process is completed for the SPUs or flowd on the node. Cold-sync monitoring is enabled by default.

When the node is rebooted, or when the SPUs or flowd come back up from failure, the priority for all the redundancy groups *x* is 0. When an SPU or flowd comes up, it tries to start the cold-sync process with its mirror SPU or flowd on the other node.

If this is the only node in the cluster, the priorities for all the redundancy groups *x* stay at 0 until a new node joins the cluster. Although the priority is at 0, the device will still be able to receive and send traffic over its interfaces. A priority of 0 implies that it will not be able to fail over in case of a failure. When a new node joins the cluster, all the SPUs or flowd, as they come up, will start the cold-sync process with the mirror SPUs or flowd of the existing node.

When the SPU or flowd of a node that is already up detects the cold-sync request from the SPU or flowd of the peer node, it posts a message to the system indicating that the cold-sync process is complete. The SPUs or flowd of the newly joined node will post a similar message. However, they post this message only after all the RTOs are learned and cold-sync is complete. On receipt of completion messages from all the SPUs or flowd, the priority for redundancy groups *x* will move to the configured priority on each node if there are no other failures of monitored components, such as interfaces. This action ensures that the existing primary node for redundancy *x* groups always moves to the configured priority first. The node joining the cluster later will move to its configured priorities only after all its SPUs or flowd have completed their cold-sync process. This action in turn guarantees that the newly added node is ready with all the RTOs before it takes over mastership.

Related Topics

- Redundancy Group Interface Monitoring on page 874

Understanding Redundant Ethernet Interfaces

A redundant Ethernet interface is a pseudointerface that includes a physical interface from each node of the cluster.

Before You Begin
Read “Understanding Chassis Cluster Formation” on page 869 and “Understanding Redundancy Groups” on page 870

A redundant Ethernet interface can contain either a pair of Fast Ethernet interfaces or a pair of Gigabit Ethernet interfaces that are referred to as child interfaces of the redundant Ethernet interface (the redundant parent). Each redundant Ethernet interface can contain only two interfaces because a cluster contains only two nodes. A single redundant Ethernet interface might include a Fast Ethernet from node 0 and a Fast Ethernet from node 1 or a Gigabit Ethernet from node 0 and a Gigabit Ethernet from node 1. Although a redundant Ethernet interface's interfaces must be the same kind—either Fast Ethernet or Gigabit Ethernet—they do not need to be in the same slots on each node.

NOTE: A redundant Ethernet interface is referred to as a *reth* in configuration commands.

The maximum number of redundant Ethernet interfaces that you can configure varies, depending on the device type you are using, as shown in Table 109 on page 880. Note that the number of redundant Ethernet interfaces configured determines the number of redundancy groups that can be configured.

Table 109: Maximum Number of Redundant Ethernet Interfaces Allowed

Device	Maximum Number of Redundant Ethernet Interfaces
SRX100	8
SRX210	8
SRX240	24
SRX650	68
SRX3400	128
SRX3600	128
SRX5600	128
SRX5800	128
J2320	128
J2350	128
J4350	128
J6350	128

A redundant Ethernet interface's child interface is associated with the redundant Ethernet interface as part of the child interface configuration. The redundant Ethernet interface child interface inherits most of its configuration from its parent.

A redundant Ethernet interface inherits its failover properties from the redundancy group *x* that it belongs to. A redundant Ethernet interface remains active as long as its primary child interface is available/active. For example, if reth0 is associated with redundancy group 1 and redundancy group 1 is active on node 0, then reth0 will be up as long as the node 0 child of reth0 is up.

Related Topics

- Understanding Chassis Cluster Formation on page 869
- Understanding Redundancy Groups on page 870
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912

- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885
- Understanding Failover on page 888

Understanding the Control Plane

The control plane software, which operates in active/backup mode, is an integral part of JUNOS Software that is active on the primary node of a cluster. It achieves redundancy by communicating state, configuration, and other information to the inactive Routing Engine on the secondary node. If the master Routing Engine fails, the secondary one is ready to assume control.

Before You Begin
Read "Understanding Chassis Cluster Formation" on page 869

This topic includes:

- About the Control Link on page 882
- About Dual Control LInks on page 882
- About Heartbeats on page 883
- About Control Link Failure and Recovery on page 884
- Related Topics on page 885

The control plane software:

- Runs on the Routing Engine and oversees the entire chassis cluster system, including interfaces on both nodes
- Manages system and data plane resources, including the Packet Forwarding Engine (PFE) on each node
- Synchronizes the configuration over the control link
- Establishes and maintains sessions, including authentication, authorization, and accounting (AAA) functions
- Manages application-specific signaling protocols
- Establishes and maintains management sessions, such as Telnet connections
- Handles asymmetric routing
- Manages routing state, Address Resolution Protocol (ARP) processing, and Dynamic Host Configuration Protocol (DHCP) processing

Information from the control plane software follows two paths:

- On the primary node (where the Routing Engine is active), control information flows from the Routing Engine to the local Packet Forwarding Engine.

- Control information flows across the control link to the secondary node's Routing Engine and Packet Forwarding Engine.

The control plane software running on the master Routing Engine maintains state for the entire cluster, and only processes running on its node can update state information. The master Routing Engine synchronizes state for the secondary node and also processes all host traffic.

About the Control Link

The control link relies on a proprietary protocol to transmit session state, configuration, and liveliness signals across the nodes.

SRX Series Devices

On SRX5600 and SRX5800 devices, by default, all control ports are disabled. Each SPC in a device has two control ports, and each device can have multiple SPCs plugged into it. To set up the control link in a chassis cluster with SRX5600 or SRX5800 devices, you connect and configure the control ports that you will use on each device (fpc*n* and fpc*n*) and then initialize the device in cluster mode.

For SRX3400 and SRX3600 devices, there are dedicated chassis cluster (HA) control ports on the switch fabric board. No control link configuration is needed for SRX3400 and SRX3600 devices.

For SRX650 and SRX240 devices, the ge-0/0/1 interface is used for the control link.

For SRX210 and SRX100 devices, the fe-0/0/7 interface is used for the control link.

J Series Devices

In a J Series chassis cluster, the control link is a physical connection between the ge-0/0/3 ports on each device, with both transformed into fxp1.

To set up the control link on J Series devices, you connect the control interfaces on the two devices back-to-back. When you initialize a device in cluster mode, JUNOS Software renames the control interface to fxp1 and uses that interface for the cluster control link. To enable the control link to transmit data, the system provides each fxp1 control link interface with an internal IP address.

About Dual Control LInks

Dual control links, where two pairs of control link interfaces are connected between each device in a cluster, are supported for the SRX5000 and SRX3000 lines. To use this functionality, each device in the cluster requires a second Routing Engine. For the SRX5000 line, this means installing a second Routing Engine, as well as a second Switch Control Board (SCB) to house the Routing Engine, and for the SRX3000 line, installing a second Routing Engine Control Board (RECB).

NOTE: The second Routing Engine must be running JUNOS Release 10.0 or later.

The second Routing Engine does not provide backup functionality. Its purpose is only to initialize the switch on the SCB (SRX5000 line) or on the RECB (SRX3000 line). The second Routing Engine does not need to be upgraded, even when there is a software upgrade of the master Routing Engine on the same node. Note the following:

- You cannot run the CLI or enter configuration mode on the second Routing Engine.
- You do not need to set the chassis ID and cluster ID on the second Routing Engine.
- You need only a console connection to the second Routing Engine. (A console connection is not needed unless you want to check that the second Routing Engine booted up or to upgrade a software image.)
- You cannot log in to the second Routing Engine from the master Routing Engine.

NOTE: As long as the first Routing Engine is installed (even if it is rebooting or failing), the second Routing Engine cannot take over the chassis mastership; that is, it cannot control all the hardware on the chassis. The second Routing Engine can only become the master when the master Routing Engine is not present.

Having two control links helps to avoid a single point of failure by reducing the number of disabled cases that are caused by control link failure. (For a chassis cluster with one control link, if the control link goes down, all redundancy groups on the secondary node go to the ineligible state and eventually to the disabled state.)

About Heartbeats

JUNOS Software transmits heartbeat signals over the control link at a configured interval. The system uses heartbeat transmissions to determine the “health” of the control link. If the number of missed heartbeats has reached the configured threshold, the system assesses whether a failure condition exists.

You specify the heartbeat threshold and heartbeat interval when you configure the chassis cluster.

The system monitors the control link's status by default.

For dual control links, which are supported on the SRX5000 and SRX3000 lines, the Juniper Services Redundancy Protocol process (jsrpd) sends and receives the control heartbeat messages on both control links. As long as heartbeats are received on one of the control links, JUNOS Software considers the other node to be alive.

About Control Link Failure and Recovery

If the control link fails, JUNOS Software disables the secondary node to prevent the possibility of each node becoming primary for all redundancy groups, including redundancy group 0.

In the event of a legitimate control link failure, redundancy group 0 remains primary on the node on which it is currently primary, inactive redundancy groups *x* on the primary node become active, and the secondary node enters a disabled state in which it is not handling traffic.

NOTE: When the secondary node is disabled, you can still log in to the management port and run diagnostics.

To determine if a legitimate control link failure has occurred, the system relies on redundant liveliness signals sent across the control link and the data link.

The system periodically transmits probes over the fabric data link and heartbeat signals over the control link. Probes and heartbeat signals share a common sequence number that maps them to a unique time event. The software identifies a legitimate control link failure if the following two conditions exist:

- The threshold number of heartbeats were lost.
- At least one probe with a sequence number corresponding to that of a missing heartbeat signal was received on the data link.

When a legitimate control link failure occurs, the following conditions apply:

- Redundancy group 0 remains primary on the node on which it is presently primary (and thus its Routing Engine remains active), and all redundancy groups *x* on the node become primary.

  If the system cannot determine which Routing Engine is primary, the node with the higher priority value for redundancy group 0 is primary and its Routing Engine is active. (You configure the priority for each node when you configure the redundancy-group statement for redundancy group 0.)

- The system disables the secondary node.

  To recover a device from disabled mode, you must reboot the device. When you reboot the disabled node, the node will synchronize its dynamic state with the primary node.

NOTE: If you make any changes to the configuration while the secondary node is disabled, execute the commit command to synchronize the configuration after you reboot the node. If you did not make configuration changes, the configuration file remains synchronized with that of the primary node.

You cannot enable preemption for redundancy group 0. If you want to change the primary node for redundancy group 0, you must do a manual failover.

When you use dual control links (supported on the SRX5000 and SRX3000 lines), note the following:

- Host inbound or outbound traffic can be impacted for up to 3 seconds during a control link failure. For example, consider a case where redundancy group 0 is primary on node 0 and there is a Telnet session to the Routing Engine through a network interface port on node 1. If the currently active control link fails, the Telnet session will lose packets for 3 seconds, until this failure is detected.
- A control link failure that occurs while the commit process is running across two nodes might lead to commit failure. In this situation, run the commit command again after 3 seconds.

NOTE: Dual control links require a second Routing Engine on each node of the chassis cluster. For more information, see “About Dual Control LInks” on page 882.

You can specify that control link recovery be done automatically by the system by setting the control-link-recovery statement. In this case, once the system determines that the control link is healthy, it issues an automatic reboot on the disabled node. When the disabled node reboots, the node joins the cluster again.

Related Topics

- Understanding the Control Plane on page 881
- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879
- Understanding Failover on page 888

Understanding the Data Plane

The data plane software, which operates in active/active mode, manages flow processing and session state redundancy and processes transit traffic. All packets belonging to a particular session are processed on the same node to ensure that the same security treatment is applied to them. The system identifies the node on which a session is active and forwards its packets to that node for processing. (After a packet is processed, the Packet Forwarding Engine transmits the packet to the node on which its egress interface exists if that node is not the local one.)

To provide for session (or flow) redundancy, the data plane software synchronizes its state by sending special payload packets called runtime objects (RTOs) from one node to the other across the fabric data link. By transmitting information about a session between the nodes, RTOs ensure the consistency and stability of sessions if a failover were to occur, and thus they enable the system to continue to process traffic belonging to existing sessions. To ensure that session information is always

synchronized between the two nodes, the data plane software gives RTOs transmission priority over transit traffic.

Before You Begin
Read “Understanding Chassis Cluster Formation” on page 869

This topic includes:

- About Session RTOs on page 886
- About the Fabric Data Link on page 886
- About Data Forwarding on page 887
- About Fabric Data Link Failure and Recovery on page 887
- Related Topics on page 888

About Session RTOs

The data plane software creates RTOs for UDP and TCP sessions and tracks state changes. It also synchronizes traffic for IPv4 passthrough protocols such as Generic Routing Encapsulation (GRE) and IPsec.

RTOs for synchronizing a session include:

- Session creation RTOs on the first packet
- Session deletion and ageout RTOs
- Change-related RTOs, which include:
  - TCP state changes
  - Timeout synchronization request and response messages
  - RTOs for creating and deleting temporary openings in the firewall (pinholes) and child session pinholes

About the Fabric Data Link

The data link is referred to as the fabric interface. It is used by the cluster's Packet Forwarding Engines to transmit transit traffic and to synchronize data plane dynamic runtime state. When the system creates the fabric interface, the software assigns it an internally derived IP address to be used for packet transmission.

The fabric is a physical connection between two nodes of a cluster and is formed by connecting a pair of Ethernet interfaces back-to-back (one from each node).

Unlike for the control link, whose interfaces are determined by the system, you specify the physical interfaces to be used for the fabric data link in the configuration.

For SRX Series chassis clusters, the fabric link can be any pair of Ethernet interfaces spanning the cluster; for J Series chassis clusters, any pair of Gigabit Ethernet interfaces.

Table 110: Supported Fabric Interface Types for SRX Series Devices

SRX5000 line	SRX3000 line	SRX650	SRX240	SRX210	SRX100
Fast Ethernet	Fast Ethernet	Fast Ethernet	Fast Ethernet	Fast Ethernet	Fast Ethernet
Gigabit Ethernet	Gigabit Ethernet	Gigabit Ethernet	Gigabit Ethernet	Gigabit Ethernet	
10-Gigabit Ethernet	10-Gigabit Ethernet				

The fabric data link does not support fragmentation. To accommodate this state, jumbo frame support is enabled by default on the link with an MTU size of 8980 bytes. To ensure that traffic that transits the data link does not exceed this size, we recommend that no other interface exceed the fabric data link's MTU size.

About Data Forwarding

For JUNOS Software, flow processing occurs on a single node on which the session for that flow was established and is active. This approach ensures that the same security measures are applied to all packets belonging to a session.

A chassis cluster can receive traffic on an interface on one node and send it out an interface on the other node. (In active/active mode, the ingress interface for traffic might exist on one node and its egress interface on the other.)

This traversal is required in the following situations:

- When packets are processed on one node, but need to be forwarded out an egress interface on the other node
- When packets arrive on an interface on one node, but must be processed on the other node

  If the ingress and egress interfaces for a packet are on one node, but the packet must be processed on the other node because its session was established there, it must traverse the data link twice. This can be the case for some complex media sessions, such as voice-over-IP (VoIP) sessions.

About Fabric Data Link Failure and Recovery

NOTE: Intrusion Detection and Prevention (IDP) services do not support failover. For this reason, IDP services will not be applied for sessions that were present prior to the failover. IDP services will be applied for new sessions created on the new primary node.

The fabric data link is vital to the chassis cluster. If the link is unavailable, traffic forwarding and RTO synchronization are affected, which can result in loss of traffic and unpredictable system behavior.

To eliminate this possibility, JUNOS Software detects fabric faults and disables one node of the cluster. It determines that a fabric fault has occurred if a fabric probe is not received but the fabric interface is active.

To recover from this state, you must reboot the disabled node. When you reboot it, the node synchronizes its state and RTOs with the primary node.

NOTE: If you make any changes to the configuration while the secondary node is disabled, execute the commit command to synchronize the configuration after you reboot the node. If you did not make configuration changes, the configuration file remains synchronized with that of the primary node.

Related Topics

- Understanding the Control Plane on page 881
- Understanding Chassis Cluster Formation on page 869
- Understanding Redundancy Groups on page 870

Understanding Failover

Chassis cluster employs a number of highly efficient failover mechanisms that promote high availability to increase your system's overall reliability and productivity.

Before You Begin

For background information, read:

- Understanding Chassis Cluster Formation on page 869
- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885

This topic includes:

- About Redundancy Group Failover on page 888
- About Manual Failover on page 889

About Redundancy Group Failover

A redundancy group is a collection of objects that fail over as a group. Each redundancy group monitors a set of objects (physical interfaces), and each monitored

object is assigned a weight. Each redundancy group has an initial threshold of 255. When a monitored object fails, the weight of the object is subtracted from the threshold value of the redundancy group. When the threshold value reaches zero, the redundancy group fails over to the other node. As a result, all the objects associated with the redundancy group will fail over as well. Graceful restart of the routing protocols enables the SRX Series device to minimize traffic disruption during a failover. For more information, see "Redundancy Group Interface Monitoring" on page 874.

Because back-to-back redundancy group 0 failovers that occur too quickly can cause a cluster to exhibit unpredictable behavior, a dampening time between failovers is needed. On a failover, the previous primary node moves to the secondary-hold state and stays there until the hold-down interval expires, after which it moves to the secondary state.

The default dampening time is 300 seconds (5 minutes) for redundancy group 0 and is configurable to up to 1800 seconds with the hold-down-interval statement. Redundancy groups *x* (redundancy groups numbered 1 through 128) have a default dampening time of 1 second, with a range of 0 through 1800 seconds. The hold-down interval affects manual failovers, as well as automatic failovers associated with monitoring failures.

About Manual Failover

You can initiate a redundancy group *x* failover manually. A manual failover applies until a failback event occurs.

For example, suppose that the user manually does a redundancy group 1 failover from node 0 to node 1. Then an interface that redundancy group 1 is monitoring fails, dropping the threshold value of the new primary redundancy group to zero. This event is considered a failback event, and the system returns control to the original redundancy group.

You can also initiate a redundancy group 0 failover manually if you want to change the primary node for redundancy group 0. You cannot enable preemption for redundancy group 0.

When you do a manual failover for redundancy group 0, the node in the primary state transitions to the secondary-hold state. The node will stay in the secondary-hold state for the default or configured time (a minimum of 300 seconds) and then transition to the secondary state.

State transitions in cases where one node is in the secondary-hold state and the other node reboots, or the control link connection or fabric link connection is lost to that node, are described as follows:

- Reboot case—The node in the secondary-hold state transitions to the primary state; the other node goes dead (inactive).
- Control link failure case—The node in the secondary-hold state transitions to the ineligible state and then to a disabled state; the other node transitions to the primary state.
- Fabric link failure case—The node in the secondary-hold state transitions directly to the disabled state.

Keep in mind that during an in-service software upgrade (ISSU), the transitions described above will not happen. Instead, the other (primary) node will transition directly to the secondary state because Juniper releases earlier than 10.0 do not interpret the secondary-hold state. While you start an ISSU, if one of the nodes has one or more redundancy groups in the secondary-hold state, you must wait for them to move to the secondary state before you can do manual failovers to make all the redundancy groups be primary on one node. For more information about ISSUs, see "Low-Impact ISSU Chassis Cluster Upgrades" on page 939.

CAUTION: Be cautious and judicious in your use of redundancy group 0 manual failovers. A redundancy group 0 failover implies a Routing Engine failover, in which case all processes running on the primary node are killed and then spawned on the new master Routing Engine. This failover could result in loss of state, such as routing state, and degrade performance by introducing system churn.

Hardware Setup for SRX Series Chassis Clusters

An SRX Series chassis cluster is created by physically connecting two identical cluster-supported SRX Series devices together using a pair of the same type of Ethernet connections. The connection is made for both a control link and a fabric (data) link between the two devices.

NOTE: For the SRX5000 and SRX3000 lines, you can connect two control links between the two devices in the cluster to reduce the chance of control link failure. Dual control links require a second Routing Engine on each node of the cluster. For more information, see "About Dual Control LInks" on page 882.

This topic includes:

- Connecting Dual Control Links on page 893

Control links in a chassis cluster are made using specific ports.

You must use the following ports to form the control link on the SRX Series branch devices:

- For SRX100 devices, connect the fe-0/0/7 port to the fe-1/0/7 port
- For SRX210 devices, connect the fe-0/0/7 port to the fe-2/0/7 port
- For SRX240 devices, connect the ge-0/0/1 port to the ge-5/0/1 port
- For SRX650 devices, connect the ge-0/0/1 port to the ge-9/0/1 port

For a device from the SRX3000 line, the connection that serves as the control link must be between the built-in control ports on each device.

SRX5000 line devices do not have built-in ports, so the control link for these gateways must be the control ports on their Services Processing Cards (SPCs) with a slot numbering offset of 6 for SRX5600 devices and 12 for SRX5800 devices.

NOTE: For the SRX3000 and SRX5000 lines, the control link ports are a one-to-one mapping with the Routing Engine slot. If your Routing Engine is in slot 0, you must use control port 0 to link the Routing Engines.

The fabric link connection for the SRX100 must be a pair of Fast Ethernet interfaces and for the SRX210 must be a pair of either Fast Ethernet or Gigabit Ethernet interfaces. The fabric link connection must be any pair of either Gigabit Ethernet or 10-Gigabit Ethernet interfaces on all other SRX Series devices.

Figure 156 on page 891, Figure 157 on page 891, Figure 158 on page 892, Figure 159 on page 892, Figure 160 on page 892, and Figure 161 on page 893 show pairs of SRX Series devices with the fabric links and control links connected.

Figure 156: Connecting SRX Series Devices in a Cluster (SRX5800 Devices)

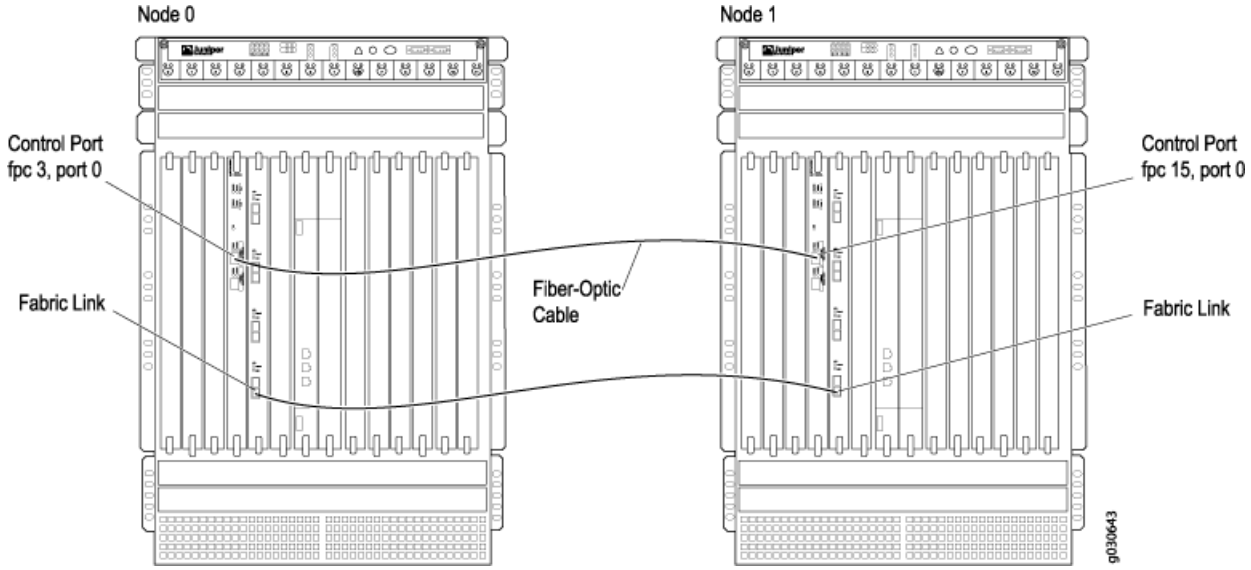

Figure 157: Connecting SRX Series Devices in a Cluster (SRX3400 Devices)

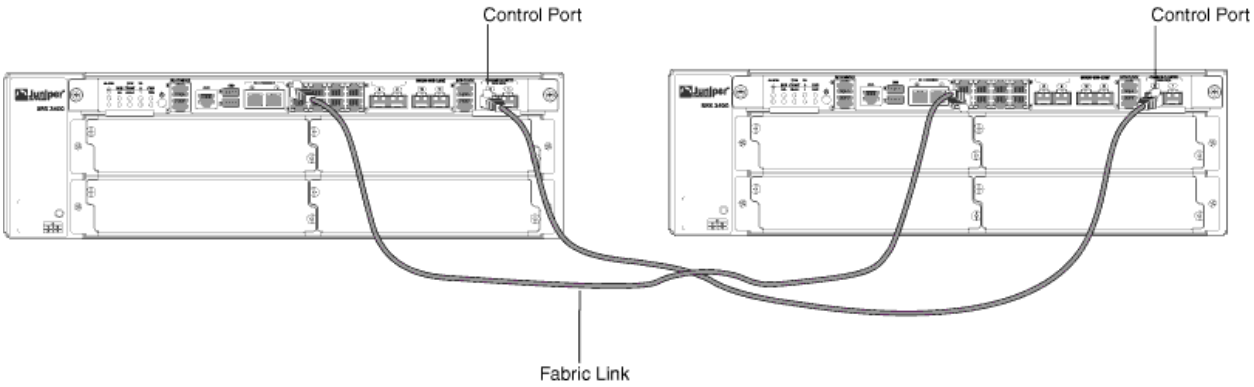

Figure 158: Connecting SRX Series Devices in a Cluster (SRX650 Devices)

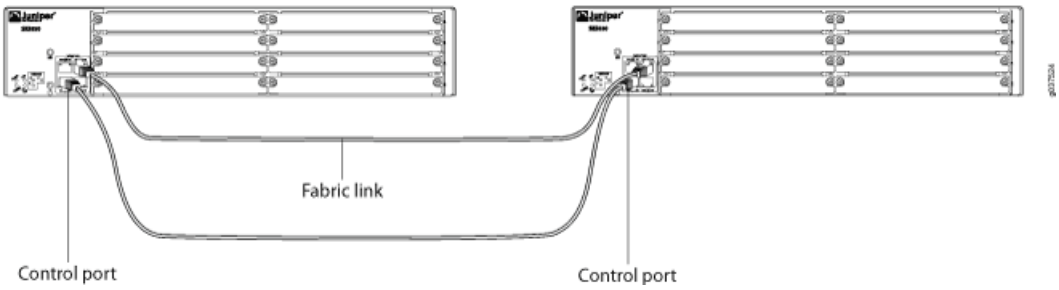

Figure 159: Connecting SRX Series Devices in a Cluster (SRX240 Devices)

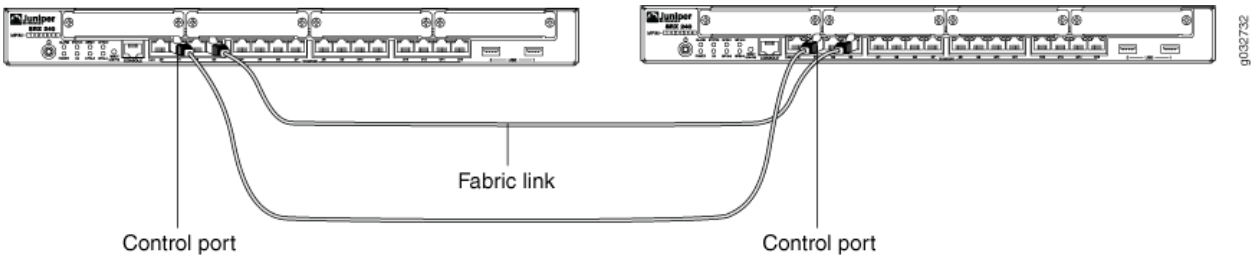

Figure 160: Connecting SRX Series Devices in a Cluster (SRX210 Devices)

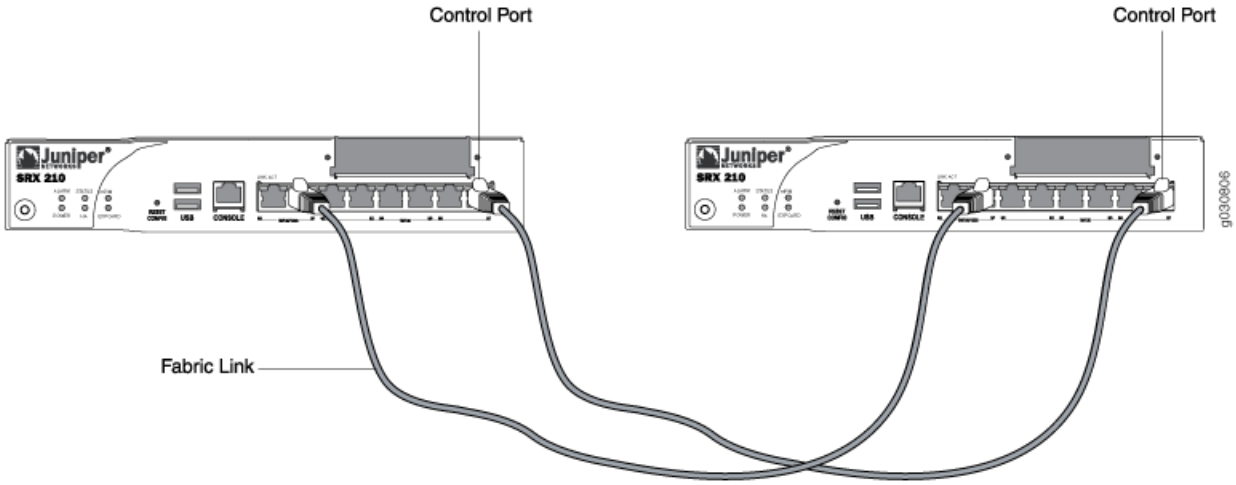

Figure 161: Connecting SRX Series Devices in a Cluster (SRX100 Devices)

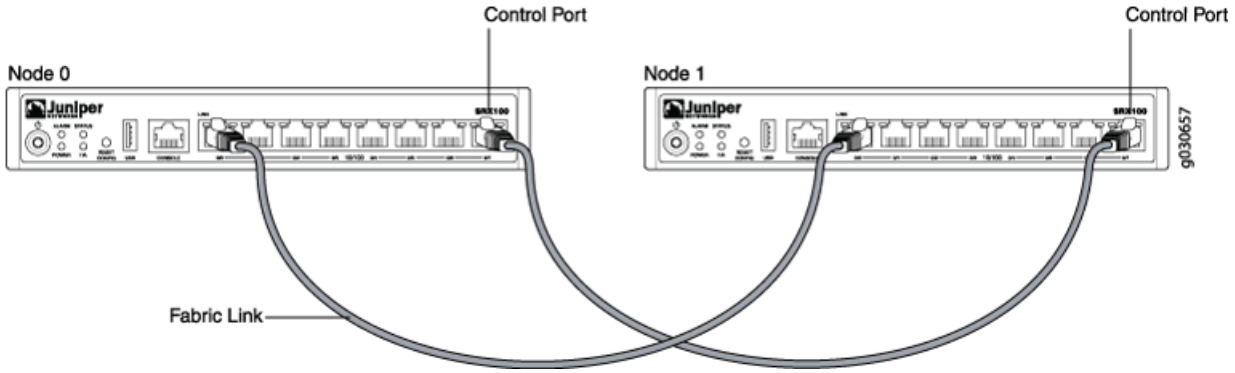

Connecting Dual Control Links

For the SRX5000 and SRX3000 lines, you can connect two control links between the two devices, effectively reducing the chance of control link failure.

For devices in the SRX5000 line, connect two pairs of the same type of Ethernet ports. For each device, you can use ports on the same Services Processing Card (SPC), but we recommend that they be on two different SPCs to provide high availability. Figure 162 on page 893 shows a pair of SRX5800 devices with dual control links connected. In this example, control port 0 and control port 1 are connected on different SPCs.

Figure 162: Connecting Dual Control Links (SRX5800 Devices)

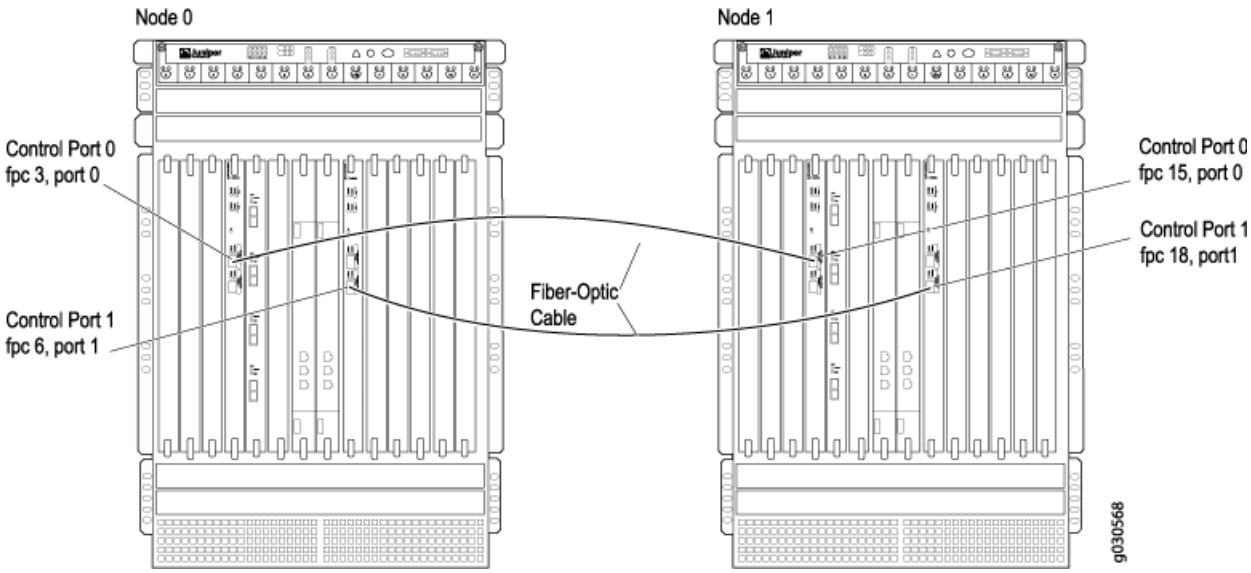

For devices in the SRX3000 line, connect two pairs of the same type of Ethernet ports. For each device, use both available built-in ports. Figure 163 on page 894 shows a pair of SRX3400 devices with dual control links connected.

Figure 163: Connecting Dual Control Ports (SRX3400 Devices)

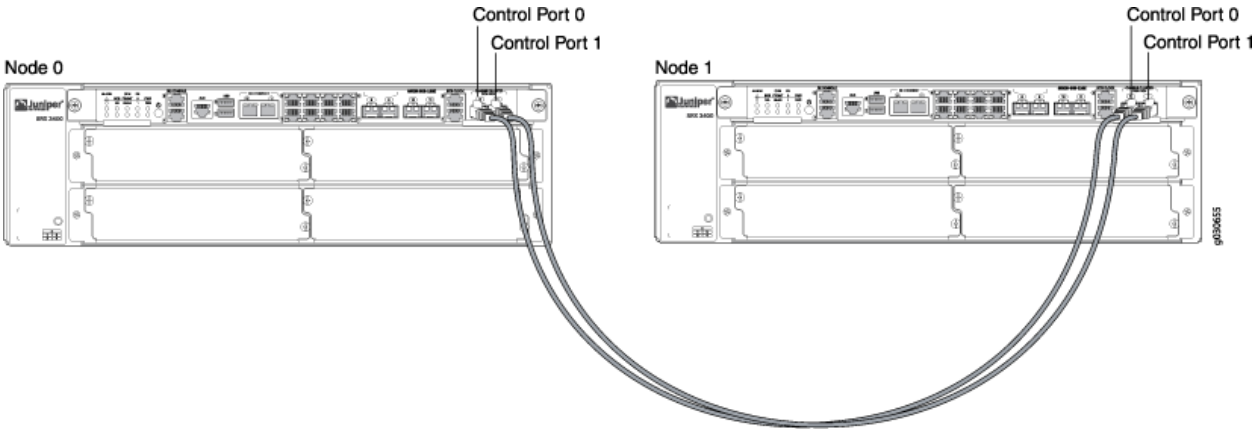

NOTE: For devices in both the SRX5000 and SRX3000 lines, you must connect control port 0 to control port 0 on one node and control port 1 to control port 1 on the other node. If you connect control port 0 to control port 1 and vice versa, the nodes will not be able to receive control heartbeat packets over the control link.

Hardware Setup for J Series Chassis Clusters

To create a J Series chassis cluster, you must physically connect a pair of the same kind of supported J Series devices back-to-back over a pair of Gigabit Ethernet connections. The connection that serves as the control link must be the built-in interface ge-0/0/3. The fabric link connection can be a combination of any pair of Gigabit Ethernet interfaces on the devices. Figure 164 on page 894 shows two J Series devices connected using the built-in interfaces for both the fabric and control links.

Figure 164: Connecting J Series Devices in a Cluster (J6350 Devices)

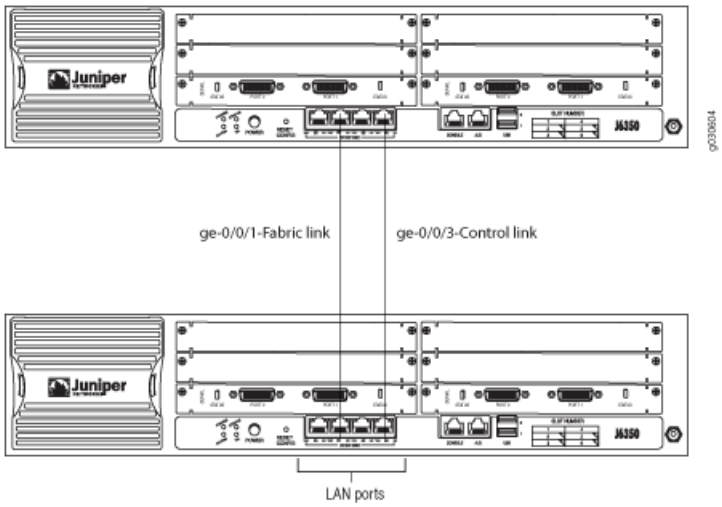

NOTE: When chassis clustering is enabled on a J Series router, two interface ports are used to link the two devices: the ge-0/0/3 interface (fxp1 port) is used for the control interface and one port is used for the fabric link (using either one of the built-in interfaces (ge-0/0/0 or ge-0/0/1) or one of the ports of a uPIM). Also, the ge-0/0/2 interface (fxp0 port) is used for the management link. This means that three of the four onboard Gigabit Ethernet ports will be in use; if additional ports are required for transit traffic, then a PIM or uPIM will be required.

What Happens When You Enable Chassis Cluster

After wiring the two devices together as described in "Hardware Setup for SRX Series Chassis Clusters" on page 890 or "Hardware Setup for J Series Chassis Clusters" on page 894, you use CLI *operational mode* commands to enable chassis clustering by assigning a cluster ID and node ID on each chassis in the cluster. The cluster ID is the same on both nodes.

To do this, you connect to the console port on the device that will be the primary, give it a node ID, and identify the cluster it will belong to, then reboot the system. You then connect the console port to the other device, give it a node ID, and assign it the same cluster ID you gave to the first node, then reboot the system. In both instances, you can cause the system to boot automatically by including the reboot parameter in the CLI command line. (For further explanation of primary and secondary nodes, see "Understanding Redundancy Groups" on page 870.)

WARNING: After fabric interfaces have been configured on a chassis cluster, removing the fabric configuration on either node will cause the redundancy group 0 (RG0) secondary node to move to a disabled state. (Resetting a device to the factory default configuration removes the fabric configuration and thereby causes the RG0 secondary node to move to a disabled state.) After the fabric configuration is committed, do not reset either device to the factory default configuration.

Figure 165 on page 899 shows how the FPC slots are numbered on two nodes in an SRX5000 line chassis cluster. Other figures show slot numbering on both nodes in other SRX Series chassis clusters. Figure 171 on page 903 shows how the PIM slots are numbered on two nodes in a J Series chassis cluster.

This topic includes:

- Node Interfaces on SRX Series Chassis Clusters on page 896
- Node Interfaces on J Series Chassis Clusters on page 901
- Management Interfaces on page 903
- Fabric Interface on page 903
- Control Interfaces on page 903
- Related Topics on page 904

Node Interfaces on SRX Series Chassis Clusters

Normally, on SRX Series devices, the built-in interfaces are numbered as follows:

For Most SRX Series Devices	ge-0/0/0	ge-0/0/1	ge-0/0/2	ge-0/0/3	. . .
For SRX210 Devices	ge-0/0/0	ge-0/0/1	fe-0/0/2	fe-0/0/3	. . .
For SRX100 Devices	fe-0/0/0	fe-0/0/1	fe-0/0/2	fe-0/0/3	. . .

For chassis clustering, all SRX Series devices have a built-in management interface named fxp0. For most SRX Series devices, the fxp0 interface is a dedicated port. For SRX210 and SRX100 devices, after you enable chassis clustering and reboot the system, the built-in interface named fe-0/0/6 is repurposed as the management interface and is automatically renamed fxp0.

For the SRX5000 line, control interfaces are configured on SPCs. For the SRX3000 line and the SRX650 and SRX240 devices, control interfaces are dedicated Gigabit Ethernet ports. For SRX210 and SRX100 devices, after you enable chassis clustering and reboot the system, the built-in interface named fe-0/0/7 is repurposed as the control interface and is automatically renamed fxp1.

After the devices are connected as a cluster, the slot numbering on one device and thus the interface numbering will change. The slot number for each slot in both nodes is determined using the following formula:

cluster slot number = (node ID * maximum slots per node) + local slot number

In chassis cluster mode, the interfaces on the secondary node are renumbered internally. For example, the management interface port on the front panel of each SRX210 device is still labeled fe-0/0/6, but internally, the node 1 port is referred to as fe-2/0/6.

Table 111 on page 897 shows the slot numbering, as well as the port and interface numbering, for both of the SRX Series devices that become node 0 and node 1 of the chassis cluster after the cluster is formed.

Table 111: SRX Series Chassis Cluster Slot Numbering and Port/Interface Naming

Model	Chassis	Maximum Slots Per Node	Slot Numbering in a Cluster	Management Physical Port/Logical Interface	Control Physical Port/Logical Interface	Fabric Physical Port/Logical Interface
5800	Node 0	12 (FPC slots)	0 — 11	Dedicated Gigabit Ethernet port	Control port on an SPC	Any Ethernet port
				fxp0	em0	fab0
	Node 1		12 — 23	Dedicated Gigabit Ethernet port	Control port on an SPC	Any Ethernet port
				fxp0	em0	fab1
5600	Node 0	6 (FPC slots)	0 — 5	Dedicated Gigabit Ethernet port	Control port on an SPC	Any Ethernet port
				fxp0	em0	fab0
	Node 1		6 — 11	Dedicated Gigabit Ethernet port	Control port on an SPC	Any Ethernet port
				fxp0	em0	fab1
3600	Node 0	13 (CFM slots)	0 — 12	Dedicated Gigabit Ethernet port	Dedicated Gigabit Ethernet port	Any Ethernet port
				fxp0	em0	fab0
	Node 1		13 — 25	Dedicated Gigabit Ethernet port	Dedicated Gigabit Ethernet port	Any Ethernet port
				fxp0	em0	fab1
3400	Node 0	8 (CFM slots)	0 — 7	Dedicated Gigabit Ethernet port	Dedicated Gigabit Ethernet port	Any Ethernet port
				fxp0	em0	fab0
	Node 1		8 — 15	Dedicated Gigabit Ethernet port	Dedicated Gigabit Ethernet port	Any Ethernet port
				fxp0	em0	fab1

Table 111: SRX Series Chassis Cluster Slot Numbering and Port/Interface Naming *(continued)*

Model	Chassis	Maximum Slots Per Node	Slot Numbering in a Cluster	Management Physical Port/Logical Interface	Control Physical Port/Logical Interface	Fabric Physical Port/Logical Interface
650	Node 0	9 (PIM slots)	0 — 8	ge-0/0/0	ge-0/0/1	Any Ethernet port
				fxp0	fxp1	fab0
	Node 1		9 — 17	ge-9/0/0	ge-9/0/1	Any Ethernet port
				fxp0	fxp1	fab1
240	Node 0	5 (PIM slots)	0 — 4	ge-0/0/0	ge-0/0/1	Any Ethernet port
				fxp0	fxp1	fab0
	Node 1		5 — 9	ge-5/0/0	ge-5/0/1	Any Ethernet port
				fxp0	fxp1	fab1
210	Node 0	2 (PIM slots)	0 and 1	fe-0/0/6	fe-0/0/7	Any Ethernet port
				fxp0	fxp1	fab0
	Node 1		2 and 3	fe-2/0/6	fe-2/0/7	Any Ethernet port
				fxp0	fxp1	fab1
100	Node 0	1(PIM slot)	0	fe-0/0/6	fe-0/0/7	Any Ethernet port
				fxp0	fxp1	fab0
	Node 1		1	fe-1/0/6	fe-1/0/7	Any Ethernet port
				fxp0	fxp1	fab1

Information about cluster slot numbering is also provided in Figure 165 on page 899, Figure 166 on page 900, Figure 167 on page 900, Figure 168 on page 900, Figure 169 on page 901, and Figure 170 on page 901

NOTE: See the appropriate *Services Gateway Hardware Guide* for details about SRX Series devices. The *JUNOS Software Interfaces and Routing Configuration Guide* provides a full discussion of the interface naming convention.

After you enable chassis clustering, the two chassis joined together cease to exist as individuals and now represent a single system. As a single system, the cluster now has twice as many slots. (See Figure 165 on page 899, Figure 166 on page 900, Figure 167 on page 900, Figure 168 on page 900, Figure 169 on page 901, and Figure 170 on page 901.)

Figure 165: FPC Slot Numbering in an SRX Series Chassis Cluster (SRX5800 Devices)

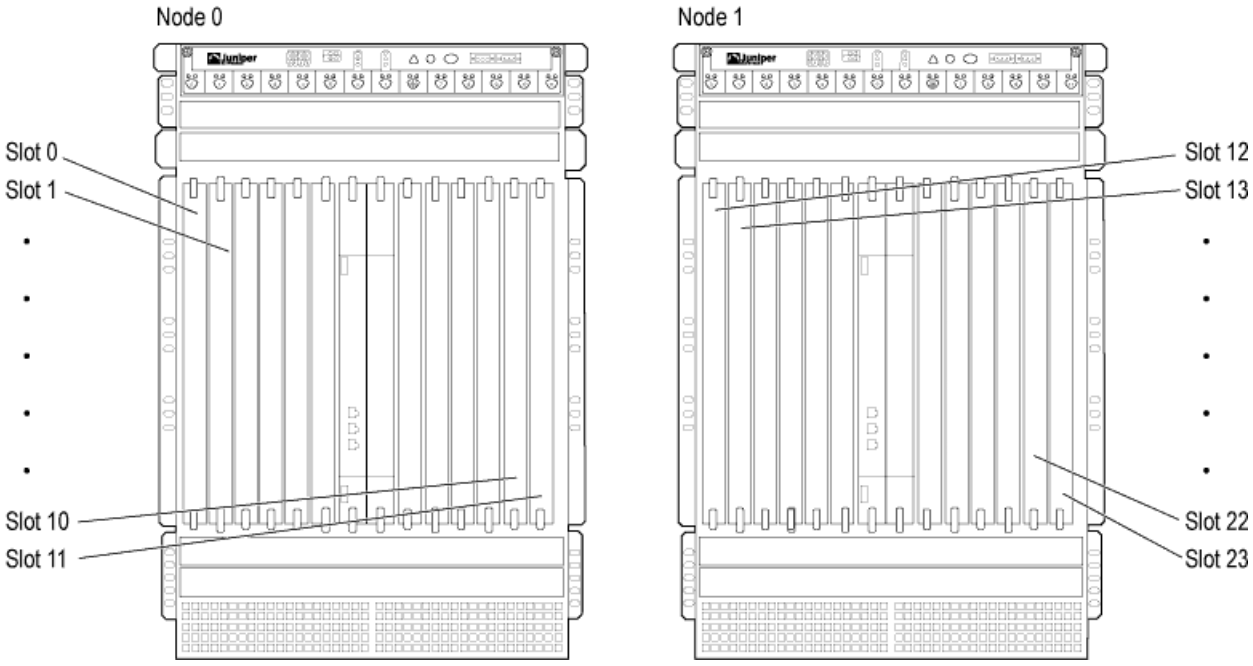

NOTE: SRX5600 and SRX5800 devices have Flex I/O Cards (flex IOCs) that have two slots to accept the following port modules:

- SRX-IOC-4XGE-XFP 4-Port XFP
- SRX-IOC-16GE-TX 16-Port RJ-45
- SRX-IOC-16GE-SFP 16-Port SFP

You can use these port modules to add from 4 to 16 Ethernet ports to your SRX Series device. Port numbering for these modules is

slot/port module/port

where *slot* is the number of the slot in the device in which the flex IOC is installed; *port module* is 0 for the upper slot in the flex IOC or 1 for the lower slot when the card is vertical, as in an SRX5800 device; and *port* is the number of the port on the port module. When the card is horizontal, as in an SRX5600 device, *port module* is 0 for the left-hand slot or 1 for the right-hand slot.

For more information, see the *Services Gateway Hardware Guide* for your specific SRX Series model.

Figure 166: Slot Numbering in an SRX Series Chassis Cluster (SRX3400 Devices)

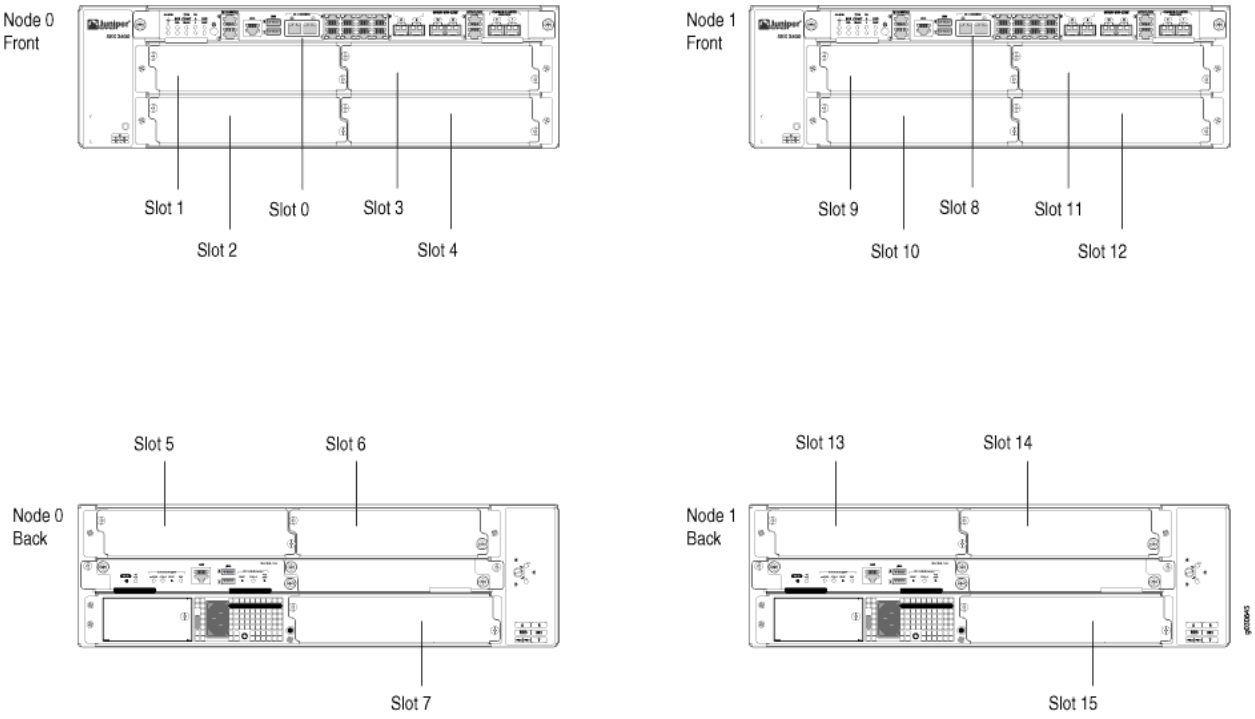

Figure 167: Slot Numbering in an SRX Series Chassis Cluster (SRX650 Devices)

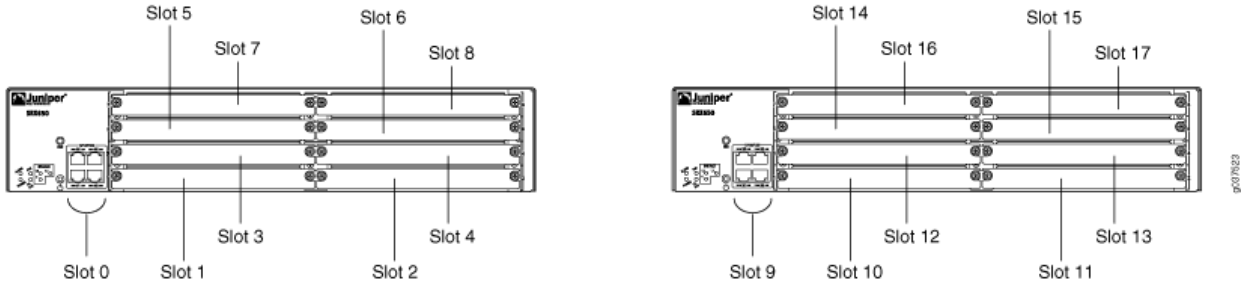

Figure 168: Slot Numbering in an SRX Series Chassis Cluster (SRX240 Devices)

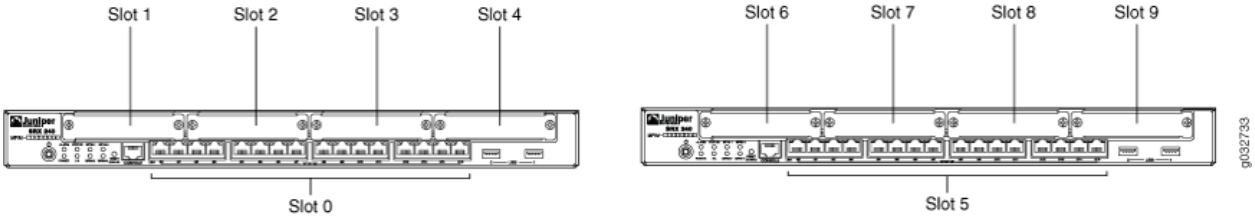

Figure 169: Slot Numbering in an SRX Series Chassis Cluster (SRX210 Devices)

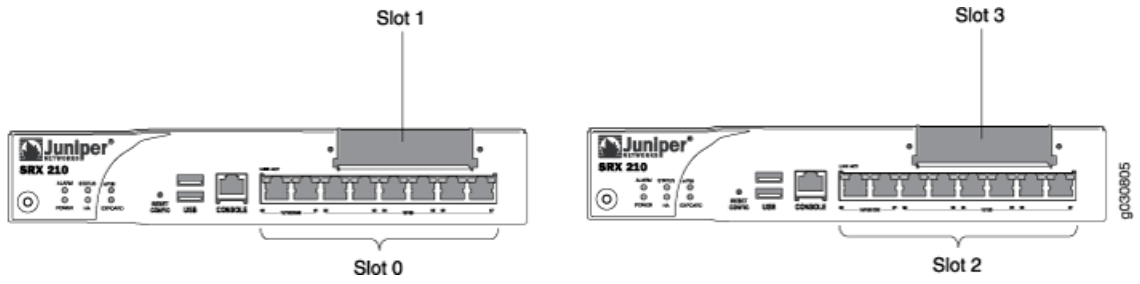

Figure 170: Slot Numbering in an SRX Series Chassis Cluster (SRX100 Devices)

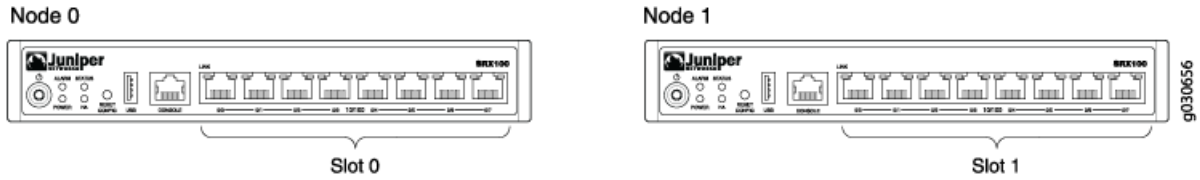

Node Interfaces on J Series Chassis Clusters

Normally, on J Series devices, the built-in interfaces are numbered as follows:

ge-0/0/0 ge-0/0/1 ge-0/0/2 ge-0/0/3 . . .

After you enable chassis clustering and reboot the system, the built-in interface named ge-0/0/2 is repurposed as the management interface and is automatically renamed fxp0. Likewise, the built-in interface named ge-0/0/3 is repurposed as the control interface and is automatically renamed fxp1.

After the devices are connected as a cluster, the slot numbering and thus the interface numbering will change for one device. The cluster determines the slot number for each slot in both nodes using the following formula:

cluster slot number = (node ID * maximum slots per node) + local slot number

In chassis cluster mode, the interfaces on the secondary node are renumbered internally. For example, the management interface port on the front panel of each J2320 device is still labeled ge-0/0/2, but internally, the node 1 port is referred to as ge-4/0/2.

Table 112 on page 902 shows the slot numbering, as well as the port and interface numbering, for both of the J Series devices that become node 0 and node 1 of the cluster after the cluster is formed.

Table 112: J Series Chassis Cluster Slot Numbering and Port/Interface Naming

Model	Chassis	Maximum Slots Per Node	Slot Numbering in a Cluster	Management Physical Port/Logical Interface	Control Physical Port/Logical Interface	Fabric Physical Port/Logical Interface
J2320	Node 0	4 (PIM slots); *includes one pre-set slot*	0 — 3	ge-0/0/2	ge-0/0/3	Any Gigabit Ethernet port
				fxp0	fxp1	fab0
	Node 1		4 — 7	ge-4/0/2	ge-4/0/3	Any Gigabit Ethernet port
				fxp0	fxp1	fab1
J2350	Node 0	6 (PIM slots); *includes one pre-set slot*	0 — 5	ge-0/0/2	ge-0/0/3	Any Gigabit Ethernet port
				fxp0	fxp1	fab0
	Node 1		6 — 11	ge-6/0/2	ge-6/0/3	Any Gigabit Ethernet port
				fxp0	fxp1	fab1
J4350 and J6350	Node 0	7 (PIM slots); *includes one pre-set slot*	0 — 6	ge-0/0/2	ge-0/0/3	Any Gigabit Ethernet port
				fxp0	fxp1	fab0
	Node 1		7 — 13	ge-7/0/2	ge-7/0/3	Any Gigabit Ethernet port
				fxp0	fxp1	fab1

Information about cluster slot numbering is also provided in Figure 171 on page 903.

NOTE: See the *J Series Services Routers Hardware Guide* for details about J Series devices. The *JUNOS Software Interfaces and Routing Configuration Guide* provides a full discussion of the interface naming convention.

After you enable chassis clustering, the two chassis joined together cease to exist as individuals and now represent a single system. As a single system, the cluster now has twice as many PIM slots. (See Figure 171 on page 903.)

Figure 171: PIM Slot Numbering in a J Series Chassis Cluster (J6350 Devices)

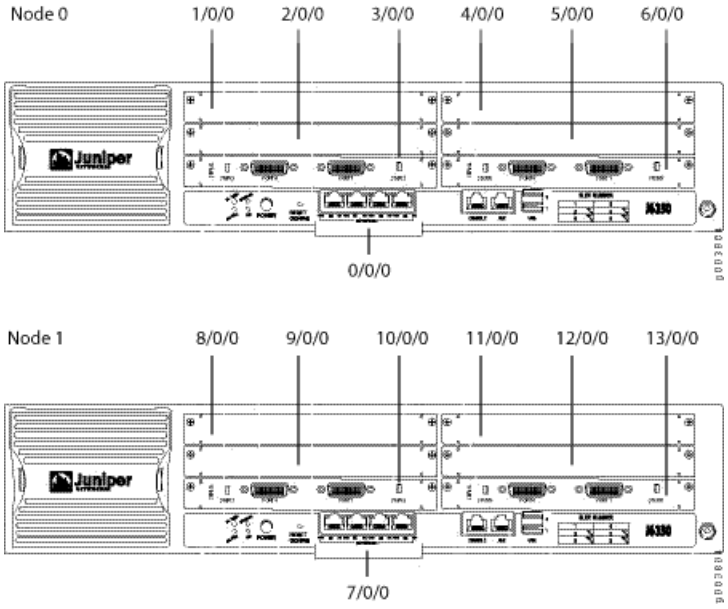

Management Interfaces

The fxp0 interfaces function like standard management interfaces on SRX Series and J Series devices and allow network access to each node in the cluster. You must, however, first connect to each node through the console port and assign a unique IP address to each fxp0 interface.

Fabric Interface

The fabric is the data link between the nodes and is used to forward traffic between the chassis. Traffic arriving on a node that needs to be processed on the other is forwarded over the fabric data link. Similarly, traffic processed on a node that needs to exit through an interface on the other node is forwarded over the fabric. The fabric also provides for synchronization of session state objects created by operations such as authentication, Network Address Translation (NAT), Application Layer Gateways (ALGs), and IP Security (IPsec) sessions. For SRX Series chassis clusters, the fabric link can be any pair of Ethernet interfaces spanning the cluster; for J Series chassis clusters, any pair of Gigabit Ethernet interfaces.

Control Interfaces

The control interfaces provide the control link between the two nodes in the cluster and are used for routing updates and for control plane signal traffic, such as heartbeat and threshold information that triggers node failover. The control link is also used to synchronize the configuration between the nodes. When you submit configuration statements to the cluster, the configuration is automatically synchronized over the control link.

NOTE: For the SRX5000 and SRX3000 lines, you can connect two control links between the two nodes in the cluster to reduce the chance of control link failure. For more information, see “Connecting Dual Control Links” on page 893.

Related Topics

- Understanding Chassis Cluster Formation on page 869
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885
- Disabling Chassis Cluster on page 942

Creating an SRX Series Chassis Cluster—Overview

This section provides an overview of the basic steps to create an SRX Series chassis cluster.

For the basic steps to set up a J Series chassis cluster, see “Creating a J Series Chassis Cluster—Overview” on page 906.

Before You Begin

For background information, read:

- Understanding Chassis Cluster Formation on page 869
- Hardware Setup for SRX Series Chassis Clusters on page 890
- What Happens When You Enable Chassis Cluster on page 895

NOTE: For SRX5000 line chassis clusters, the placement and type of SPCs must match in the two devices. For SRX3000 line chassis clusters, the placement and type of SPCs, IOCs, and NPCs must match in the two devices. For SRX650, SRX240, SRX210, and SRX100 chassis clusters, the placement and type of GPIMs, XGPIMs, XPIMs, and Mini-PIMs (as applicable) must match in the two devices.

To create an SRX Series chassis cluster:

1. Physically connect a pair of the same kind of supported SRX Series devices together:
   a. To create the fabric link between two nodes in a cluster, connect any pair of Ethernet interfaces. For most SRX devices, the only requirement is that both interfaces be Gigabit Ethernet interfaces (or 10-Gigabit Ethernet interfaces). For SRX210 devices, both interfaces must be of a similar type (that is, a pair of Fast Ethernet interfaces or a pair of Gigabit Ethernet interfaces). For SRX100 devices, connect a pair of Fast Ethernet interfaces.

Figure 156 on page 891 shows nodes connected using built-in I/O ports for the fabric link.

b. Connect the control ports that you will use on each device (for example, fpc3 and fpc15, as shown in Figure 156 on page 891). For SRX3600, SRX3400, SRX650, and SRX240 devices, the control ports are dedicated Gigabit Ethernet ports. For SRX210 and SRX100 devices, the control port is the highest numbered port (fe-0/0/7).

When using dual control link functionality (SRX5000 and SRX3000 lines only), connect the two pairs of control ports that you will use on each device (for example, fpc3 and fpc15 for the first pair and fpc6 and fpc18 for the second, as shown in Figure 162 on page 893). For more information, see "Connecting Dual Control Links" on page 893.

For SRX5600 and SRX5800 devices, control ports should be on corresponding slots in the two devices, with the following slot numbering offsets:

Device	Offset
SRX5800	12 (for example, fpc3 and fpc15)
SRX5600	6 (for example, fpc3 and fpc9)

2. On the first device to be initialized in the cluster—this is the node that will form the cluster—connect to the console port.

For connection instructions, see the appropriate *Services Gateway Getting Started Guide*.

3. Configure the control ports (SRX5000 line only).

See "Configuring the Control Ports" on page 909.

4. Use CLI operational mode commands to enable clustering:
   a. Identify the cluster by giving it the cluster ID.
   b. Identify the node by giving it its own node ID and then reboot the system.

See "Setting the Node ID and Cluster ID" on page 910.

5. On the other device, connect to the console port and use CLI operational mode commands to enable clustering:
   a. Identify the cluster that the device is joining by setting the same cluster ID you set on the first node.
   b. Identify the node by giving it its own node ID and then reboot the system.

6. Configure the management interfaces on the cluster.

See "Configuring the Management Interface" on page 911.

7. Configure the cluster:

To use J-Web Quick Configuration, see:

a. Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912

b. Configuring Redundant Ethernet Interfaces—Quick Configuration on page 915

To configure the cluster with the CLI, see:

a. Configuring Chassis Cluster Information on page 919

b. Configuring the Fabric on page 920

c. Configuring Redundancy Groups on page 921

d. Configuring Redundant Ethernet Interfaces on page 923

e. Configuring Interface Monitoring on page 925

8. To initiate manual failover, see "Initiating a Manual Redundancy Group Failover" on page 927.
9. To configure conditional route advertisement over redundant Ethernet interfaces, see "Configuring Conditional Route Advertising" on page 929.
10. To verify the configuration, see "Verifying the Chassis Cluster Configuration" on page 931.

Related Topics

- Understanding Chassis Cluster Formation on page 869
- What Happens When You Enable Chassis Cluster on page 895
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912
- Configuring Redundant Ethernet Interfaces—Quick Configuration on page 915
- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920
- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879

Creating a J Series Chassis Cluster—Overview

This section provides an overview of the basic steps to create a J Series chassis cluster.

For the basic steps to set up an SRX Series chassis cluster, see “Creating an SRX Series Chassis Cluster—Overview” on page 904.

Before You Begin

For background information, read:

- Understanding Chassis Cluster Formation on page 869
- Hardware Setup for J Series Chassis Clusters on page 894
- What Happens When You Enable Chassis Cluster on page 895

NOTE: For J Series chassis clusters, the two nodes in a cluster must be identical models, but can have any combination of PIMs installed.

To create a J Series chassis cluster:

1. Physically connect a pair of the same kind of supported J Series devices together:
   a. To create the fabric link between two nodes in a cluster, connect any pair of Gigabit Ethernet interfaces, either the built-in interfaces or interfaces on the PIMs. The only requirement is that both interfaces be Gigabit Ethernet interfaces.

      Figure 164 on page 894 shows nodes connected using the built-in ge-0/0/1 interface for the fabric link.
   b. Connect the ge-0/0/3 interfaces together to create the control link.
2. On the first device to be initialized in the cluster—this is the node that will form the cluster—connect to the console port.

   For connection instructions, see the *J Series Services Routers Hardware Guide*.
3. Use CLI operational mode commands to enable clustering:
   a. Identify the cluster by giving it a cluster ID.
   b. Identify the node by giving it its own node ID and then reboot the system.

   See “Setting the Node ID and Cluster ID” on page 910.
4. On the other device, connect to the console port and use CLI operational mode commands to enable clustering:
   a. Identify the cluster that the device is joining by setting the same cluster ID you set on the first node.
   b. Identify the node by giving it its own node ID and then reboot the system.
5. Configure the management interfaces on the cluster.

   See “Configuring the Management Interface” on page 911.
6. Configure the cluster:

To use J-Web Quick Configuration, see:

a. Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912

b. Configuring Redundant Ethernet Interfaces—Quick Configuration on page 915

To configure the cluster with the CLI, see:

a. Configuring Chassis Cluster Information on page 919

b. Configuring the Fabric on page 920

c. Configuring Redundancy Groups on page 921

d. Configuring Redundant Ethernet Interfaces on page 923

e. Configuring Interface Monitoring on page 925

7. To initiate manual failover, see "Initiating a Manual Redundancy Group Failover" on page 927.
8. To configure conditional route advertisement over redundant Ethernet interfaces, see "Configuring Conditional Route Advertising" on page 929.
9. To verify the configuration, see "Verifying the Chassis Cluster Configuration" on page 931.

Related Topics

- Understanding Chassis Cluster Formation on page 869
- What Happens When You Enable Chassis Cluster on page 895
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912
- Configuring Redundant Ethernet Interfaces—Quick Configuration on page 915
- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920
- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879

Configuring the Control Ports

On SRX5600 and SRX5800 devices, by default, all control ports are disabled. You need to configure the control ports that you will use on each device to set up the control link or, optionally, the dual control links in a chassis cluster.

Before You Begin

For background information, read:

- About the Control Link on page 882
- About Dual Control LInks on page 882

This topic includes:

- CLI Configuration on page 909
- Related Topics on page 910

CLI Configuration

Use the following commands to configure the control ports for use as the control link for the chassis cluster:

```
user@host# set chassis cluster control-ports fpc4 port 0
user@host# set chassis cluster control-ports fpc10 port 0
```

Use the following commands to configure the control ports for use as dual control links for the chassis cluster.

```
user@host# set chassis cluster control-ports fpc4 port 0
user@host# set chassis cluster control-ports fpc10 port 0
user@host# set chassis cluster control-ports fpc6 port 1
user@host# set chassis cluster control-ports fpc12 port 1
```

Related Topics

- Verifying the Chassis Cluster Configuration on page 931

Setting the Node ID and Cluster ID

After connecting the two devices together, you configure a cluster ID and a node ID. A cluster ID identifies the cluster that the two nodes belong to. A node ID identifies a unique node within a cluster.

Before You Begin

Read:

- Understanding Chassis Cluster Formation on page 869
- Understanding Redundancy Groups on page 870

The system uses node IDs and cluster IDs to apply the correct configuration for each node when you use the apply-group command described in "Configuring Chassis Cluster Information" on page 919. The node ID and cluster ID statements are written to the EPROM, and when the system is rebooted, they take effect.

You can deploy up to 15 clusters in a Layer 2 domain. Each cluster is defined by a cluster-id value within the range of 1 through 15. A device can belong to only one cluster at any given time. Nodes in a cluster are numbered 0 and 1.

This topic includes:

- CLI Configuration on page 910
- Related Topics on page 911

CLI Configuration

To set the node IDs and cluster IDs, connect to each device through the console port and enter the following operational commands, then reboot the system.

- Enter the cluster ID and node ID information for the first node. If you want redundancy groups to be primary on this node when priority settings for both nodes are the same, make it node 0.

  ```
  user@host> set chassis cluster cluster-id 1 node 0
  warning: A reboot is required for chassis cluster to be enabled
  ```

- Enter the cluster ID and node ID for the other node. If you want redundancy groups to be secondary on this node when priority settings for both nodes are the same, make it node 1.

  ```
  user@host> set chassis cluster cluster-id 1 node 1 reboot
  Successfully enabled chassis cluster. Going to reboot now.
  ```

After you set the cluster ID and node ID for each node and the system reboots, the built-in interfaces are automatically renamed (see Table 112 on page 902). Use the show chassis cluster status operational command to view node status.

```
{primary:node1}
user@host# show chassis cluster status
Cluster ID: 3
Node name                  Priority     Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                   254         secondary no       no
    node1                   2           primary   no       no

Redundancy group: 1 , Failover count: 1
    node0                   101         Secondary no       no
    node1                   99          primary   no       no
```

When you complete the chassis cluster basic configuration, any subsequent configuration changes you make are automatically synchronized on both nodes.

Related Topics

- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Conditional Route Advertising on page 929
- Verifying the Chassis Cluster Configuration on page 931

Configuring the Management Interface

You must assign a unique IP address to each node in the cluster to provide network management access. This configuration is not replicated across the two nodes.

Before You Begin

Read:

- Understanding Chassis Cluster Formation on page 869
- Setting the Node ID and Cluster ID on page 910

In most SRX Series chassis clusters, the fxp0 interface is a dedicated port. For SRX210, SRX100, and all J Series chassis clusters, the fxp0 interface is repurposed from a built-in interface.

In a J Series chassis cluster, you configure management access to the cluster by defining a unique hostname for each node and assigning a unique IP address to the fxp0 interface on each node. The fxp0 interface is created when the system reboots the devices after you designate one node as the primary device and the other as the secondary device.

NOTE: If you try to access the nodes in a cluster over the network before you configure the fxp0 interface, you will lose access to the cluster.

This topic includes:

- CLI Configuration on page 912
- Related Topics on page 912

CLI Configuration

From the console port connection to the device you want to designate as the primary node, in configuration mode enter the following commands to name the node node0-router and assign IP address 10.1.1.1/24 to it:

```
user@host# set groups node0 system host-name node0-router
user@host# set groups node0 interfaces fxp0 unit 0 family inet address 10.1.1.1/24
```

From the console port connection of the device you want to designate the secondary node, in configuration mode enter the following commands to name the node node1-router and assign IP address 10.1.1.2/24 to it:

```
user@host# set groups node1 system host-name node1-router
user@host# set groups node1 interfaces fxp0 unit 0 family inet address 10.1.1.2/24
```

Enter the following command in configuration mode to apply these unique configurations to the appropriate node. (If you are migrating from a device to a cluster, delete the hostname from the configuration and then use the apply-groups command.)

This configuration is not replicated across the two nodes.

```
user@host# set apply-groups "${node}"
```

Related Topics

- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Conditional Route Advertising on page 929
- Verifying the Chassis Cluster Configuration on page 931

Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration

You can use J-Web Quick Configuration to quickly configure chassis clusters, as well as their redundancy groups. See Figure 172 on page 913.

CAUTION: Before you can use J-Web Quick Configuration to configure the chassis cluster parameters, you must:

- Set the cluster ID and node ID for each node.
- Configure management interfaces.

Before You Begin

Read:

- Setting the Node ID and Cluster ID on page 910 for details on information you must configure first
- Configuring Chassis Cluster Information on page 919

Figure 172 on page 913 shows the Chassis Cluster page.

Figure 172: Chassis Cluster Page

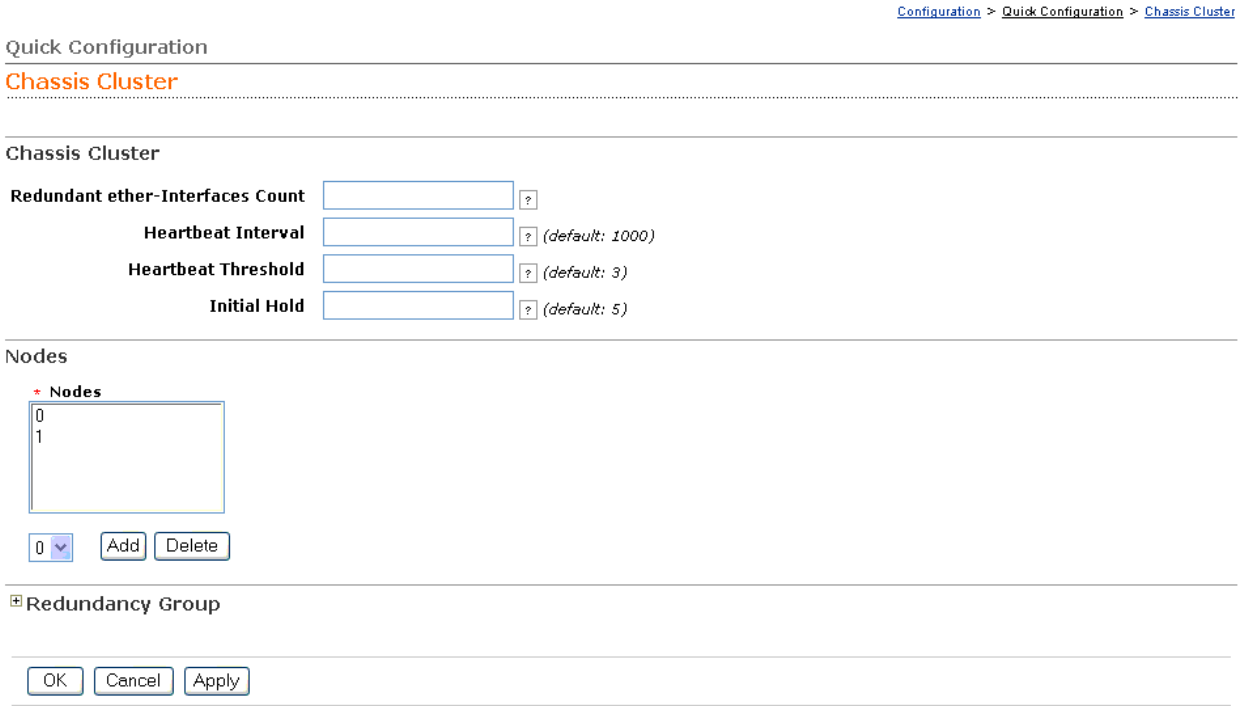

To configure chassis cluster with J-Web Quick Configuration:

1. Select **Configure > System Properties > Chassis Cluster**.
2. Fill in the parameter settings as described in Table 113 on page 914.
3. Click one of the following buttons:

- To apply the configuration and stay on the quick configuration Chassis Cluster page, click **Apply**.
- To apply the configuration and return to the main Configuration page, click **OK**.
- To cancel your entries and return to the main page, click **Cancel**.

Go on to "Configuring Redundant Ethernet Interfaces—Quick Configuration" on page 915.

Table 113: Chassis Cluster Options

Field	Function	Action
Chassis Cluster		
Redundant ether-Interface Count	Specifies the number of redundant Ethernet Interfaces to be created for the cluster.	Specify a value between **1** and **128**. For information about the maximum number of redundant Ethernet interfaces allowed for each device type, see Table 109 on page 880.
Heartbeat Interval	Specifies, in milliseconds, how often heartbeats are sent out to the other nodes in the cluster. The default heartbeat interval is 1000.	Specify a value between **1000** and **1200**.
Heartbeat Threshold	Specifies the number of consecutive missed heartbeats that indicate a cluster failure. The default number is 3.	Specify a value between **3** and **8**.
Nodes	Specifies the nodes to be used for the chassis cluster.	To specify a node, select the node from the list and click **Add.** To delete a node, select the node you want to delete and click **Delete**.

Group Number	Displays the selected redundancy group information or creates or modifies a redundancy group.	Select a group number from the list of configured redundancy groups.
		■ To modify information for a redundancy group, select the corresponding group number. Specify a value between **1** and **128**.
		■ To create a new redundancy group:
		  ■ Assign a group ID to the redundancy group. Specify a value between **1** and **128**.
		  ■ Configure the redundancy group with information described in the rest of this table.
		
		The maximum number of redundancy groups is equal to the number of redundant Ethernet interfaces that you configure.

Table 113: Chassis Cluster Options *(continued)*

Field	Function	Action
Gratuitous ARP Count	For a redundancy group, specifies the number of gratuitous Address Resolution Protocol (ARP) requests that a newly elected master sends out on the active redundant Ethernet interface child links to notify network devices of a change in mastership on the redundant Ethernet interface links. The default value is 4.	Specify a value between **1** and **16**.
Preempt	For a redundancy group, allows a node with a better priority to initiate a failover. By default, this feature is disabled. When disabled, a node with a better priority does not initiate a redundancy group failover (unless some other factor, such as faulty network connectivity identified for monitored interfaces, causes a failover).	Select the check box.
Priorities	For a redundancy group, specifies the priority value of each node. The redundancy group becomes primary on the node with the best (higher) priority. By default, this feature is disabled.	Select a node and a priority number, then click **Add**. The value is a number between **1** and **225**.
Interface Monitor	For a redundancy group, specifies the interfaces to be monitored by the redundancy group and their respective weights.	Select the interface from the list and enter its weight. The value is a number between **1** and **255**.

Related Topics

- Understanding Redundancy Groups on page 870
- Setting the Node ID and Cluster ID on page 910
- Configuring Redundant Ethernet Interfaces on page 923
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885

Configuring Redundant Ethernet Interfaces—Quick Configuration

You can use J-Web Quick Configuration to quickly configure redundant Ethernet (**reth**) interfaces. A redundant Ethernet interface is a pseudointerface that manages two “child” physical interfaces, one on each node of the cluster. Configuration parameters set for a redundant Ethernet interface are inherited by its child interfaces. A redundant Ethernet interface allows the chassis cluster to share one IP address across two links. When a redundancy group that the redundant Ethernet interface

belongs to fails over, its redundant Ethernet interfaces fail over with it and their interfaces on the new node become active.

Before You Begin

For background information, read “Understanding Redundant Ethernet Interfaces” on page 879

Before configuring redundant Ethernet interfaces, you must specify the reth-count so that reth interfaces will show in the configuration or J-Web interfaces screen. For example, to specify that there will be five redundant Ethernet interfaces, enter:

```
{primary:node1}
user@host# set chassis cluster reth-count 5
```

Note that the number of reth interfaces configured determines the number of redundancy groups that can be configured. For information about the maximum number of redundant Ethernet interfaces allowed for each device type, see Table 109 on page 880.

Figure 173 on page 917 shows the Interface page.

Figure 173: Interface Page

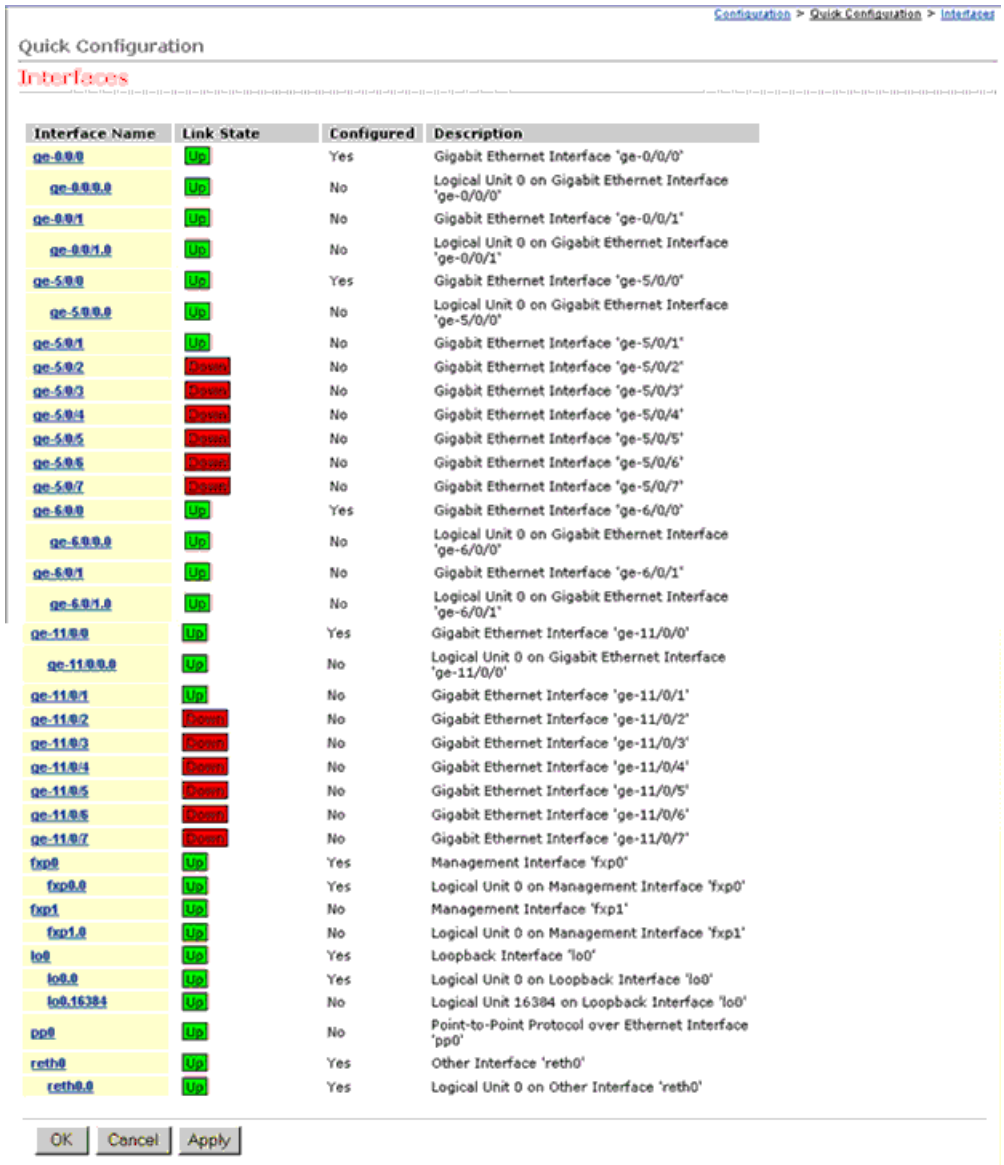

Figure 174 on page 918 shows the Redundant Ethernet Interface Configuration page.

Figure 174: Redundant Ethernet Interface Configuration page

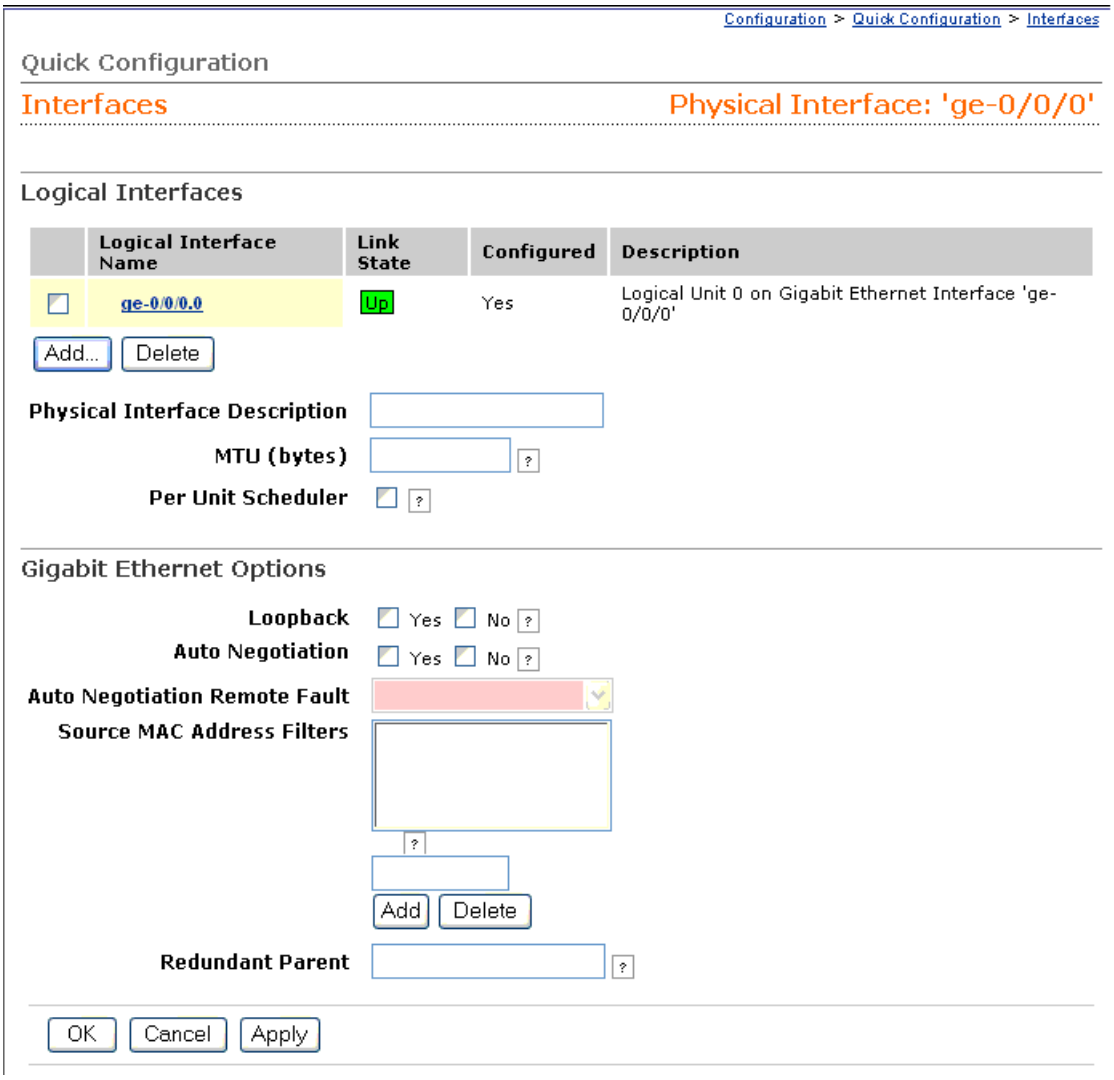

To configure redundant Ethernet interfaces with J-Web Quick Configuration:

1. Select **Configure > Interfaces**.
2. Click an interface name by which to group physical Ethernet interfaces for redundancy. See Figure 173 on page 917.
3. To add an interface to a redundant Ethernet interface, click **Add**.
4. Fill in the parameter settings for the logical interfaces as described in Table 114 on page 919. For details, see the *JUNOS Software Interfaces and Routing Configuration Guide*.
5. Fill in the information for **Redundant Parent** to specify the redundant parent Ethernet interface of the child physical interface.
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.

- To apply the configuration and return to the main Configuration page, click **OK**.
- To cancel your entries and return to the main page, click **Cancel**.

Table 114: Redundant Ethernet Interface Options

Field	Function	Action
Logical Interfaces		
Add Logical Interfaces	Defines one or more logical units that you connect to this physical redundant Ethernet interface. You must define at least one logical unit for a redundant Ethernet interface.	Click **Add**. To delete a logical interface, select the check box corresponding to the interface you want to delete and click **Delete**.
High Availability		
Redundancy Number	Specifies the number of the redundancy group to which the redundant interface belongs. Failover properties of the interface are inherited from the redundancy group.	Select a number from **0** through **225**.
Loop Back	Enables or disables the loopback option.	By default, the loopback is disabled. Select **Yes** to enable loopback mode.
Flow Control	Enables flow control on the Ethernet interface.	Select **Yes**.
Sources Filtering	Enables the filtering of media access control (MAC) source addresses to block all incoming packets to that interface.	By default, the source address filtering is disabled. Select **Yes**.
Redundant Parent	Specifies the name of the redundant Ethernet interface that a physical interface is associated with to form a redundant Ethernet interface pair.	Specify a redundant Ethernet interface name.

Configuring Chassis Cluster Information

For the chassis cluster configuration, you specify the number of redundant Ethernet interfaces that the cluster contains and the information used to monitor the "health" of the cluster.

Before You Begin

Read:

- Understanding Chassis Cluster Formation on page 869
- Setting the Node ID and Cluster ID on page 910
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 912

You must configure the redundant Ethernet interfaces count for the cluster in order for the redundant Ethernet interfaces that you configure to be recognized.

This topic includes:

- CLI Configuration on page 920
- Related Topics on page 920

CLI Configuration

Use the following command in configuration mode to define the number of redundant Ethernet interfaces for the cluster:

```
{primary:node1}
user@host# set chassis cluster reth-count 3
```

Related Topics

- Configuring the Fabric on page 920
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Conditional Route Advertising on page 929
- Verifying the Chassis Cluster Configuration on page 931

Configuring the Fabric

The fabric is the back-to-back data connection between the nodes in a cluster. Traffic on one node that needs to be processed on the other node or to exit through an interface on the other node passes over the fabric. Session state information also passes over the fabric.

Before You Begin

Read:

- Understanding Chassis Cluster Formation on page 869
- Setting the Node ID and Cluster ID on page 910
- Understanding the Control Plane on page 881

In most SRX Series devices in a chassis cluster, you can configure any pair of Gigabit Ethernet interfaces or any pair of 10-Gigabit interfaces to serve as the fabric between nodes. Interfaces on SRX210 devices are Fast Ethernet or Gigabit Ethernet (the paired interfaces must be of a similar type) and all interfaces on SRX100 devices are Fast Ethernet interfaces.

In a J Series chassis cluster, you can configure any pair of Gigabit Ethernet interfaces to serve as the fabric between nodes.

For details about port and interface usage for management, control, and fabric links, see Table 111 on page 897 and Table 112 on page 902.

You cannot configure filters, policies, or services on the fabric interface. Fragmentation is not supported on the fabric link. The maximum transmission unit (MTU) size is 8,980 bytes. We recommend that no interface in the cluster exceed this MTU. Jumbo frame support on the member links is enabled by default.

This topic includes:

- CLI Configuration on page 921
- Related Topics on page 921

CLI Configuration

Enter the following commands to join ge-0/0/1 on one node in the cluster and ge-7/0/1 on the other to form the fabric:

```
{primary:node0}
user@host# set interfaces fab0 fabric-options member-interfaces ge-0/0/1
{secondary:node1}
user@host# set interfaces fab1 fabric-options member-interfaces ge-7/0/1
```

Related Topics

- Configuring Chassis Cluster Information on page 919
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Conditional Route Advertising on page 929
- Verifying the Chassis Cluster Configuration on page 931

Configuring Redundancy Groups

A redundancy group is an abstract entity that includes and manages a collection of objects. A redundancy group can be primary on only one node at a time.

Before You Begin

Read:

- Understanding Redundancy Groups on page 870
- Setting the Node ID and Cluster ID on page 910
- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920

Before you can create redundant Ethernet interfaces, you must first create their redundancy groups.

This topic includes:

- CLI Configuration on page 922
- Related Topics on page 923

CLI Configuration

Use the following command in configuration mode to specify the number of gratuitous Address Resolution Protocol (ARP) requests that an interface can send to notify other network devices of its presence after the redundancy group it belongs to has failed over:

```
{primary:node1}
user@host# set chassis cluster redundancy-group 1 gratuitous-arp-count 4
```

Use the following command in configuration mode to identify an interface to be monitored by a specific redundancy group and give it a weight. You can configure a redundancy group to monitor any interfaces, not only those belonging to its redundant Ethernet interfaces.

```
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor fe-3/1/1/1
  weight 100
```

Use the following command in configuration mode to specify an IP address to be monitored by a specific redundancy group and give it a weight. You can configure a redundancy group to monitor any IP address reachable through a redundant Ethernet interface.

```
{primary:node0}
user@host# set chassis cluster redundancy-group 1 ip-monitoring family inet
  10.1.1.10 weight 100
```

Use the following command in configuration mode to set a monitoring weight and threshold for IP monitoring on the redundancy group. The weight will be deducted from the redundancy group’s threshold when the IP monitoring threshold value reaches 0.

```
{primary:node0}
user@host# set chassis cluster redundancy-group 1 ip-monitoring weight 100
  threshold 200
```

Use the following commands in configuration mode to specify a redundancy group's priority for primacy on each node of the cluster. The higher number takes precedence.

```
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 1 priority 100
{secondary:node0}
user@host# set chassis cluster redundancy-group 1 node 0 priority 200
```

Use the following command in configuration mode to specify if a node with a better (higher) priority can initiate a failover to become primary for the redundancy group:

```
{primary:node1}
user@host# set chassis cluster redundancy-group 1 preempt
```

Use the following command in configuration mode to specify the minimum interval to be allowed between back-to-back failovers for the redundancy group:

```
{primary:node1}
user@host# set chassis cluster redundancy-group 1 hold-down-interval 6
```

Related Topics

- Configuring Redundant Ethernet Interfaces on page 923
- Verifying the Chassis Cluster Configuration on page 931

Configuring Redundant Ethernet Interfaces

A redundant Ethernet interface is a pseudointerface that contains two physical interfaces, one from each node of the cluster. To create a redundant Ethernet interface, you configure the two physical interfaces independently. You configure the rest of the configuration that pertains to them at the level of the redundant Ethernet interface, and each of the child interfaces inherits this configuration.

NOTE: A redundant Ethernet interface is referred to as a *reth* in configuration commands.

Before You Begin

Read:

- Setting the Node ID and Cluster ID on page 910
- Configuring Chassis Cluster Information on page 919
- Configuring the Fabric on page 920
- Configuring Redundancy Groups on page 921
- Understanding Redundant Ethernet Interfaces on page 879

This topic includes:

- CLI Configuration on page 923
- Related Topics on page 924

CLI Configuration

Use the following commands to bind redundant child physical interfaces to reth1:

```
{primary:node1}
user@host# set interfaces ge-0/0/0 gigether-options redundant-parent reth1
{primary:node1}
user@host# set interfaces ge-7/0/0 gigether-options redundant-parent reth1
{primary:node1}
user@host# set interfaces fe-1/0/0 fast-ether-options redundant-parent reth1
```

```
{primary:node1}
user@host# set interfaces fe-8/0/0 fast-ether-options redundant-parent reth1
```

Use the following commands to:

- Add reth 1 to redundancy group 1.
- Set the MTU size to 1500 bytes.
- Assign IP address 10.1.1.3/24 to reth1.

```
{primary:node1}
user@host# set interfaces reth1 redundant-ether-options redundancy-group 1
{primary:node1}
user@host# set interfaces reth1 unit 0 family inet mtu 1500
{primary:node1}
user@host# set interfaces reth1 unit 0 family inet address 10.1.1.3/24
```

Use the following command to associate reth1.0 with a security zone named Trust. Security zone configuration is the same for redundant Ethernet interfaces as for any other interface.

```
{primary:node1}
user@host# set security zones security-zone Trust interfaces reth1.0
```

Related Topics

- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Conditional Route Advertising on page 929
- Verifying the Chassis Cluster Configuration on page 931

Configuring Control Link Recovery

Although control link recovery is disabled by default, you can configure it.

Before You Begin

For background information, read:

- About the Control Link on page 882
- About Control Link Failure and Recovery on page 884

This topic includes:

- CLI Configuration on page 925
- Related Topics on page 925

CLI Configuration

Use the following command to configure control link recovery:

```
{primary:node1}
user@host# set chassis cluster control-link-recovery
```

Control link recovery is done automatically by the system. After the control link recovers, the system takes the following actions:

- It checks whether it receives at least 30 consecutive heartbeats on the control link, or in the case of dual control links (SRX5000 and SRX3000 lines only), on either control link. This is to ensure that the control link is not flapping and is perfectly healthy.
- After it determines that the control link is healthy, the system issues an automatic reboot on the node that was disabled when the control link failed. When the disabled node reboots, it can rejoin the cluster. There is no need for any manual intervention.

Related Topics

- Verifying the Chassis Cluster Configuration on page 931

Configuring Interface Monitoring

Redundancy group failover is triggered by the results from monitoring the health of interfaces that belong to the redundancy group. When you assign a weight to an interface to be monitored, the system monitors the interface for availability. If a physical interface fails, the weight is deducted from the corresponding redundancy group's threshold. Every redundancy group has a threshold of 255. If the threshold hits 0, a failover is triggered. Failover is triggered even if the redundancy group is in manual failover mode and preempt is not enabled.

Before You Begin

For background information, read:

- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879

This topic includes:

- CLI Configuration on page 926
- Related Topics on page 926

CLI Configuration

Use the following command to set interface monitoring on ge-7/0/3:

```
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor ge-7/0/3
   weight 255
```

Related Topics

- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Redundancy Groups on page 921

Configuring a Dampening Time Between Back-to-Back Redundancy Group Failovers

Although back-to-back redundancy group failovers have a default dampening time, you can configure it. For redundancy group 0, the default interval is 5 minutes (300 seconds), with a range of 300 through 1800 seconds. For redundancy group *x* (redundancy groups 1 through 128), the default interval is 1 second, with a range of 0 through 1800 seconds.

Before You Begin

For background information, read:

- Understanding Redundancy Groups on page 870
- Understanding Redundant Ethernet Interfaces on page 879
- Understanding Failover on page 888

This topic includes:

- CLI Configuration on page 926
- Related Topics on page 926

CLI Configuration

Use the following command to set the minimum interval (in seconds) to be allowed between back-to-back failovers for redundancy group 0:

```
{primary:node1}
user@host# set chassis cluster redundancy-group 0 hold-down-interval 420
```

Related Topics

- Configuring Redundant Ethernet Interfaces on page 923
- Configuring Redundancy Groups on page 921

■ Initiating a Manual Redundancy Group Failover on page 927

Initiating a Manual Redundancy Group Failover

You can initiate a failover manually with the request command. A manual failover bumps up the priority of the redundancy group for that member to 255.

Before You Begin

For background information, read:

- Configuring Redundancy Groups on page 921
- Configuring Redundant Ethernet Interfaces on page 923
- Configuring a Dampening Time Between Back-to-Back Redundancy Group Failovers on page 926

CAUTION: Be cautious and judicious in your use of redundancy group 0 manual failovers. A redundancy group 0 failover implies a Routing Engine (RE) failover, in which case all processes running on the primary node are killed and then spawned on the new master Routing Engine (RE). This failover could result in loss of state, such as routing state, and degrade performance by introducing system churn.

This topic includes:

■ CLI Configuration on page 927

CLI Configuration

Use the show command to display the status of nodes in the cluster:

```
{primary:node0}
user@host> show chassis cluster status redundancy-group 0
Cluster ID: 9
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                   254         primary        no       no
    node1                   1           secondary      no       no
```

Output to this command indicates that node 0 is primary.

Use the request command to trigger a failover and make node 1 primary:

```
{primary:node0}
user@host> request chassis cluster failover redundancy-group 0 node 1
- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
Initiated manual failover for redundancy group 0
```

Use the show command to display the new status of nodes in the cluster:

```
{secondary-hold:node0}
user@host> show chassis cluster status redundancy-group 0
Cluster ID: 9
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 2
    node0                   254         secondary-hold no       yes
    node1                   255         primary        no       yes
```

Output to this command shows that node 1 is now primary and node 0 is in the secondary-hold state. After 5 minutes, node 0 will transition to the secondary state.

You can reset the failover for redundancy groups by using the request command. This change is propagated across the cluster.

```
{secondary-hold:node0}
user@host> request chassis cluster failover reset redundancy-group 0 node 0
node0:
--------------------------------------------------------------------------
No reset required for redundancy group 0.

node1:
--------------------------------------------------------------------------
Successfully reset manual failover for redundancy group 0
```

You cannot trigger a back-to-back failover until the 5-minute interval expires.

```
{secondary-hold:node0}
user@host> request chassis cluster failover redundancy-group 0 node 0
node0:
--------------------------------------------------------------------------
Manual failover is not permitted as redundancy-group 0 on node0 is in
secondary-hold state.
```

Use the show command to display the new status of nodes in the cluster:

```
{secondary-hold:node0}
user@host> show chassis cluster status redundancy-group 0
Cluster ID: 9
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 2
    node0                   254         secondary-hold no       no
    node1                   1           primary        no       no
```

Output to this command shows that a back-to-back failover has not occurred for either node.

After doing a manual failover, you must issue the reset failover command before requesting another failover.

When the primary node fails and comes back up, election of the primary node is done based on regular criteria (priority and preempt).

Configuring Conditional Route Advertising

Route advertisement over redundant Ethernet interfaces in a chassis cluster is complicated by the fact that the active node in the cluster can change dynamically. Conditional route advertisement enables you to advertise routes in such a way that incoming traffic from the core network is attracted to the Border Gateway Protocol (BGP) interface that exists on the same node as the currently active redundant Ethernet interface. In this way, traffic is processed by the active node and does not traverse the fabric interface between nodes. You do this by manipulating the BGP attribute at the time routes are advertised by BGP.

Before You Begin

For background information, read "Understanding Chassis Cluster Formation" on page 869

The goal of conditional route advertisement in a chassis cluster is to ensure that incoming traffic from the upstream network arrives on the node that is on the currently active redundant Ethernet interface. To understand how this works, bear in mind that in a chassis cluster, each node has its own set of interfaces. Figure 175 on page 930 shows a typical scenario, with a redundant Ethernet interface connecting the corporate LAN, through a chassis cluster, to an external network segment.

Figure 175: Conditional Route Advertising

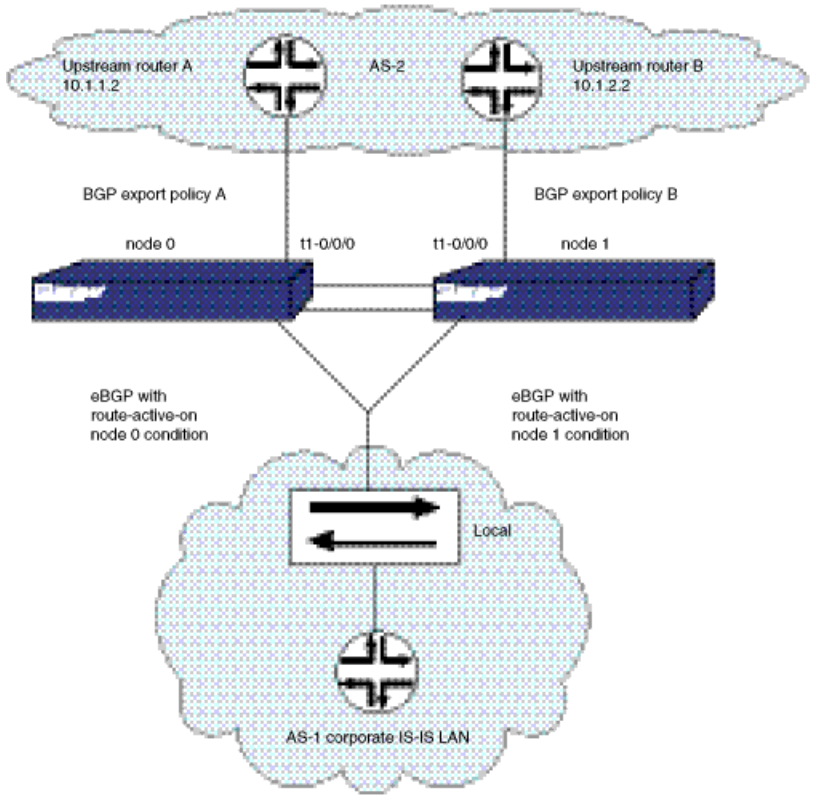

CLI Configuration

As illustrated in Figure 175 on page 930, routing prefixes learned from the redundant Ethernet interface through the interior gateway protocol (IGP) are advertised toward the network core using BGP. Two BGP sessions are maintained, one off interface t1-1/0/0 and one off t1-1/0/1 for BGP multihoming. All routing prefixes are advertised on both sessions. Thus, for a route advertised by BGP, learned over a redundant Ethernet interface, if the active redundant Ethernet interface is on the same node as the BGP session, you advertise the route with a “good” BGP attribute.

To achieve this behavior, you apply a policy to BGP before exporting routes. An additional term in the policy match condition determines the current active redundant Ethernet interface child interface of the next hop before making the routing decision. When the active status of a child redundant Ethernet interface changes, BGP reevaluates the export policy for all routes affected.

```
{primary:node1}
user@host# set policy-options policy-statement reth-nh-active-on-0 term ospf-on-0
   from protocol ospf
{primary:node1}
```

```
user@host# set policy-options policy-statement reth-nh-active-on-0 term ospf-on-0
   from condition reth-nh-active-on-0
{primary:node1}
user@host#set policy-options policy-statement reth-nh-active-on-0 term ospf-on-0
   then metric 10
{primary:node1}
user@host# set policy-options policy-statement reth-nh-active-on-0 term ospf-on-0
   then accept
{primary:node1}
user@host# set policy-options condition reth-nh-active-on-0 route-active-on node0
```

The condition statement in this configuration works as follows. The command states that any routes evaluated against this condition will pass only if:

- The routes have a redundant Ethernet interface as their next-hop interface.
- The current child interface of the redundant Ethernet interface above is active at node 0 (as specified by the route-active-on node0 keyword).

```
{primary:node1}
user@host# set policy-options condition reth-nh-active-on-0 route-active-on
   node0
```

Note that a route might have multiple equal-cost next hops, and those next hops might be redundant Ethernet interfaces, regular interfaces, or a combination of both. The route still satisfies the requirement that it has a redundant Ethernet interface as its next hop.

If you use the BGP export policy set for node 0 in the previous example command, only OSPF routes that satisfy the following requirements will be advertised through the session:

- The OSPF routes have a redundant Ethernet interface as their next hop.
- The current child interface of the redundant Ethernet interface is currently active at node 0.

You must also create and apply a separate policy statement for the other BGP session by using this same process.

In addition to the BGP MED attribute, you can define additional BGP attributes, such as origin-code, as-path, and community.

Related Topics

"Verifying the Chassis Cluster Configuration" on page 931

Verifying the Chassis Cluster Configuration

To verify the chassis cluster configuration, perform the following tasks:

- Verifying the Chassis Cluster on page 932
- Verifying Chassis Cluster Interfaces on page 932

- Verifying Chassis Cluster Statistics on page 933
- Verifying Chassis Cluster Status on page 936
- Verifying Chassis Cluster Redundancy Group Status on page 937

Verifying the Chassis Cluster

Purpose Display chassis cluster verification options.

Action From the CLI, enter the show chassis cluster ? command:

```
{primary:node1}
user@host> show chassis cluster ?
Possible completions:
  interfaces           Display chassis-cluster interfaces
  statistics           Display chassis-cluster traffic statistics
  status               Display chassis-cluster status
```

What it Means The output shows a list of all chassis cluster verification parameters. Verify the following information:

- Interfaces—Displays information about chassis cluster interfaces.
- Statistics—Displays information about chassis cluster services and interfaces.
- Status—Displays failover status about nodes in a cluster.

Related Topics

- Verifying Chassis Cluster Interfaces on page 932
- Verifying Chassis Cluster Statistics on page 933
- Verifying Chassis Cluster Status on page 936
- Verifying Chassis Cluster Redundancy Group Status on page 937

Verifying Chassis Cluster Interfaces

Purpose Display information about chassis cluster interfaces.

Action From the CLI, enter the show chassis cluster interfaces command:

```
{primary:node1}
user@host> show chassis cluster interfaces
Control link 0 name: em0

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          1

Interface Monitoring:
    Interface         Weight    Status    Redundancy-group
    ge-6/0/0          200       Up        1

{primary:node1}
```

```
user@host> show chassis cluster interfaces
Control link 0 name: em0
Control link 1 name: em1

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          1

Interface Monitoring:
    Interface         Weight    Status    Redundancy-group
    ge-6/0/0          200       Up        1
```

What it Means The first output shows the state of fxp1, the control link between the nodes, and provides information about the link state; it also shows the state of the fabric interface between the nodes and provides information about traffic on that link. The second output shows this information for a chassis cluster that uses dual control links (SRX5000 and SRX3000 lines only).

Verifying Chassis Cluster Statistics

Purpose Display information about chassis cluster services and interfaces.

Action From the CLI, enter the show chassis cluster statistics command:

```
{primary:node1}
user@host> show chassis cluster statistics

Control link statistics:
    Control link 0:
        Heartbeat packets sent: 798
        Heartbeat packets received: 784
Fabric link statistics:
    Probes sent: 793
    Probes received: 0
Services Synchronized:
    Service name                              RTOs sent    RTOs received
    Translation context                       0            0
    Incoming NAT                              0            0
    Resource manager                          0            0
    Session create                            0            0
    Session close                             0            0
    Session change                            0            0
    Gate create                               0            0
    Session ageout refresh requests           0            0
    Session ageout refresh replies            0            0
    IPSec VPN                                 0            0
    Firewall user authentication              0            0
    MGCP ALG                                  0            0
    H323 ALG                                  0            0
    SIP ALG                                   0            0
    SCCP ALG                                  0            0
    PPTP ALG                                  0            0
    RTSP ALG                                  0            0

{primary:node1}
user@host> show chassis cluster statistics

Control link statistics:
```

```
    Control link 0:
        Heartbeat packets sent: 258689
        Heartbeat packets received: 258684
    Control link 1:
        Heartbeat packets sent: 258689
        Heartbeat packets received: 258684
Fabric link statistics:
    Probes sent: 258681
    Probes received: 258681
Services Synchronized:
    Service name                              RTOs sent    RTOs received
    Translation context                       0            0
    Incoming NAT                              0            0
    Resource manager                          0            0
    Session create                            1            0
    Session close                             1            0
    Session change                            0            0
    Gate create                               0            0
    Session ageout refresh requests           0            0
    Session ageout refresh replies            0            0
    IPSec VPN                                 0            0
    Firewall user authentication              0            0
    MGCP ALG                                  0            0
    H323 ALG                                  0            0
    SIP ALG                                   0            0
    SCCP ALG                                  0            0
    PPTP ALG                                  0            0
    RPC ALG                                   0            0
    RTSP ALG                                  0            0
    RAS ALG                                   0            0
    MAC address learning                      0            0
    GPRS GTP                                  0            0

{primary:node1}
user@host> show chassis cluster statistics

Control link statistics:
    Control link 0:
        Heartbeat packets sent: 82371
        Heartbeat packets received: 82321
    Control link 1:
        Heartbeat packets sent: 0
        Heartbeat packets received: 0
```

What it Means The first output shows the control link statistics (heartbeats sent and received), the fabric link statistics (probes sent and received), and the number of RTOs sent and received for services. The second output shows these statistics for a chassis cluster that uses dual control links (SRX5000 and SRX3000 lines only). The output for the SRX5000 and SRX3000 lines will always show control link 0 and control link 1 statistics, even though only one control link is active or working, as shown in the third (partial) output.

Purpose Clear displayed information about chassis cluster services and interfaces.

Action From the CLI, enter the clear chassis cluster statistics command:

```
{primary:node1}
user@host> clear chassis cluster statistics
```

```
Cleared control-plane statistics
Cleared data-plane statistics
```

Purpose Display chassis cluster control-plane statistics.

Action From the CLI, enter the show chassis cluster control-plane statistics command:

```
{primary:node1}
user@host> show chassis cluster control-plane statistics

Control link statistics:
    Control link 0:
        Heartbeat packets sent: 124
        Heartbeat packets received: 125
Fabric link statistics:
    Probes sent: 124
    Probes received: 125

{primary:node1}
user@host> show chassis cluster control-plane statistics

Control link statistics:
    Control link 0:
        Heartbeat packets sent: 258698
        Heartbeat packets received: 258693
    Control link 1:
        Heartbeat packets sent: 258698
        Heartbeat packets received: 258693
Fabric link statistics:
    Probes sent: 258690
    Probes received: 258690
```

What it Means The first output shows the control link statistics (heartbeats sent and received) and the fabric link statistics (probes sent and received). The second output shows these statistics for a chassis cluster that uses dual control links (SRX5000 and SRX3000 lines only).

Purpose Clear displayed chassis cluster control plane statistics.

Action From the CLI, enter the clear chassis cluster control—plane statistics command:

```
{primary:node1}
user@host> clear chassis cluster control-plane statistics

Cleared control-plane statistics
```

Purpose Display chassis cluster data plane statistics.

Action From the CLI, enter the show chassis cluster data-plane statistics command:

```
{primary:node1}
user@host> show chassis cluster data-plane statistics

Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
```

```
    Resource manager                          0           0
    Session create                            0           0
    Session close                             0           0
    Session change                            0           0
    Gate create                               0           0
    Session ageout refresh requests           0           0
    Session ageout refresh replies            0           0
    IPSec VPN                                 0           0
    Firewall user authentication              0           0
    MGCP ALG                                  0           0
    H323 ALG                                  0           0
    SIP ALG                                   0           0
    SCCP ALG                                  0           0
    PPTP ALG                                  0           0
    RTSP ALG                                  0           0
```

What it Means The output shows the number of RTOs sent and received for services.

Purpose Clear displayed chassis cluster data plane statistics.

Action From the CLI, enter the clear chassis cluster data-plane statistics command:

```
{primary:node1}
user@host> clear chassis cluster data-plane statistics

Cleared data-plane statistics
```

Verifying Chassis Cluster Status

Purpose Display the failover status of a chassis cluster.

Action From the CLI, enter the show chassis cluster status command:

```
{primary:node1}
user@host> show chassis cluster status
Cluster ID: 3
    Node name               Priority    Status    Preempt  Manual failover

 Redundancy-group: 0, Failover count: 1
    node0                   254         primary   no       no
    node1                   2           secondary no       no

 Redundancy-group: 1, Failover count: 1
    node0                   254         primary   no       no
    node1                   1           secondary no       no

{primary:node1}
user@host> show chassis cluster status
Cluster ID: 15
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 5
    node0                   200         primary        no       no
    node1                   0           lost           n/a      n/a

Redundancy group: 1 , Failover count: 41
    node0                   101         primary        no       no
    node1                   0           lost           n/a      n/a
```

```
{primary:node1}
user@host> show chassis cluster status
Cluster ID: 15
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 5
    node0                   200         primary        no       no
    node1                   0           unavailable    n/a      n/a

Redundancy group: 1 , Failover count: 41
    node0                   101         primary        no       no
    node1                   0           unavailable    n/a      n/a
```

What it Means The three outputs show the failover status of the chassis cluster in addition to information about the chassis cluster redundancy groups. The second output shows node 1 as *lost*, meaning that node 1 lost contact with node 0 through the control link. Loss of contact is most likely to occur when both nodes are in a cluster and due to a control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. The third output shows node 1 as *unavailable*, meaning that node 0 is not aware of node 1 in the cluster. That is, node 0 has not received a single heartbeat over the control link from node 1 since node 1 booted up. Unavailability is most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.

Purpose Clear the failover status of a chassis cluster.

Action From the CLI, enter the clear chassis cluster failover-count command:

```
{primary:node1}
user@host> clear chassis cluster failover-count
Cleared failover-count for all redundancy-groups
```

Verifying Chassis Cluster Redundancy Group Status

Purpose Display the failover status of a chassis cluster redundancy group.

Action From the CLI, enter the show chassis cluster status redundancy-group command:

```
{primary:node1}
user@host> show chassis cluster status redundancy-group 2
Cluster ID: 14
    Node name          Priority    Status    Preempt  Manual failover

 Redundancy-Group: 2, Failover count: 1
    node0              50          secondary no       no
    node1              100         primary   no       no
```

What it Means The output shows state and priority of both nodes in a cluster and indicates whether the primary has been preempted or whether there has been a manual failover.

SNMP Failover Traps

Chassis clustering supports SNMP traps, which are triggered whenever there is a redundancy group failover.

Before You Begin

For background information, read about enterprise-specific MIBs for JUNOS Software in the *JUNOS Network Management Configuration Guide*.

The trap message can help you troubleshoot failovers. It contains the following information:

- The cluster ID and node ID
- The reason for the failover
- The redundancy group that is involved in the failover
- The redundancy group's previous state and current state

These are the different states that a cluster can be in at any given instant: hold, primary, secondary-hold, secondary, ineligible, and disabled. Traps are generated for the following state transitions (only a transition from a hold state does not trigger a trap):

- primary <–> secondary
- primary –> secondary-hold
- secondary-hold –> secondary
- secondary –> ineligible
- ineligible –> disabled
- ineligible –> primary
- secondary –> disabled

A transition can be triggered due to any event, such as interface monitoring, SPU monitoring, failures, and manual failovers.

The trap will be forwarded over the control link if the outgoing interface is on a node different from the node on the Routing Engine that generates the trap.

You can specify that a trace log be generated by setting the traceoptions flag snmp statement.

Related Topics

- Understanding Failover on page 888

Upgrading Devices in a Chassis Cluster

Devices in a chassis cluster can be upgraded separately one at a time; some models allow one device after the other to be upgraded using failover and an in-service software upgrade (ISSU) to reduce the operational impact of the upgrade.

Upgrading Each Device Separately

To upgrade each device in a chassis cluster separately:

NOTE: During this type of chassis cluster upgrade, a service disruption of about 3 to 5 minutes will occur.

1. Load the new image file on node 0.
2. Perform the image upgrade without rebooting the node by entering:

   user@host> request system software add *image_name*

3. Load the new image file on node 1.
4. Repeat Step 2.
5. Reboot both nodes simultaneously.

Low-Impact ISSU Chassis Cluster Upgrades

For some platforms, devices in a chassis cluster can be upgraded without a service disruption using an in-service software upgrade (ISSU). The chassis cluster ISSU feature allows both devices in a cluster to be upgraded or downgraded from supported JUNOS versions with a traffic impact similar to that of redundancy group failovers.

Before You Begin

1. Note that in-service software upgrades (ISSUs) are available only for JUNOS Release 9.6 and later.
2. Before starting an ISSU, you should fail over all redundancy groups so that they are all active on only one device. See “Initiating a Manual Redundancy Group Failover” on page 927 for more information.
3. We also recommend that routing protocols graceful restart be enabled prior to starting an ISSU.

Once all redundancy groups are active on one device, the upgrade is initiated by using a request command:

1. Fail over all redundancy groups to one device.
2. Start the ISSU by entering the following command:

   user@host> request system software in-service-upgrade *image_name* reboot

   If reboot is not included in the command, you will need to manually reboot each device as the ISSU completes updating the software image.
3. Wait for both devices to complete the upgrade, then verify that all policies, zones, redundancy groups, and other runtime objects (RTOs) return to their correct states. Also verify that both devices in the cluster are running the new JUNOS Software build.

NOTE: During the upgrade, both devices might experience redundancy group failovers, but traffic will not be disrupted. Each device validates the package and checks version compatibility before doing the upgrade. If the system finds that the new package is not version compatible with the currently installed version, the device will refuse the upgrade, or prompt you to take corrective action. Sometimes a single feature is not compatible, in which case the upgrade software will prompt you to either abort the upgrade or turn off the feature before doing the upgrade.

If the ISSU fails to complete and only one device in the cluster has been upgraded, you can roll back to the previous configuration on that device alone by using the following commands on the upgraded device:

1. request chassis cluster in-service-upgrade abort
2. request system software rollback
3. request system reboot

The ISSU command has one option: *no-old-master-upgrade*. This option leaves the current master device in a nonupgraded state, which is a precaution against service failure. The no-old-master-upgrade option allows routing control to be quickly returned to the old master device if the newly upgraded device does not operate correctly.

Use of the no-old-master-upgrade option will require you to run a standard upgrade on the old master device after the ISSU is completed on the backup device.

If you use the no-old-master-upgrade option, when the backup device completes its upgrade and you are confident that the new build is operating as expected, then upgrade the old master as follows:

1. Run request system software add *image_name*
2. Run request chassis cluster in-service-upgrade abort to stop the ISSU process.
3. Run request system reboot

This feature is available only via the command-line interface. See the "request system software in-service-upgrade" section of the *JUNOS Software CLI Reference*.

Automatic Failback

If you want redundancy groups to automatically return to node 0 as the primary after the ISSU is complete, you must set the redundancy group priority such that node 0 is primary and enable the preempt option. Note that this method will work for all redundancy groups except redundancy group 0. You must manually fail over redundancy group 0. To set the redundancy group priority and enable the preempt option, see "Configuring Redundancy Groups" on page 921. To manually fail over a redundancy group, see "Initiating a Manual Redundancy Group Failover" on page 927.

Troubleshooting ISSU Failure

Certain circumstances might cause an ISSU attempt to fail. This section explains two of them.

- If you attempt to upgrade a device pair running a JUNOS Software image earlier than 9.6, the ISSU will fail without changing anything about either device in the cluster. Devices running JUNOS releases earlier than 9.6 must be upgraded separately using individual device upgrade procedures.
- If the secondary device experiences a power-off condition before it boots up using the new image specified when the ISSU is initiated, when power is restored the newly upgraded device will still be waiting to end the ISSU. To end the ISSU on the secondary device, run request chassis cluster in-service-upgrade abort followed by reboot to abort the ISSU on that device.

Mismatched Control Link Statistics During an ISSU

When using dual control links (supported on the SRX5000 and SRX3000 lines only), mismatched control link statistics might be reported with the show chassis cluster statistics and show chassis cluster control-plane statistics commands while you run an ISSU with nodes on devices running different releases. (ISSUs are available in JUNOS Release 9.6 and later and dual control links are available in JUNOS Release 10.0 and later.) For example, assume that one node on a device is running JUNOS Release 9.6 and another node on a device is running JUNOS Release 10.0. In this example, a mismatch might occur because the latter device will be sending heartbeats on both control links, but the other device will receive heartbeats only on one control link.

Related Topics

- Understanding Chassis Cluster Formation on page 869
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885
- What Happens When You Enable Chassis Cluster on page 895
- Initiating a Manual Redundancy Group Failover on page 927

Disabling Chassis Cluster

To disable chassis cluster, enter the following command:

```
{primary:node1}
user@host> set chassis cluster disable reboot
Successfully disabled chassis cluster. Going to reboot now.
```

After the system reboots, the chassis cluster is disabled.

Related Topics

- Understanding Chassis Cluster Formation on page 869
- Understanding the Control Plane on page 881
- Understanding the Data Plane on page 885
- What Happens When You Enable Chassis Cluster on page 895

Chassis Cluster Configuration Scenarios

This section describes the following deployment scenarios:

- Active/Passive Chassis Cluster Scenario on page 942
- Active/Passive Chassis Cluster with IPsec Tunnel Scenario on page 947
- Asymmetric Routing Chassis Cluster Scenario on page 954

Before You Begin

For background information, read:

- Understanding Chassis Cluster Formation on page 869
- Hardware Setup for SRX Series Chassis Clusters on page 890
- Hardware Setup for J Series Chassis Clusters on page 894
- What Happens When You Enable Chassis Cluster on page 895
- Creating an SRX Series Chassis Cluster—Overview on page 904
- Creating a J Series Chassis Cluster—Overview on page 906

Active/Passive Chassis Cluster Scenario

In this case, a single device in the cluster is used to route all traffic while the other device is used only in the event of a failure (see Figure 176 on page 943). When a failure occurs, the backup device becomes master and controls all forwarding.

Figure 176: Active/Passive Chassis Cluster Scenario (J Series Devices)

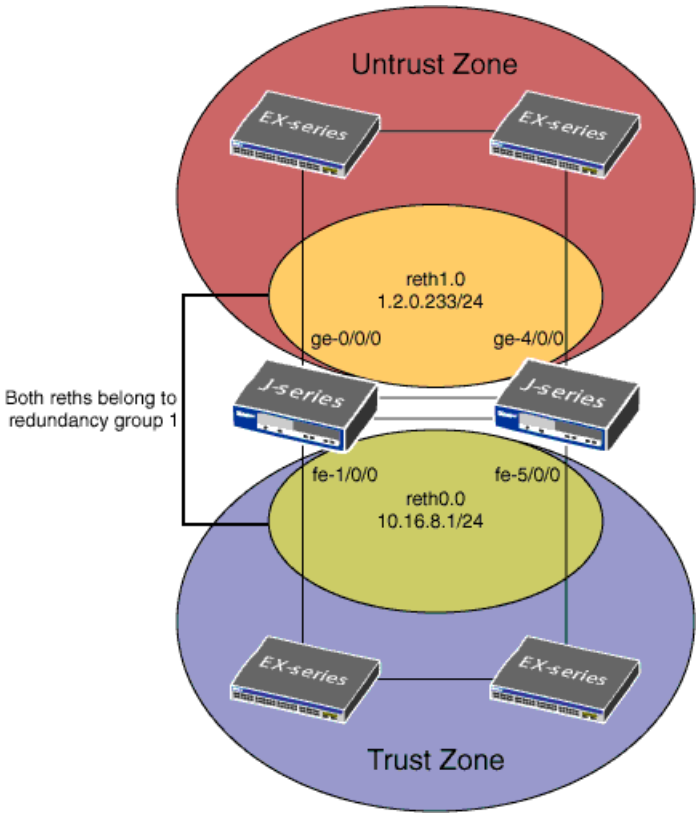

An active/passive chassis cluster can be achieved by using redundant Ethernet interfaces (*reths*) that are all assigned to the same redundancy group. If any of the interfaces in an active group in a node fails, the group is declared inactive and all the interfaces in the group will fail over to the other node.

This configuration minimizes the traffic over the fabric link because only one node in the cluster will forward traffic at any given time.

CLI Configuration

1. Enabling clustering

   In node 0:

   ```
   user@host> set chassis cluster node 0 cluster-id 1
   warning: A reboot is required for chassis cluster to be enabled
   ```

   In node 1:

   ```
   user@host> set chassis cluster node 1 cluster-id 1 reboot
   Successfully enabled chassis cluster. Going to reboot now.
   ```

2. Management interface

In a cluster, the configuration is shared among the cluster members. Member-specific configurations (such as the IP address of the management port of each member) are entered using configuration groups.

```
{primary:node1}
user@host# set groups node0 system host-name J2320-A
{primary:node1}
user@host# set groups node0 interfaces fxp0 unit 0 family inet address
  192.168.3.110/24
{primary:node1}
user@host#set groups node1 system host-name J2320-B
{primary:node1}
user@host# set groups node1 interfaces fxp0 unit 0 family inet address
  192.168.3.111/24
{primary:node1}
user@host# set apply-groups "${node}"
```

3. Fabric interface

```
{primary:node1}
user@host# set interfaces fab0 fabric-options member-interfaces ge-0/0/1
{primary:node1}
user@host# set interfaces fab1 fabric-options member-interfaces ge-4/0/1
```

4. Redundancy groups

```
{primary:node1}
user@host# set chassis cluster reth-count 2
{primary:node1}
user@host# set chassis cluster heartbeat-interval 1000
{primary:node1}
user@host# set chassis cluster heartbeat-threshold 3
{primary:node1}
user@host# set chassis cluster node 0
{primary:node1}
user@host# set chassis cluster node 1
{primary:node1}
user@host# set chassis cluster redundancy-group 0 node 0 priority 100
{primary:node1}
user@host# set chassis cluster redundancy-group 0 node 1 priority 1
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 0 priority 100
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 1 priority 1
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor fe-1/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor fe-5/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor ge-0/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor ge-4/0/0
  weight 255
```

5. Redundant Ethernet interfaces

```
{primary:node1}
user@host# set interfaces ge-0/0/0 gigether-options redundant-parent reth1
{primary:node1}
user@host# set interfaces ge-4/0/0 gigether-options redundant-parent reth1
{primary:node1}
user@host# set interfaces fe-1/0/0 fastether-options redundant-parent reth0
{primary:node1}
user@host# set interfaces fe-5/0/0 fastether-options redundant-parent reth0
{primary:node1}
user@host# set interfaces reth0 redundant-ether-options redundancy-group 1
{primary:node1}
user@host# set interfaces reth0 unit 0 family inet address 10.16.8.1/24
{primary:node1}
user@host# set interfaces reth1 redundant-ether-options redundancy-group 1
{primary:node1}
user@host# set interfaces reth1 unit 0 family inet address 1.2.0.233/24
```

6. Security zones

```
{primary:node1}
user@host# set security zones security-zone Untrust interfaces reth1.0
{primary:node1}
user@host# set security zones security-zone Trust interfaces reth0.0
```

7. Security policies

```
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  match source-address any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  match destination-address any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  match application any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  then permit
```

J-Web Configuration

1. Enabling clustering

   See Step 1 in "CLI Configuration" on page 943.

2. Management interface

   See Step 2 in "CLI Configuration" on page 943.

3. Fabric interface

   See Step 3 in "CLI Configuration" on page 943.

4. Redundancy groups

- Select **Configure > System Properties > Chassis Cluster**.
- Enter the following, then click **Apply**:

  Redundant ether-Interface Count: **2**

  Heartbeat Interval: **1000**

  Heartbeat Threshold: **3**

  Nodes: **0**

  Group Number: **0**

  Priorities: **100**

- Enter the following, then click **Apply**:

  Nodes: **0**

  Group Number: **1**

  Priorities: **1**

- Enter the following, then click **Apply**:

  Nodes: **1**

  Group Number: **0**

  Priorities: **100**

- Enter the following, then click **Apply**:

  Nodes: **1**

  Group Number: **1**

  Priorities: **1**

  Interface Monitor—Interface: **fe-1/0/0**

  Interface Monitor—Weight: **255**

  Interface Monitor—Interface: **fe-5/0/0**

  Interface Monitor—Weight: **255**

  Interface Monitor—Interface: **ge-0/0/0**

  Interface Monitor—Weight: **255**

  Interface Monitor—Interface: **ge-4/0/0**

  Interface Monitor—Weight: **255**

5. Redundant Ethernet interfaces
   - Select **Configure > System Properties > Chassis Cluster**.
   - Select **ge-0/0/0**.

- Enter the following:

  Redundant Parent: **reth1**

- Click **Apply**.
- Select **ge-4/0/0**.
- Enter the following:

  Redundant Parent: **reth1**

- Click **Apply**.
- Select **fe-1/0/0**.
- Enter the following:

  Redundant Parent: **reth0**

- Click **Apply**.
- Select **fe-5/0/0**.
- Enter the following:

  Redundant Parent: **reth0**

- Click **Apply**.
- For last four configuration settings of Step 5, see Step 5 in "CLI Configuration" on page 943.

6. Security zones

   See Step 6 in "CLI Configuration" on page 943.

7. Security policies

   See Step 7 in "CLI Configuration" on page 943.

Active/Passive Chassis Cluster with IPsec Tunnel Scenario

In this case, a single device in the cluster terminates in an IPsec tunnel and is used to process all traffic while the other device is used only in the event of a failure (see Figure 177 on page 948). When a failure occurs, the backup device becomes master and controls all forwarding.

Figure 177: Active/Passive Chassis Cluster with IPsec Tunnel Scenario (SRX Series Devices)

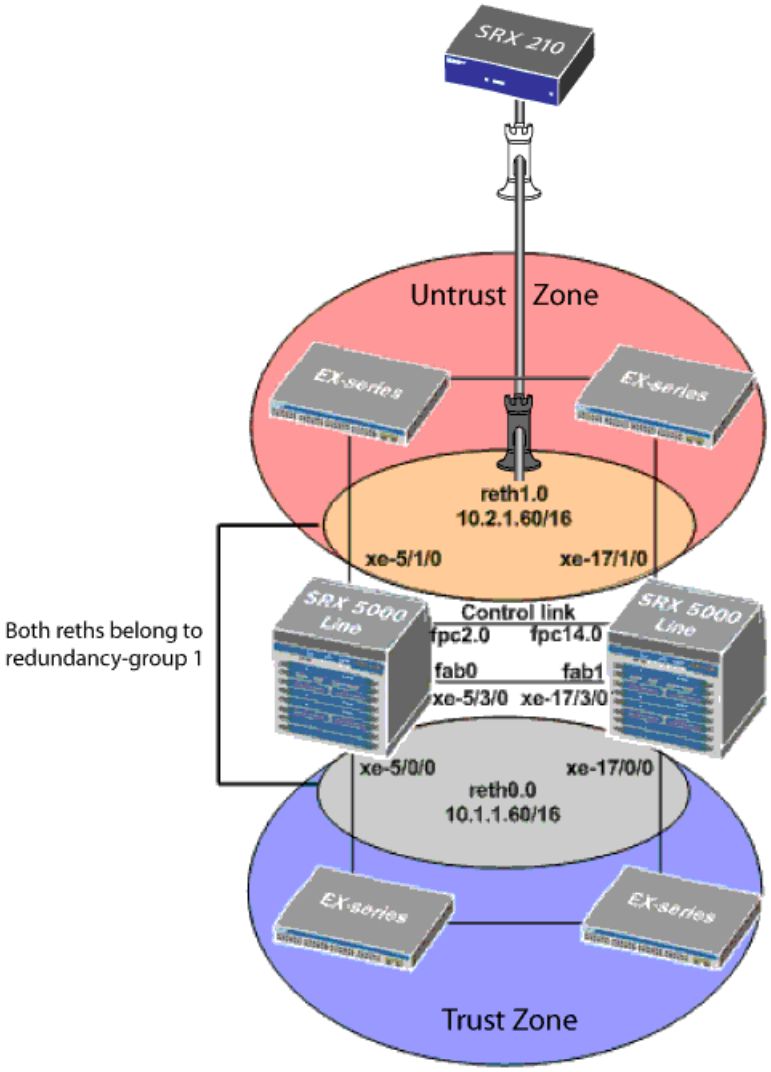

An active/passive chassis cluster can be achieved by using redundant Ethernet interfaces (*reths*) that are all assigned to the same redundancy group. If any of the interfaces in an active group in a node fails, the group is declared inactive and all the interfaces in the group will fail over to the other node.

This configuration provides a way for a site-to-site IPsec tunnel to terminate in an active/passive cluster where a redundant Ethernet interface is used as the tunnel endpoint. In the event of a failure, the redundant Ethernet interface in the backup SRX Series device will become active, forcing the tunnel to change endpoints to terminate in the new active SRX Series device. Because tunnel keys and session information are synchronized between the members of the chassis cluster, a failover will not require the tunnel to be renegotiated and all established sessions will be maintained.

CLI Configuration

1. Enabling clustering

   In node 0:

   ```
   user@host> set chassis cluster node 0 cluster-id 1
   warning: A reboot is required for chassis cluster to be enabled
   ```

   In node 1:

   ```
   user@host> set chassis cluster node 1 cluster-id 1 reboot
   Successfully enabled chassis cluster. Going to reboot now.
   {primary:node1}
   user@host# set chassis cluster control-ports fpc 2 port 0
   {primary:node1}
   user@host# set chassis cluster control-ports fpc 14 port 0
   ```

2. Management interface

   In a cluster, the configuration is shared among the cluster members. Member-specific configurations (such as the IP address of the management port of each member) are entered using configuration groups.

   ```
   {primary:node1}
   user@host# set groups node0 system host-name SRX5800-1
   {primary:node1}
   user@host# set groups node0 interfaces fxp0 unit 0 family inet address
     172.19.100.50/24
   {primary:node1}
   user@host#set groups node1 system host-name SRX5800-2
   {primary:node1}
   user@host# set groups node1 interfaces fxp0 unit 0 family inet address
     172.19.100.51/24
   {primary:node1}
   user@host# set apply-groups "${node}"
   ```

3. Fabric interface

   ```
   {primary:node1}
   user@host# set interfaces fab0 fabric-options member-interfaces xe-5/3/0
   {primary:node1}
   user@host# set interfaces fab1 fabric-options member-interfaces xe-17/3/0
   ```

4. Redundancy groups

   ```
   {primary:node1}
   user@host# set chassis cluster reth-count 2
   {primary:node1}
   user@host# set chassis cluster heartbeat-interval 1000
   {primary:node1}
   user@host# set chassis cluster heartbeat-threshold 3
   {primary:node1}
   user@host# set chassis cluster node 0
   ```

```
{primary:node1}
user@host# set chassis cluster node 1
{primary:node1}
user@host# set chassis cluster redundancy-group 0 node 0 priority 254
{primary:node1}
user@host# set chassis cluster redundancy-group 0 node 1 priority 1
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 0 priority 254
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 1 priority 1
{primary:node1}
user@host# set chassis cluster redundancy-group 1 preempt
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor xe-5/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor xe-5/1/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor xe-17/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor xe-17/1/0
  weight 255
```

5. Redundant Ethernet interfaces

```
{primary:node1}
user@host# set interfaces xe-5/1/0 gigether-options redundant-parent reth1
{primary:node1}
user@host# set interfaces xe-17/1/0 gigether-options redundant-parent reth1
{primary:node1}
user@host# set interfaces xe-5/0/0 gigether-options redundant-parent reth0
{primary:node1}
user@host# set interfaces xe-17/0/0 gigether-options redundant-parent reth0
{primary:node1}
user@host# set interfaces reth0 redundant-ether-options redundancy-group 1
{primary:node1}
user@host# set interfaces reth0 unit 0 family inet address 10.1.1.60/16
{primary:node1}
user@host# set interfaces reth1 redundant-ether-options redundancy-group 1
{primary:node1}
user@host# set interfaces reth1 unit 0 family inet address 10.2.1.60/16
```

6. IPsec configuration

```
{primary:node1}
user@host# set interfaces st0 unit 0 multipoint family inet address 10.10.1.1/30
{primary:node1}
user@host# set interfaces st0 family inet address 10.10.1.1/30
{primary:node1}
user@host# set security ike policy preShared mode main
{primary:node1}
user@host# set security ike policy preShared proposal-set standard
{primary:node1}
```

```
user@host# set security ike policy preShared pre-shared-key ascii-text
  "juniper"## Encrypted password
{primary:node1}
user@host# set security ike gateway SRX210-1 ike-policy preShared
{primary:node1}
user@host# set security ike gateway SRX210-1 address 10.1.1.90
{primary:node1}
user@host# set security ike gateway SRX210-1 external-interface reth0.0
{primary:node1}
user@host# set security ipsec policy std proposal-set standard
{primary:node1}
user@host# set security ipsec vpn SRX210-1 bind-interface st0.0
{primary:node1}
user@host# set security ipsec vpn SRX210-1 vpn-monitor optimized
{primary:node1}
user@host# set security ipsec vpn SRX210-1 ike gateway SRX210-1
{primary:node1}
user@host# set security ipsec vpn SRX210-1 ike ipsec-policy std
{primary:node1}
user@host# set security ipsec vpn SRX210-1 establish-tunnels immediately
```

7. Static routes

```
{primary:node1}
user@host# set routing-options static route 0.0.0.0/0 next-hop 10.2.1.1
{primary:node1}
user@host# set routing-options static route 10.3.0.0/16 next-hop 10.10.1.2
```

8. Security zones

```
{primary:node1}
user@host# set security zones security-zone Untrust host-inbound-traffic
  system-services all
{primary:node1}
user@host# set security zones security-zone Untrust host-inbound-traffic
  protocols all
{primary:node1}
user@host# set security zones security-zone Untrust interfaces reth1.0
{primary:node1}
user@host# set security zones security-zone Trust host-inbound-traffic
  system-services all
{primary:node1}
user@host# set security zones security-zone Trust host-inbound-traffic protocols
  all
{primary:node1}
user@host# set security zones security-zone Trust interfaces reth0.0
{primary:node1}
user@host# set security zones security-zone vpn host-inbound-traffic
  system-services all
{primary:node1}
user@host# set security zones security-zone vpn host-inbound-traffic protocols
  all
{primary:node1}
user@host# set security zones security-zone vpn interfaces st0.0
```

9. Security policies

```
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
   match source-address any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
   match destination-address any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
   match application any
{primary:node1}
user@host# set security policies from-zone Trust to-zone vpn policy ANY then
   permit
```

J-Web Configuration

1. Enabling clusters

   See Step 1 in "CLI Configuration" on page 949.

2. Management interface

   See Step 2 in "CLI Configuration" on page 949.

3. Fabric interface

   See Step 3 in "CLI Configuration" on page 949.

4. Redundancy groups
   - Select **Configure > System Properties > Chassis Cluster**.
   - Enter the following, then click **Apply**:

     Redundant ether-Interfaces Count: **2**

     Heartbeat Interval: **1000**

     Heartbeat Threshold: **3**

     Nodes: **0**

     Group Number: **0**

     Priorities: **254**

   - Enter the following, then click **Apply**:

     Nodes: **0**

     Group Number: **1**

     Priorities: **254**

   - Enter the following, then click **Apply**:

     Nodes: **1**

     Group Number: **0**

Priorities: **1**

- Enter the following, then click **Apply**:

  Nodes: **1**

  Group Number: **1**

  Priorities: **1**

  Preempt: Select the check box.

  Interface Monitor—Interface: **xe-5/0/0**

  Interface Monitor—Weight: **255**

  Interface Monitor—Interface: **xe-5/1/0**

  Interface Monitor—Weight: **255**

  Interface Monitor—Interface: **xe-17/0/0**

  Interface Monitor—Weight: **255**

  Interface Monitor—Interface: **xe-17/1/0**

  Interface Monitor—Weight: **255**

5. Redundant Ethernet interfaces
   - Select **Configure > System Properties > Chassis Cluster**.
   - Select **xe-5/1/0**.
   - Enter the following:

     Redundant Parent: **reth1**

   - Click **Apply**.
   - Select **xe-17/1/0**.
   - Enter the following:

     Redundant Parent: **reth1**

   - Click **Apply**.
   - Select **xe-5/0/0**.
   - Enter the following:

     Redundant Parent: **reth0**

   - Click **Apply**.
   - Select **xe-17/0/0**.
   - Enter the following:

Redundant Parent: **reth0**

- Click **Apply**.
- For last four configuration settings of Step 5, see Step 5 in "CLI Configuration" on page 949.

6. IPsec configuration

   See Step 6 in "CLI Configuration" on page 949.

7. Static routes
   - Select **Configure > Routing > Static Routing**.
   - Click **Add**.
   - Enter the following, then click **Apply**:

     Static Route Address: **0.0.0.0/0**

     Next-Hop Addresses: **10.2.1.1**

   - Enter the following, then click **Apply**:

     Static Route Address: **10.3.0.0/16**

     Next-Hop Addresses: **10.10.1.2**

8. Security zones

   See Step 8 in "CLI Configuration" on page 949.

9. Security policies

   See Step 9 in "CLI Configuration" on page 949.

Asymmetric Routing Chassis Cluster Scenario

In this case, chassis cluster makes use of its asymmetric routing capability (see Figure 178 on page 955). Traffic received by a node is matched against that node's session table. The result of this lookup determines whether or not that node should process the packet or forward it to the other node over the fabric link. Sessions are anchored on the egress node for the first packet that created the session. If traffic is received on the node in which the session is not anchored, those packets are forwarded over the fabric link to the node where the session is anchored.

NOTE: The anchor node for the session can change if there are changes in routing during the session.

Figure 178: Asymmetric Routing Chassis Cluster Scenario (J Series Devices)

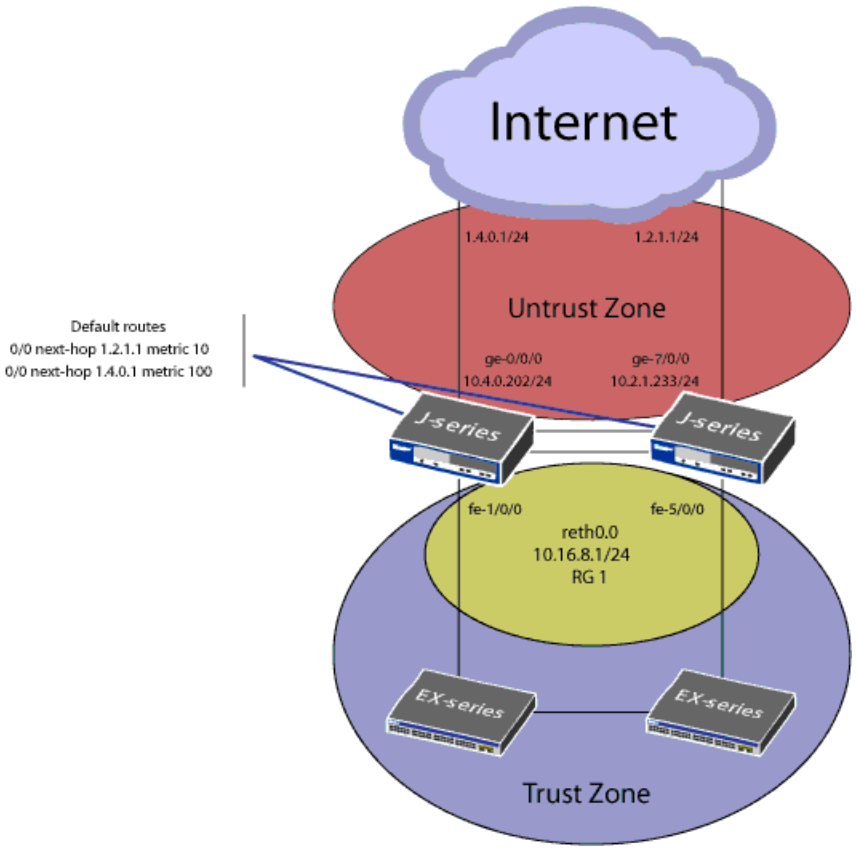

In this scenario, two Internet connections are used, with one being preferred. The connection to the Trust zone is done by using a redundant Ethernet interface to provide LAN redundancy for the devices in the Trust zone. This scenario describes two failover cases in which sessions originate in the Trust zone with a destination of the Internet (Untrust zone).

Case 1: Failures in the Trust Zone reth

Under normal operating conditions, traffic flows from the Trust zone interface fe-1/0/0, belonging to reth0.0, to the Internet. Because the primary Internet connection is on node 0, sessions are both created in node 0 and synced to node 1. However, session are only active on node 0.

A failure in interface fe-1/0/0 triggers a failover of the redundancy group, causing interface fe-5/0/0 in node 1 to become active. After the failover, traffic arrives at node 1. After session lookup, the traffic is sent to node 0 because the session is active on this node. Node 0 then processes the traffic and forwards it to the Internet. The return traffic follows a similar process. The traffic arrives at node 0 and gets processed for security purposes—for example., antispam scanning, antivirus scanning, and application of security policies—on node 0 because the session is anchored to node 0. The packet is then sent to node 1 via the fabric interface for egress processing and eventual transmission out of node 1 via interface fe-5/0/0.

Case 2: Failures in the Untrust Zone Interfaces

In this case, sessions are migrated from node to node. Under normal operating conditions, traffic is processed by only node 0. A failure of interface ge-0/0/0 on node 0 causes a change in the routing table, so that it now points to interface ge-7/0/0 in node 1. After the failure, sessions in node 0 become inactive, and the passive sessions in node 1 become active. Traffic arriving from the Trust zone is still received on interface fe-1/0/0, but is forwarded to node 1 for processing. After traffic is processed in node 1, it is forwarded to the Internet through interface ge-7/0/0.

In this chassis cluster configuration, redundancy group 1 is used to control the redundant Ethernet interface connected to the Trust zone. As configured in this scenario, redundancy group 1 fails over only if interface fe-1/0/0 or fe-5/0/0 fails, but not if the interfaces connected to the Internet fail. Optionally, the configuration could be modified to permit redundancy group 1 to monitor all interfaces connected to the Internet and fail over if an Internet link were to fail. So, for example, the configuration can allow redundancy group 1 to monitor ge-0/0/0 and make fe-5/0/0 active for reth0 if the ge-0/0/0 Internet link fails. (This option is not described in the configuration examples below.)

CLI Configuration

NOTE: First, do basic chassis cluster and management interfaces setup.

1. Fabric interface

```
{primary:node1}
user@host# set interfaces fab0 fabric-options member-interfaces ge-0/0/1
{primary:node1}
user@host# set interfaces fab1 fabric-options member-interfaces ge-7/0/1
```

2. Redundancy groups

```
{primary:node1}
user@host# set chassis cluster reth-count 1
user@host# set chassis cluster heartbeat-interval 1000
{primary:node1}
user@host# set chassis cluster heartbeat-threshold 3
{primary:node1}
user@host# set chassis cluster node 0
{primary:node1}
user@host# set chassis cluster node 1
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 0 priority 100
{primary:node1}
user@host# set chassis cluster redundancy-group 1 node 1 priority 1
{primary:node1}
user@host# set chassis cluster redundancy-group 1 interface-monitor fe-1/0/0
  weight 255
{primary:node1}
```

```
user@host# set chassis cluster redundancy-group 1 interface-monitor fe-5/0/0
  weight 255
```

3. Redundant Ethernet interfaces

```
{primary:node1}
user@host# set interfaces ge-0/0/0 unit 0 family inet address 1.4.0.202/24
{primary:node1}
user@host# set interfaces fe-1/0/0 fastether-options redundant-parent reth0
{primary:node1}
user@host# set interfaces ge-7/0/0 unit 0 family inet address 1.2.1.233/24
{primary:node1}
user@host# set interfaces fe-5/0/0 fastether-options redundant-parent reth0
{primary:node1}
user@host# set interfaces reth0 unit 0 family inet address 10.16.8.1/24
```

4. Static routes (one to each ISP, with preferred route through ge-0/0/0)

```
{primary:node1}
user@host# set routing-options static route 0.0.0.0/0 qualified-next-hop 1.4.0.1
  metric 10
{primary:node1}
user@host# set routing-options static route 0.0.0.0/0 qualified-next-hop 1.2.1.1
  metric 100
```

5. Security zones

```
{primary:node1}
user@host# set security zones security-zone Untrust interfaces ge-0/0/0.0
{primary:node1}
user@host# set security zones security-zone Untrust interfaces ge-7/0/0.0
```

6. Security policies

```
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  match source-address any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  match destination-address any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  match application any
{primary:node1}
user@host# set security policies from-zone Trust to-zone Untrust policy ANY
  then permit
```

J-Web Configuration

NOTE: First, enable clustering and do management interfaces setup.

1. Fabric interface

   See Step 1 in "CLI Configuration" on page 956.

2. Redundancy groups
   - Select **Configure > System Properties > Chassis Cluster**.
   - Enter the following, then click **Apply**:

     Redundant ether-Interface Count: **1**

     Heartbeat Interval: **1000**

     Heartbeat Threshold: **3**

     Nodes: **0**

     Group Number: **1**

     Priorities: **100**

   - Enter the following, then click **Apply**:

     Nodes: **1**

     Group Number: **1**

     Priorities: **1**

     Interface Monitor—Interface: **fe-1/0/0**

     Interface Monitor—Weight: **255**

     Interface Monitor—Interface: **fe-5/0/0**

     Interface Monitor—Weight: **255**

3. Redundant Ethernet interfaces
   - Select **Configure > Interfaces**.
   - Select **fe-1/0/0**.
   - Enter the following:

     Redundant Parent: **reth0**

   - Click **Apply**.
   - Select **fe-5/0/0**.
   - Enter the following:

Redundant Parent: **reth0**

- Click **Apply**.
- For other configuration settings in Step 3, see Step 3 in "CLI Configuration" on page 943.

4. Static routes (one to each ISP, with preferred route through **ge-0/0/0**)
   - Select **Configure > Routing > Static Routing**.
   - Click **Add**.
   - Enter the following, then click **Apply**:

     Static Route Address: **0.0.0.0/0**

     Next-Hop Addresses: **1.4.0.1**, **1.2.1.1**

5. Security zones

   See Step 5 in "CLI Configuration" on page 956.

6. Security policies

   See Step 6 in "CLI Configuration" on page 956.

Part 11
Network Address Translation

- Network Address Translation on page 963

Chapter 29

Network Address Translation

Network Address Translation (NAT) is a method by which IP addresses in a packet are mapped from one group to another and, optionally, port numbers in the packet are translated into different port numbers.

NAT is described in RFC 1631 to solve IP (version 4) address depletion problems. Since then, NAT has been found to be a useful tool for firewalls, traffic redirect, load sharing, network migrations, and so on.

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This chapter includes the following topics:

- NAT Overview on page 964
- Understanding NAT Rule Sets and Rules on page 965
- Understanding Static NAT on page 967
- Static NAT Configuration Overview on page 968
- Example: Configuring Static NAT on page 969
- Understanding Destination NAT on page 970
- Destination NAT Configuration Overview on page 970
- Example: Configuring Destination NAT on page 972
- Configuring Destination NAT—Quick Configuration on page 973
- Understanding Source NAT on page 974
- Source NAT Configuration Overview on page 975
- Example: Configuring Source NAT Pools on page 977
- Understanding NAT Source Pools with PAT on page 978
- Disabling Port Randomization for Source NAT on page 979
- Understanding NAT Source Pools Without PAT on page 979
- Example: Configuring Source NAT on page 980
- Configuring Source NAT—Quick Configuration on page 982
- Example: Configuring a Persistent Address on page 984
- Understanding Persistent NAT on page 985
- Understanding Session Traversal Utilities for NAT (STUN) Protocol on page 986

- Persistent NAT Configuration Overview on page 987
- Example: Configuring Persistent NAT with Source NAT Address Pool on page 988
- Example: Configuring Persistent NAT with Interface NAT on page 989
- Configuring Proxy ARP on page 991
- Verifying NAT Configuration on page 991

NAT Overview

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses.

This section describes the types of NAT supported on Juniper Networks devices.

- Source NAT on page 964
- Destination NAT on page 965
- Static NAT on page 965
- Related Topics on page 965

Source NAT

Source NAT allows connections to be initiated only for outgoing network connections—for example, from a private network to the Internet. Source NAT is commonly used to perform the following translations:

- Translate a single IP address to another address (for example, to provide a single device in a private network with access to the Internet).
- Translate a contiguous block of addresses to another block of addresses of the same size.
- Translate a contiguous block of addresses to another block of addresses of smaller size.
- Translate a contiguous block of addresses to a single IP address or a smaller block of addresses using port translation.
- Translate a contiguous block of addresses to the address of the egress interface.

Translation to the address of the egress interface does not require an address pool; all other source NAT translations require configuration of an address pool. One-to-one and many-to-many translations for address blocks of the same size do not require port translation because there is an available address in the pool for every address that would be translated.

If the size of the address pool is smaller than the number of addresses that would be translated, either the total number of concurrent addresses that can be translated is limited by the size of the address pool or port translation must be used. For example, if a block of 253 addresses is translated to an address pool of 10 addresses,

a maximum of 10 devices can be connected concurrently unless port translation is used.

Destination NAT

Destination NAT allows connections to be initiated only for incoming network connections—for example, from the Internet to a private network. Destination NAT is commonly used to perform the following actions:

- Translate a single IP address to another address (for example, to allow a device on the Internet to connect to a host on a private network).
- Translate a contiguous block of addresses to another block of addresses of the same size (for example, to allow access to a group of servers).
- Translate a destination IP address and port to another destination IP address and port (for example, to allow access to multiple services using the same IP address but different ports).

Static NAT

Static NAT allows connections to be originated from either side of the network, but translation is limited to one-to-one or between blocks of addresses of the same size. For each private address, a public address must be allocated. No address pools are necessary.

Related Topics

- Understanding NAT Rule Sets and Rules on page 965
- Understanding Static NAT on page 967
- Understanding Destination NAT on page 970
- Understanding Source NAT on page 974

Understanding NAT Rule Sets and Rules

NAT processing centers on the evaluation of NAT rule sets and rules. A rule set determines the overall direction of the traffic to be processed. For example, a rule set can select traffic from a particular interface or to a specific zone. A rule set can contain multiple rules. Once a rule set is found that matches specific traffic, each rule in the rule set is evaluated for a match. Each rule in the rule set further specifies the traffic to be matched and the action to be taken when traffic matches the rule.

NAT Rule Sets

A rule set specifies a general set of matching conditions for traffic. For static NAT and destination NAT, a rule set specifies one of the following:

- Source interface
- Source zone

- Source routing instance

For source NAT rule sets, you configure both source and destination conditions:

- Source interface, zone, or routing instance
- Destination interface, zone, or routing instance

It is possible for a packet to match more than one rule set; in this case, the rule set with the more specific match is used. An interface match is considered more specific than a zone match, which is more specific than a routing instance match. If a packet matches both a destination NAT rule set that specifies a source zone and a destination NAT rule set that specifies a source interface, the rule set that specifies the source interface is the more specific match.

Source NAT rule set matching is more complex because you specify both source and destination conditions in a source NAT rule set. In the case where a packet matches more than one source NAT rule set, the rule set chosen is based on the following source/destination conditions (in order of priority):

1. Source interface/destination interface
2. Source zone/destination interface
3. Source routing instance/destination interface
4. Source interface/destination zone
5. Source zone/destination zone
6. Source routing instance/destination zone
7. Source interface/destination routing instance
8. Source zone/destination routing instance
9. Source routing instance/destination routing instance

For example, you can configure rule set A, which specifies a source interface and a destination zone, and rule set B, which specifies a source zone and a destination interface. If a packet matches both rule sets, rule set B is the more specific match.

NOTE: You cannot specify the same source and destination conditions for source NAT rule sets.

NAT Rules

Once a rule set that matches the traffic has been found, each rule in the rule set is evaluated in order for a match. NAT rules can match on the following packet information:

- Destination address (for static NAT only)
- Source and destination address (for destination and source NAT)
- Destination port (for destination and source NAT)

The first rule in the rule set that matches the traffic is used. If a packet matches a rule in a rule set during session establishment, traffic is processed according to the action specified by that rule.

Rule Processing

The NAT type determines the order in which NAT rules are processed. During the first packet processing for a flow, NAT rules are applied in the following order:

1. Static NAT rules
2. Destination NAT rules
3. Route lookup
4. Security policy lookup
5. Reverse mapping of static NAT rules
6. Source NAT rules

Static NAT and destination NAT rules are processed before route and security policy lookup. Static NAT rules take precedence over destination NAT rules. Reverse mapping of static NAT rules takes place after route and security policy lookup and takes precedence over source NAT rules. Source NAT rules are processed after route and security policy lookup and after reverse mapping of static NAT rules.

The configuration of rules and rule sets is basically the same for each type of NAT—source, destination, or static. But because both destination and static NAT are processed before route lookup, you cannot specify the destination zone, interface or routing instance in the rule set.

Related Topics

- Static NAT Configuration Overview on page 968
- Destination NAT Configuration Overview on page 970
- Source NAT Configuration Overview on page 975

Understanding Static NAT

Static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes destination IP address translation in one direction and source IP address translation in the reverse direction. From the NAT device, the original destination address is the virtual host IP address while the mapped-to address is the real host IP address.

NOTE: The original destination address, along with other addresses in source and destination NAT pools, must not overlap within the same routing instance.

Static NAT does not perform port address translation (PAT) and no address pools are needed for static NAT.

In NAT rule lookup, static NAT rules take precedence over destination NAT rules and reverse mapping of static NAT rules take precedence over source NAT rules.

Related Topics

- Static NAT Configuration Overview on page 968

Static NAT Configuration Overview

The main configuration tasks for static NAT are as follows:

- Configure static NAT rules that align with your network and security requirements.
- Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Static NAT Rules

Static NAT rules specify two layers of match conditions:

- Traffic direction—Allows you to specify from interface, from zone, or from routing-instance.
- Packet information—Destination IP address.

If multiple static NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 and rule B specifies traffic from interface ge-0/0/0, rule B is used to perform static NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

For the static NAT rule action, specify the translated address and (optionally) the routing instance.

In NAT lookup, static NAT rules take precedence over destination NAT rules and reverse mapping of static NAT rules takes precedence over source NAT rules.

Related Topic

- Example: Configuring Static NAT on page 969

Example: Configuring Static NAT

JUNOS Software allows you to define static NAT rules to perform one-to-one static mapping from one IP address or subnet to another IP address or subnet without port address translation.

Before You Begin

For background information, read:

- Understanding Static NAT on page 967

In this example, you perform the following tasks:

- Define a static NAT rule for traffic from zone red to address 20.1.1.100. Translate the destination IP to 10.0.0.200 and routing-instance ri-green.
- Define a static NAT rule for traffic from zone red to address 20.1.1.101. Translate the destination IP to 10.0.1.200 and routing-instance ri-blue.
- Define a static NAT rule for traffic from interface ge-0/0/0.0 to address 30.1.1.3. Translate the destination IP to 10.0.2.200.
- Define a static NAT rule for traffic from routing-instance ri-red to address subnet 30.1.1.0/24. Translate the destination address subnet to 172.16.0.0/24.

CLI Configuration

```
user@host# set security nat static rule-set rs3 from zone red
user@host# set security nat static rule-set rs3 rule r1 match destination-address
  20.1.1.100
user@host# set security nat static rule-set rs3 rule r1 then static-nat prefix
  10.0.0.200
user@host# set security nat static rule-set rs3 rule r1 then static-nat prefix
  10.0.0.200 routing-instance ri-green

user@host# set security nat static rule-set rs3 from zone red
user@host# set security nat static rule-set rs3 rule r2 match destination-address
  20.1.1.101
user@host# set security nat static rule-set rs3 rule r2 then static-nat prefix
  10.0.1.200
user@host# set security nat static rule-set rs3 rule r2 then static-nat prefix
  10.0.1.200 routing-instance ri-blue

user@host#set security nat static rule-set rs3 from interface ge-0/0/0.0
```

```
user@host#set security nat static rule-set rs3 rule r3 match destination address
    30.1.1.3
user@host#set security nat static rule-set rs3 rule r3 then static-nat prefix
    10.0.2.200

user@host#set security nat static rule-set rs4 from routing-instance ri-red
user@host#set security nat static rule-set rs4 rule r3 match destination-address
    30.1.1.0/24
user@host#set security nat static rule-set rs4 rule r3 then static-nat prefix
    172.16.0.0/24
```

Understanding Destination NAT

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

NOTE: When destination NAT is performed, the destination IP address is translated according to configured destination NAT rules and then security policies are applied.

The following types of destination NAT are supported:

- Translation of the original destination IP address to an IP address from a user-defined pool. This type of translation does not include Port Address Translation (PAT). If the original destination IP address range is larger than the address range in the user-defined address pool, any untranslated packets are dropped.
- Translation of the original destination IP address (and optional port number) to one specific IP address (and port number) from a user-defined pool.

Related Topics

- Destination NAT Configuration Overview on page 970

Destination NAT Configuration Overview

The main configuration tasks for destination NAT are as follows:

- Configure a destination NAT address pool that aligns with your network and security requirements.
- Configure destination NAT rules that align with your network and security requirements.
- Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Destination NAT Address Pool

For destination NAT address pools, specify the following:

- Name of the destination NAT address pool
- Destination address or address range

NOTE: Do not overlap NAT addresses for source NAT, destination NAT, and static NAT within one routing instance.

- Destination port that is used for port forwarding
- Routing instance to which the pool belongs (the default is the main inet.0 routing instance)

Destination NAT Rules

Destination NAT rules specify two layers of match conditions:

- Traffic direction—Allows you to specify from interface, from zone, or from routing-instance.
- Packet information—Can be source IP addresses, destination IP address or subnet, or a single destination port number.

If multiple destination NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 and rule B specifies traffic from interface ge-0/0/0, rule B is used to perform destination NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

The actions you can specify for a destination NAT rule are:

- off—Do not perform destination NAT.
- pool—Use the specified user-defined address pool to perform destination NAT.

Destination NAT rules are applied to traffic in the first packet that is processed for the flow or in the fast path for the ALG. Destination NAT rules are processed after static NAT rules but before source NAT rules.

Related Topic

- Example: Configuring Destination NAT on page 972

Example: Configuring Destination NAT

JUNOS Software allows you to define NAT rules to translate the destination address from one IP address to another.

Before You Begin
For background information, read: ■ Destination NAT Configuration Overview on page 970

In this example, you perform the following tasks:

- Define a destination NAT rule for traffic from routing-instance ri-1 to address 1.1.1.1. Translate the destination IP address to 2.2.2.10.
- Define a destination NAT rule for traffic from zone Z2 to address 1.1.1.1. Translate the destination IP address to 2.2.2.20.
- Define a destination NAT rule for traffic from interface fe-0/0/0.0 to address 1.1.1.1. Translate the destination IP address to 2.2.2.30.

CLI Configuration

```
user@host# set security nat destination pool dpool-1 routing-instance ri-1
user@host# set security nat destination pool dpool-1 address 2.2.2.10
user@host# set security nat destination pool dpool-2 routing-instance ri-1
user@host# set security nat destination pool dpool-2 address 2.2.2.20
user@host# set security nat destination pool dpool-3 routing-instance ri-1
user@host# set security nat destination pool dpool-3 address 2.2.2.30

user@host# set security nat destination rule-set rs2 from routing-instance ri-1
user@host# set security nat destination rule-set rs2 rule r1 match destination-address
    1.1.1.1
user@host# set security nat destination rule-set rs2 rule r1 then destination-nat
    pool dpool-1

user@host# set security nat destination rule-set rs3 from zone Z2
user@host# set security nat destination rule-set rs3 rule r2 match destination-address
    1.1.1.1
user@host# set security nat destination rule-set rs3 rule r2 then destination-nat
    pool dpool–2

user@host# set security nat destination rule-set rs4 from interface fe-0/0/0.0
user@host# set security nat destination rule-set rs4 rule r3 match destination-address
    1.1.1.1
```

```
user@host# set security nat destination rule-set rs4 rule r3 then destination-nat
    pool dpool-3
```

Configuring Destination NAT—Quick Configuration

You can use J-Web Quick Configuration to quickly configure destination NAT.

Before You Begin
For background information, read: ■ Destination NAT Configuration Overview on page 970

To configure destination NAT with Quick Configuration:

1. Select **Configure > NAT > Destination NAT**.
2. Fill in the options as shown in Table 115 on page 973.
3. Click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.

Table 115: Destination NAT Configuration Options

Field	Function	Action
Destination NAT		
Destination Address Pool	Defines one or more destination NAT pools.	Click **Add.**
Destination Rule Sets	Defines one or more rule sets.	Click **Add.**
Destination NAT Pool		
Name	Descriptive name for the pool.	Type a name for the pool.
Routing Instance	Specifies the routing instance in which to perform route lookup for the address in the pool. By default, the routing instance is inet.0.	Enter the routing instance name.
IP Address Range	Specifies an address range.	Enter the IP addresses.
IP Address	Specifies a single IP address.	Enter an IP address.
Port	Specifies the optional port number.	Enter a port number.
Destination NAT Rule Set		

Table 115: Destination NAT Configuration Options *(continued)*

Field	Function	Action
Name	Name of the rule set.	Type a name for the rule set.
From	Specifies where the packet is coming from. If a packet matches more than one rule set, the rule set with the more specific match is used. An interface match is considered more specific than a zone match, which is more specific than a routing instance match.	Routing Instance—Specify the source routing instance. Interface—Specify the source interface. Zone—Specify the source zone.
Rules	Set of conditions that are evaluated, in order, for a match to perform the specified destination NAT action.	
Rule name	Name of the destination NAT rule.	Type a name for the rule.
Source address	Specifies the source address to match. You can configure multiple addresses or subnets.	Enter one or more addresses or subnets.
Destination address	Specifies the destination address to match. You can configure multiple addresses or subnets.	Enter one or more addresses or subnets.
Destination port	Specifies the destination port to match.	Enter a single port number.
Then	Specifies the action of the destination NAT rule.	Destination nat off—No destination NAT is performed. Destination NAT pool—Use the specified pool to perform destination NAT.

Understanding Source NAT

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

The following types of source NAT are supported:

- Translation of the original source IP address to the egress interface's IP address (also called interface NAT). Port address translation is always performed.
- Translation of the original source IP address to an IP address from a user-defined address pool without port address translation. The association between the original source IP address to the translated source IP address is dynamic. However, once there is an association, the same association is used for the same original source IP address for new traffic that matches the same NAT rule.
- Translation of the original source IP address to an IP address from a user-defined address pool with port address translation. The association between the original source IP address to the translated source IP address is dynamic. Even if an

association exists, the same original source IP address may be translated to a different address for new traffic that matches the same NAT rule.

- Translation of the original source IP address to an IP address from a user-defined address pool by shifting the IP addresses. This type of translation is one-to-one, static, and without port address translation. If the original source IP address range is larger than the IP address range in the user-defined pool, untranslated packets are dropped.

Related Topics

- Source NAT Configuration Overview on page 975

Source NAT Configuration Overview

The main configuration tasks for source NAT are as follows:

- Configure a source NAT address pool that aligns with your network and security requirements (not needed for interface NAT).
- Configure pool utilization alarms (optional)—Specify thresholds for pool utilization.
- Configure address persistent (optional)—Ensures that the same IP address is assigned from the source NAT pool to a host for multiple concurrent sessions.
- Configure source NAT rules that align with your network and security requirements.
- Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Source NAT Pools

For source NAT address pools, specify the following:

- Name of the source NAT address pool.
- Up to eight address or address ranges.

NOTE: Do not overlap NAT addresses for source NAT, destination NAT, and static NAT within one routing instance.

- Routing instance to which the pool belongs (the default is the main inet.0 routing instance).
- No port translation (optional)—By default, port address translation is performed with source NAT. If you specify the port no-translation option, the number of hosts that the source NAT pool can support is limited to the number of addresses in the pool.
- Overflow pool (optional)—Packets are dropped if there are no addresses available in the designated source NAT pool. To prevent that from happening when the port no-translation option is configured, you can specify an overflow pool. Once

addresses from the original source NAT pool are exhausted, IP addresses and port numbers are allocated from the overflow pool. A user-defined source NAT pool or an egress interface can be used as the overflow pool. (When the overflow pool is used, the pool ID is returned with the address.)

- IP address shifting (optional)—A range of original source IP addresses can be mapped to another range of IP addresses by shifting the IP addresses. Specify the host-address-base option with the base address of the original source IP address range.

Pool Utilization Alarms

When the raise-threshold option is configured for source NAT, an SNMP trap is triggered if the source NAT pool utilization rises above this threshold. If the optional clear-threshold option is configured, an SNMP trap is triggered if the source NAT pool utilization drops below this threshold. If clear-threshold is not configured it is set by default to 80 percent of the raise-threshold value.

Persistent Addresses

By default, port address translation is performed with source NAT. However, an original source address may not be translated to the same IP address for different traffic that originates from the same host. The source NAT address-persistent option ensures that the same IP address is assigned from the source NAT pool to a specific host for multiple concurrent sessions.

Source NAT Rules

Source NAT rules specify two layers of match conditions:

- Traffic direction—Allows you to specify combinations of from interface, from zone, or from routing-instance and to interface, to zone, or to routing-instance. You cannot configure the same from and to contexts for different rule sets.
- Packet information—Can be source and destination IP addresses or subnets.

If multiple source NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 to zone 2 and rule B specifies traffic from zone 1 to interface ge-0/0/0, rule B is used to perform source NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match. For more information about rule set matching, see “Understanding NAT Rule Sets and Rules” on page 965.

The actions you can specify for a source NAT rule are:

- off—Do not perform source NAT.
- pool—Use the specified user-defined address pool to perform source NAT.
- interface—Use the egress interface’s IP address to perform source NAT.

Source NAT rules are applied to traffic in the first packet that is processed for the flow or in the fast path for the ALG. Source NAT rules are processed after static NAT rules, destination NAT rules, and reverse mapping of static NAT rules and after route and security policy lookup.

Related Topics

- Example: Configuring Source NAT Pools on page 977

Example: Configuring Source NAT Pools

Source NAT pools provide JUNOS Software with a supply of addresses from which to draw when translation is performed.

NOTE: When the JUNOS Software performs source NAT, security policies are applied first, and then the address in question is translated according to configured source NAT rules.

Before You Begin
For background information, read: ■ Source NAT Configuration Overview on page 975

In this example, you perform the following tasks:

- Configure a source NAT pool with port address translation.
- Configure a source NAT pool without port address translation.
- Configure a source NAT pool by specifying an IP address and prefix.
- Configure a source NAT pool with host address shifting.
- Configure pool utilization alarms and address persistent.

CLI Configuration

To configure a source NAT pool with port address translation:

```
user@host# set security nat source pool src-nat-with-pat routing-instance red
user@host# set security nat source pool src-nat-with-pat address 10.1.1.1
user@host# set security nat source pool src-nat-with-pat address 10.1.1.10 to
  10.1.1.20
```

To configure a source NAT pool without port address translation:

```
user@host# set security nat source pool src-nat-without-pat routing-instance red
user@host# set security nat source pool src-nat-without-pat address 10.1.1.100 to
  10.1.1.150
```

```
user@host# set security nat source pool src-nat-without-pat port no-translation
user@host# set security nat source pool src-nat-without-pat overflow-pool
  src-nat-with-pat
```

To configure a source NAT pool by specifying an IP address and prefix:

```
user@host# set security nat source pool src-nat-with-prefix routing-instance red
user@host# set security nat source pool src-nat-with-prefix address 10.1.1.0/24
```

To configure a source NAT pool with host address shifting:

```
user@host# set security nat source pool src-nat-by-shifting routing-instance red
user@host# set security nat source pool src-nat-by-shifting address 10.1.1.200 to
  10.1.1.250
user@host# set security nat source pool src-nat-by-shifting host-address-base
  192.168.1.1
```

To configure pool utilization alarms and address to be persistent:

```
user@host# set security nat source pool-utilization-alarm raise-threshold 80
  clear-threshold 70
user@host# set security nat source address-persistent
```

Understanding NAT Source Pools with PAT

Using the source pool with Port Address Translation (PAT), JUNOS Software translates both the source IP address and the port number of the packets. When PAT is used, multiple hosts can share the same IP address.

Before You Begin
For background information, read “Understanding Source NAT” on page 974.

JUNOS Software maintains a list of assigned port numbers to distinguish what session belongs to which host. When PAT is enabled, up to 64,500 hosts can share a single IP address. Each source pool can contain multiple IP addresses, multiple IP address ranges, or both. For a source pool with PAT, JUNOS Software may assign different addresses to a single host for different concurrent sessions, unless the source pool or JUNOS Software has the persistent address feature enabled.

This topic covers:

- Port Ranges on page 978
- Address Persistent on page 979
- Related Topics on page 979

Port Ranges

For interface source pool and source pool with PAT, range (1024, 65535) is available for port number mapping per IP address. Within range (1024, 63487) one port is

allocated at a time. In range (63488, 65535), two ports are allocated at a time for RTP/RTCP applications such as SIP, H.323, and RTSP.

Address Persistent

When a host initiates several sessions that match a policy that requires network address translation and is assigned an address from a source pool that has PAT enabled, the device assigns a different source IP address for each session. Such random address assignment can be problematic for services that create multiple sessions that require the same source IP address for each session. For example, it is important to have the same IP address for multiple sessions when using the AOL Instant Message (AIM) client.

To ensure that the router assigns the same IP address from a source pool to a host for multiple concurrent sessions, you can enable a persistent IP address per router.

Related Topics

- Example: Configuring Source NAT on page 980
- Example: Configuring a Persistent Address on page 984
- Disabling Port Randomization for Source NAT on page 979

Disabling Port Randomization for Source NAT

For pool-based source NAT and interface NAT, port numbers are allocated randomly by default. Although randomized port number allocation can provide protection from security threats such as DNS poison attacks, it can also affect performance and memory usage for pool-based source NAT.

You can disable port randomization by using the port-randomization disable statement at the [edit security nat source] hierarchy level. To re-enable port randomization, use the port-randomization statement at the [edit security nat source] hierarchy level.

CLI Configuration

```
user@host# set security nat source port-randomization disable
```

Understanding NAT Source Pools Without PAT

When you define a source pool, JUNOS Software enables PAT by default. To disable PAT, you must specify no port translation when you are defining a source pool.

Before You Begin
For background information, read “Understanding Source NAT” on page 974.

When using a source pool without PAT, JUNOS Software performs source Network Address Translation for the IP address without performing PAT for the source port number. For applications that require that a particular source port number remain fixed, you must use source pool without PAT.

The source pool can contain multiple IP addresses, multiple IP address ranges, or both. For source pool without PAT, JUNOS Software assigns one translated source address to the same host for all its concurrent sessions.

This topic covers:

- Source Pool Utilization Alarm on page 980
- Related Topics on page 980

Source Pool Utilization Alarm

Pool utilization for each source pool without PAT is computed. You can turn on pool utilization alarm by configuring alarm thresholds. An SNMP trap is triggered every time pool utilization rises above a threshold and goes below a threshold.

Related Topics

- Example: Configuring Source NAT on page 980

Example: Configuring Source NAT

When source NAT is performed, source pools provide JUNOS Software with a supply of addresses from which to draw. When a NAT rule requires NAT and references a specific source pool, JUNOS Software draws addresses from that pool when translation is performed.

NOTE: When the JUNOS Software performs source NAT, security policies are applied first, and then the address in question is translated according to configured NAT source rules.

Before You Begin

For background information, read:

- Source NAT Configuration Overview on page 975
- Example: Configuring Source NAT Pools on page 977

Source NAT rules have three available actions:

- off — Do not perform source NAT.
- pool — Use user-defined source NAT pools to perform source NAT.
- interface — Use the egress interface IP address to perform source NAT.

NOTE: The off action is a useful command for detail control when you are configuring source NAT rules. For example, you can configure a rule that says, "if rule A is from zone1 to zone2, do source NAT." However, you do not want source NAT performed if the traffic egresses from interface if2, which belongs to zone2. In that case, you can define a rule B, which is from zone1 to if2 with off as the source NAT action.

In this example, you perform the following tasks:

- Define a source NAT pool for traffic from routing-instance ri-2 to routing-instance ri-1 with any source IP address and destination IP address 30.1.1.1. Map the source IP address to 10.1.1.1.
- Define a source NAT pool for traffic from zone z3 or z4 to routing-instance ri-1 with any source IP address and destination IP address 30.1.1.2. Map the source IP address to 10.1.1.2.
- Define a source NAT pool for traffic from interface fe-0/0/0.0 or fe-0/0/1.0 to interface ge-1/0/0.0 or ge-1/0.1.0 with any source IP address and destination IP address 30.1.1.3. Map the source IP address to 10.1.1.3.
- Define a source NAT pool for traffic from routing-instance ri-2 to zone z2 with any source IP address and destination IP address 30.1.1.4. Map the source IP address to 10.1.1.4.
- Define a source NAT pool for traffic from routing-instance ri-2 to routing-instance ri-1 with any source IP address and destination IP address 30.1.1.5. Map the source IP address to 10.1.1.5.

CLI Configuration

To configure source NAT pools:

```
user@host# set security nat source pool spool-1 routing-instance ri-1
user@host# set security nat source pool spool-1 address 10.1.1.1
user@host# set security nat source pool spool-2 routing-instance ri-1
user@host# set security nat source pool spool-2 address 10.1.1.2
user@host# set security nat source pool spool-3 routing-instance ri-1
user@host# set security nat source pool spool-3 address 10.1.1.3
user@host# set security nat source pool spool-4 routing-instance ri-1
user@host# set security nat source pool spool-4 address 10.1.1.4
user@host# set security nat source pool spool-5 routing-instance ri-1
user@host# set security nat source pool spool-5 address 10.1.1.5
```

To configure source NAT rules:

```
user@host# set security nat source rule-set rs1 from routing-instance ri-2
user@host# set security nat source rule-set rs1 to routing-instance ri-1
user@host#set security nat source rule-set rs1 rule r1 match destination-address
  30.1.1.1
user@host# set security nat source rule-set rs1 rule r1 then source-nat pool spool-1
user@host# set security nat source rule-set rs1 rule r5 match destination-address
  30.1.1.5
user@host# set security nat source rule-set rs1 rule r5 then source-nat pool spool-5
```

```
user@host# set security nat source rule-set rs2 from zone [z3 z4]
user@host#set security nat source rule-set rs2 to routing-instance ri-1
user@host# set security nat source rule-set rs2 rule r2 match destination-address
  30.1.1.2
user@host# set security nat source rule-set rs2 rule r2 then source-nat pool spool-2

user@host# set security nat source rule-set rs3 from interface [fe-0/0/0.0
  fe-0/0/1.0]
user@host# set security nat source rule-set rs3 to interface [ge-1/0/0.0 ge-1/0.1.0]
user@host# set security nat source rule-set rs3 rule r3 match destination-address
  30.1.1.3
user@host# set security nat source rule-set rs3 rule r3 then source-nat spool-3

user@host# set security nat source rule-set rs4 from routing-instance ri-2
user@host# set security nat source rule-set rs4 to zone z2
user@host# set security nat source rule-set rs4 rule r4 match destination-address
  30.1.1.4
user@host# set security nat source rule-set rs4 rule r4 then source-nat pool spool-4
```

Configuring Source NAT—Quick Configuration

You can use J-Web Quick Configuration to quickly configure source NAT.

Before You Begin
For background information, read: ■ Source NAT Configuration Overview on page 975 ■ Example: Configuring Source NAT Pools on page 977

To configure source NAT with Quick Configuration:

1. Select **Configure > NAT > Source NAT**.
2. Fill in the options as shown in Table 116 on page 982.
3. Click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main Configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.

Table 116: Source NAT Configuration Options

Field	Function	Action
Source NAT		

Table 116: Source NAT Configuration Options *(continued)*

Field	Function	Action
Address Persistent	Ensures that the router assigns the same IP address from a source pool to a host for multiple concurrent sessions.	Enable Address Persistent.
Pool Utilization Alarm	Turns on pool utilization alarm by configuring alarm thresholds.	Raise Threshold—Enter a number between 50–100 to raise the threshold for the pool utilization alarm. Clear Threshold—Enter a number between 40–100 to clear the threshold for the pool utilization alarm.
Source Address Pool	Defines one or more source NAT pools.	Click **Add.**
Source Rule Sets	Defines one or more rule sets.	Click **Add.**
Source NAT Pool		
Name	Descriptive name for the pool.	Type a name for the pool.
Routing Instance	Specify the routing instance to which the pool is bound. By default, the routing instance is inet.0.	Enter the routing instance name.
Host Address Base	Specify the base address of the original source IP address range. This is used for IP shifting.	Enter the IP address.
Port	Specifies whether port translation is performed (default).	Default—Port translation is enabled. No translation—Disables port translation. Translation—Enables port translation. Enter a port range.
Overflow	Specifies a source pool to use when the current address pool is exhausted. The pool can be a user-defined pool or an interface's IP address.	Off—No overflow pool. Interface—Enter the IP address of the interface. Pool name—Select a user-defined pool.
Source Addresses	Specifies address prefixes, addresses, or a range of addresses or address prefixes.	Click **Add** or **Delete** to configure or remove address entries.
Source NAT Rule Set		
Name	Name of the rule set.	Type a name for the rule set.
From	Specifies where the packet is coming from. If a packet matches more than one rule set, the rule set with the more specific match is used. An interface match is considered more specific than a zone match, which is more specific than a routing instance match.	Routing Instance—Specify the source routing instance. Interface—Specify the source interface. Zone—Specify the source zone.

Table 116: Source NAT Configuration Options *(continued)*

Field	Function	Action
To	Specifies to where the packet is going. If a packet matches more than one rule set, the rule set with the more specific match is used. An interface match is considered more specific than a zone match, which is more specific than a routing instance match.	Routing Instance—Specify the destination routing instance. Interface—Specify the destination interface. Zone—Specify the destination zone.
Rules	Set of conditions that are evaluated, in order, for a match to perform the specified source NAT action.	
Rule name	Name of the source NAT rule.	Type a name for the rule.
Source address	Specifies the source address to match. You can configure multiple addresses or subnets.	Enter one or more addresses or subnets.
Destination address	Specifies the destination address to match. You can configure multiple addresses or subnets.	Enter one or more addresses or subnets.
Then	Specifies the action of the source NAT rule.	Source nat off—No source NAT is performed. Source NAT pool—Use the specified pool to perform source NAT. Source NAT interface—Use the egress interface address to perform source NAT.

Example: Configuring a Persistent Address

You can enable the persistent address feature on a router to ensure that the router assigns the same IP address from a source pool to a host for multiple concurrent sessions. You can reference a pool set within a policy to easily reference multiple interfaces within a zone.

Before You Begin

For background information, read:

- Address Persistent on page 979

In this example, you perform the following tasks:

- Enable a persistent address globally.
- Configure pool utilization thresholds.

To configure a persistent address and pool sets, use the CLI configuration editor.

This topic covers:

- CLI Configuration on page 985
- Related Topics on page 985

CLI Configuration

```
user@host# set security nat source address-persistent
user@host# set security nat source pool-utilization-alarm raise-threshold 90
user@host# set security nat source pool-utilization-alarm clear-threshold 80
```

Related Topics

- Example: Configuring Source NAT on page 980

Understanding Persistent NAT

Persistent NAT allows applications to use the Session Traversal Utilities for NAT (STUN) protocol when passing through NAT firewalls (see "Understanding Session Traversal Utilities for NAT (STUN) Protocol" on page 986). Persistent NAT ensures that all requests from the same internal transport address are mapped to the same external transport address by the NAT device closest to the STUN server.

The following types of persistent NAT can be configured on the Juniper Networks device:

- Any remote host—All requests from a specific internal IP address and port are mapped to the same external IP address and port. Any external host can send a packet to the internal host using the mapped external transport address when the incoming policy from external to internal is configured.
- Target host—All requests from a specific internal IP address and port are mapped to the same external IP address and port. An external host can send a packet to an internal host only if the internal host had previously sent a packet to the external host's IP address.
- Target host port—All requests from a specific internal IP address and port are mapped to the same external IP address and port. An external host can send a packet to an internal host only if the internal host had previously sent a packet to the external host's IP address and port.

You configure any of the persistent NAT types with source NAT rules. The source NAT rule action can use a source NAT pool (with or without port translation) or an egress interface. Persistent NAT is not applicable for destination NAT, because persistent NAT bindings are based on outgoing sessions from internal to external.

NOTE: Port overloading is used in JUNOS Software only for normal interface NAT traffic. Persistent NAT does not support port overloading, and you must explicitly disable port overloading with the port-overloading off option at the [edit security nat source] hierarchy level.

To configure security policies to permit or deny persistent NAT traffic, you can use two new predefined services—junos-stun and junos-persistent-nat.

NOTE: Persistent NAT is different from the persistent address feature (see "Example: Configuring a Persistent Address" on page 984). The persistent address feature applies to address mappings for source NAT pools configured on the device. The persistent NAT feature applies to address mappings on an external NAT device, and is configured for a specific source NAT pool or egress interface. Also, persistent NAT is intended for use with STUN client-server applications.

Related Topics

- Understanding Session Traversal Utilities for NAT (STUN) Protocol on page 986
- Persistent NAT Configuration Overview on page 987

Understanding Session Traversal Utilities for NAT (STUN) Protocol

Many video and voice applications do not work properly in a NAT environment. For example, Session Initiation Protocol (SIP), used with VoIP, encodes IP addresses and port numbers within application data. If a NAT firewall exists between the requestor and receiver, the translation of the IP address and port number in the data invalidates the information.

Also, a NAT firewall does not maintain a pinhole for incoming SIP messages. This forces the SIP application to either constantly refresh the pinhole with SIP messages or use an ALG to track registration, a function that may or may not be supported by the gateway device.

The Session Traversal Utilities for NAT (STUN) protocol, first defined in *RFC 3489, Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)* and then later in *RFC 5389, Session Traversal Utilities for NAT*, is a simple client-server protocol. A STUN client sends requests to a STUN server, which returns responses to the client. A STUN client is usually part of an application that requires a public IP address and/or port. STUN clients can reside in an end system such as a PC or in a network server whereas STUN servers are usually attached to the public Internet.

NOTE: Both the STUN client and STUN server must be provided by the application. Juniper Networks does not provide a STUN client or server.

The STUN protocol allows a client to:

- Discover whether the application is behind a NAT firewall.
- Determine the type of NAT binding being used (see "Understanding Persistent NAT" on page 985).

- Learn the IP address and port bindings allocated by NAT closest to the STUN server. (There may be multiple levels of NAT between the STUN client and the STUN server.)

The client application can use the IP address binding information within protocols such as SIP and H.323.

Related Topics

- Persistent NAT Configuration Overview on page 987

Persistent NAT Configuration Overview

To configure persistent NAT, specify the following with the source NAT rule action (for either a source NAT pool or an egress interface):

- The type of persistent NAT—One of the following: any remote host, target host, or target host port (see "Understanding Persistent NAT" on page 985).
- (Optional) Inactivity timeout—Time, in seconds, that the persistent NAT binding remains in the device's memory when all the sessions of the binding entry are gone. When the configured timeout is reached, the binding is removed from memory. The default value is 5 minutes. Configure a value from 60 through 7200 seconds.

  When all sessions of a persistent NAT binding are gone, the binding remains in a query state in the SRX Series device's memory for the specified inactivity timeout period. The query binding is automatically removed from memory when the inactivity timeout period expires (the default is 5 minutes). You can explicitly remove all or specific persistent NAT query bindings with the clear security nat source persistent-nat-table command.
- (Optional) Maximum session number—Maximum number of sessions with which a persistent NAT binding can be associated. The default is 30 sessions. Configure a value from 8 through 100.

For interface NAT, you need to explicitly disable port overloading with the port-overloading off option at the [edit security nat source] hierarchy level.

Finally, there are two predefined services that you can use in security policies to permit or deny STUN and persistent NAT traffic:

- junos-stun—STUN protocol traffic
- junos-persistent-nat—Persistent NAT traffic

For the any remote host persistent NAT type, the direction of the security policy is from external to internal. For target host or target host port persistent NAT types, the direction of the security policy is from internal to external.

Related Topics

- Example: Configuring Persistent NAT with Source NAT Address Pool on page 988
- Example: Configuring Persistent NAT with Interface NAT on page 989

Example: Configuring Persistent NAT with Source NAT Address Pool

You can configure any of the persistent NAT types with source NAT rules. The example in this section shows how to configure persistent NAT when source NAT is performed with a user-defined address pool.

Before You Begin

For background information, read:

- Understanding Persistent NAT on page 985
- Persistent NAT Configuration Overview on page 987

The following example configures the target host persistent NAT type when source NAT is performed. In the following configuration, the source NAT address pool sp1 consists of the address 30.1.1.5/32. The source NAT rule set srs1 configures the following:

- Traffic direction is from zone internal to zone external.
- For packets with source address in the 40.1.1.0/24 subnet (internal phones) and destination address 20.20.20.105/32 (STUN server), use the source NAT pool sp1 to perform source NAT with the target host persistent NAT type.
- Set the persistent NAT inactivity-timeout to 180 seconds.

CLI Configuration

To configure the source NAT address pool:

```
user@host# set security nat source pool sp1 address 30.1.1.5/32
```

To configure the source NAT rule set:

```
user@host# set security nat source rule-set srs1 from zone internal
user@host# set security nat source rule-set srs1 to zone external
user@host# set security nat source rule-set srs1 rule sr1 match source-address
  40.1.1.0/24
user@host# set security nat source rule-set srs1 rule sr1 match destination-address
  20.20.20.105/32
user@host# set security nat source rule-set srs1 rule sr1 then source-nat pool sp1
user@host# set security nat source rule-set srs1 rule sr1 then source-nat pool
  persistent-nat permit target-host
```

```
user@host# set security nat source rule-set srs1 rule sr1 then source-nat pool
    persistent-nat inactivity-timeout 180
```

For the target host persistent NAT type, configure a security policy to allow persistent NAT traffic from the internal network (internal zone) to the external network (external zone).

To configure a security policy to allow STUN traffic from internal SIP phones to an external STUN server:

```
user@host# set security policies from-zone internal to-zone external policy stun_traffic
    match source-address internal_phones destination-address stun_server application
    junos-stun
user@host# set security policies from-zone internal to-zone external policy stun_traffic
    then permit
```

To configure a security policy to allow SIP proxy traffic from internal SIP phones to an external SIP proxy server:

```
user@host# set security policies from-zone internal to-zone external policy
    sip_proxy_traffic match source-address internal_phones destination-address
    sip_proxy_server application junos-sip
user@host# set security policies from-zone internal to-zone external policy
    stun_proxy_traffic then permit
```

To configure a security policy to allow SIP traffic from internal to external SIP phones:

```
user@host# set security policies from-zone internal to-zone external policy sip_traffic
    match source-address internal_phones destination-address external_phones
    application junos-persistent-nat
user@host# set security policies from-zone internal to-zone external policy sip_traffic
    then permit
```

Example: Configuring Persistent NAT with Interface NAT

You can configure any of the persistent NAT types with source NAT rules. The example in this section shows how to configure persistent NAT when interface NAT is used to perform source NAT. For interface NAT, port overloading must be disabled.

Before You Begin

For background information, read:

- Understanding Persistent NAT on page 985
- Persistent NAT Configuration Overview on page 987

The following example configures the any remote host persistent NAT type when interface NAT is performed. The interface NAT rule set int1 configures the following:

- Traffic direction is from interface ge-0/0/1.0 to interface ge-0/0/2.0.

- For packets with source address 40.1.1.0/24 (internal phones) and destination address 20.20.20.105/32 (STUN server), perform interface NAT with the any remote host persistent NAT type.

You must also disable port overloading for interface NAT.

CLI Configuration

To configure the interface NAT rule set:

```
user@host# set security nat source rule-set int1 from interface ge-0/0/1.0
user@host# set security nat source rule-set int1 to interface ge-0/0/2.0
user@host# set security nat source rule-set int1 rule in1 match source-address
  40.1.1.0/24
user@host# set security nat source rule-set int1 rule in1 match destination-address
  20.20.20.105/32
user@host# set security nat source rule-set int1 rule in1 then source-nat interface
  persistent-nat permit any-remote-host
```

To disable port overloading for interface NAT:

```
user@host# set security nat source interface port-overloading off
```

For the any remote host persistent NAT type, configure a security policy to allow persistent NAT traffic from the external network (external zone) to the internal network (internal zone).

To configure a security policy to allow STUN traffic from the internal SIP phones to the external STUN server:

```
user@host# set security policies from-zone internal to-zone external policy stun_traffic
  match source-address internal_phones destination-address stun_server application
  junos-stun
user@host# set security policies from-zone internal to-zone external policy stun_traffic
  then permit
```

To configure a security policy to allow SIP proxy traffic from the internal SIP phones to the external SIP proxy server:

```
user@host# set security policies from-zone internal to-zone external policy
  sip_proxy_traffic match source-address internal_phones destination-address
  sip_proxy_server application junos-sip
user@host# set security policies from-zone internal to-zone external policy
  stun_proxy_traffic then permit
```

To configure a security policy to allow SIP traffic from external SIP phones to internal SIP phones:

```
user@host# set security policies from-zone external to-zone internal policy sip_traffic
  match source-address external_phones destination-address internal_phones
  application junos-persistent-nat
user@host# set security policies from-zone external to-zone internal policy sip_traffic
  then permit
```

Configuring Proxy ARP

You use NAT proxy ARP functionality to configure proxy ARP entries for IP addresses that require either source or destination NAT and that are in the same subnet as the ingress interface.

NOTE: On SRX Series devices, you must explicitly configure NAT proxy ARP.

When configuring NAT proxy ARP, you must specify the logical interface on which to configure proxy ARP. Then you enter an address or address range.

The device performs proxy ARP for the following conditions:

- When addresses defined in the static NAT and source NAT pool are in the same subnet as that of the ingress interface
- When addresses in the original destination address entry in the destination NAT rules are in the same subnet as that of the ingress interface

CLI Configuration

user@host# **set security nat proxy-arp interface fe-0/0/0.0 address 10.1.1.10 to 10.1.1.20**

Verifying NAT Configuration

The NAT trace options hierarchy configures trace file and flags for verification purposes. J Series and SRX Series devices have two main components. Those are the Routing Engine (RE) and the Packet Processing Engine (PFE). The PFE is divided into the ukernel portion and the real-time portion. For verification, you can turn on flags individually to debug NAT functionality on the RE, ukernel PFE, or real-time PFE. The trace data is written to/var/log/security-trace by default.

NOTE: If session logging has been enabled in the policy configurations on the device, the session logs will include specific NAT details for each session. See “Monitoring Policy Statistics” on page 94 for information on how to enable session logging and “Information Provided in Session Log Entries” on page 30 for a description of information provided in session logs.

CLI Configuration

user@host# **set security nat traceoptions flag all**
user@host# **set security nat traceoptions flag destination-nat-pfe**
user@host# **set security nat traceoptions flagdestination-nat-re**
user@host# **set security nat traceoptions flag destination-nat-rti**
user@host# **set security nat traceoptions flag destination-nat-pfe**
user@host# **set security nat traceoptions flag source-nat-pfe**

```
user@host# set security nat traceoptions flag source-nat-re
user@host# set security nat traceoptions flag source-nat-rt
user@host# set security nat traceoptions flag static-nat-pfe
user@host# set security nat traceoptions flag static-nat-re
user@host# set security nat traceoptions flag static-nat-rt
```

Part 12
GPRS

■ General Packet Radio Service on page 995

Chapter 30

General Packet Radio Service

For information about which devices support the features documented in this chapter, see the *JUNOS Software Feature Support Reference for SRX Series and J Series Devices*.

This topic includes the following:

- Understanding GPRS on page 995
- Understanding GPRS Tunneling Protocol on SRX Series Devices on page 996
- Understanding Policy-Based GPRS Tunneling Protocol on page 998
- Understanding GPRS Tunneling Protocol Inspection Object on page 1000
- Understanding GTP Message Filtering on page 1001
- Understanding GTP Information Elements on page 1006
- Understanding GGSN Redirection on page 1011

Understanding GPRS

General Packet Radio Service networks connect to several external networks including those of roaming partners, corporate customers, GPRS Roaming Exchange (GRX) providers, and the public Internet. GPRS network operators face the challenge of protecting their network while providing and controlling access to and from these external networks. Juniper Networks provides solutions to many of the security problems plaguing GPRS network operators.

In the GPRS architecture, the fundamental cause of security threats to an operator's network is the inherent lack of security in GPRS tunneling protocol (GTP). GTP is the protocol used between GPRS support nodes (GSNs). Communication between different GPRS networks is not secure because GTP does not provide any authentication, data integrity, or confidentiality protection. Implementing Internet Protocol security (IPsec) for connections between roaming partners, setting traffic rate limits, and using stateful inspection can eliminate a majority of the GTP's security risks. Juniper Networks security devices mitigate a wide variety of attacks on the Gp, Gn, and Gi interfaces. The GTP firewall features in JUNOS Software address key security issues in mobile operators' networks.

NOTE: GPRS is supported on the SRX3400, SRX3600, SRX5600, and SRX5800 SRX Series devices

This chapter describes the GTP features that JUNOS Software supports and explains how you can configure them on an SRX Series device. This chapter contains the following sections:

Understanding GPRS Tunneling Protocol on SRX Series Devices

This topic includes:

- Gp and Gn Interfaces on page 996
- Gi Interface on page 997
- Operational Modes on page 998

The GPRS Tunneling Protocol (GTP) is used to establish a GTP tunnel, for individual mobile stations (MS), between a Serving GPRS Support Node (SGSN) and a Gateway GPRS Support Node (GGSN). A GTP tunnel is a channel between GSNs through which two hosts exchange data. The SGSN receives packets from the MS and encapsulates them within a GTP header before forwarding them to the GGSN through the GTP tunnel. When the GGSN receives the packets, it decapsulates them and forwards them to the external host.

A Juniper Networks security device provides firewall protection for the following types of GPRS interfaces:

- Gn—The Gn interface is the connection between an SGSN and a GGSN within the same Public Land Mobile Network (PLMN).
- Gp—The Gp interface is the connection between two Public Land Mobile Network PLMNs.
- Gi—The Gi interface is the connection between a GGSN and the Internet or destination networks connected to a PLMN.

NOTE: The term *interface* has different meanings in JUNOS Software and in GPRS technology. In JUNOS Software, an interface is like a doorway to a security zone and allows traffic to enter and exit the zone. In GPRS, an interface is a connection, or a reference point, between two components of a GPRS infrastructure, for example, an SGSN and a GGSN.

Gp and Gn Interfaces

You implement a security device on the Gn interface to protect core network assets such as the SGSN and GGSN. To secure GTP tunnels on the Gn interface, you place the security device between SGSNs and GGSNs within a common PLMN.

When you implement a security device to the Gp interface, you protect a PLMN from another PLMN. To secure GTP tunnels on the Gp interface, you place the SGSNs and GGSNs of a PLMN behind the security device so that all traffic, incoming and outgoing, goes through the firewall.

Figure 179 on page 997 illustrates the placement of Juniper Networks SRX Series devices used to protect PLMNs on the Gp and Gn interfaces

Figure 179: Gp and Gn Interfaces

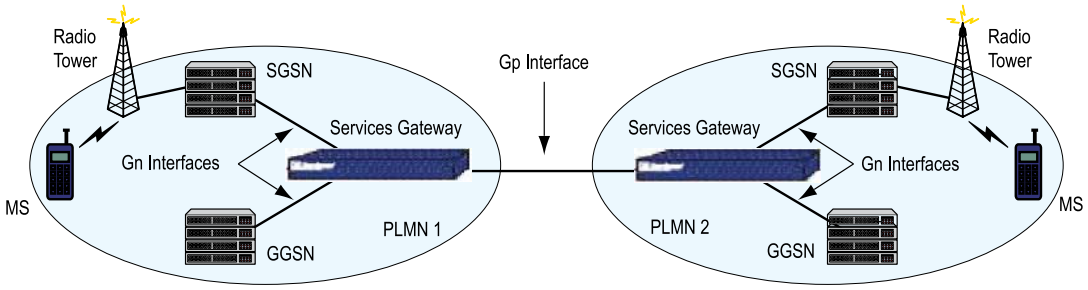

Gi Interface

When you implement a security device on the Gi interface, you can simultaneously control traffic for multiple networks, protect a PLMN against the Internet and external networks, and protect mobile users from the Internet and other networks. JUNOS Software provides a great number of virtual routers, making it possible for you to use one virtual router per customer network and thereby allow the separation of traffic for each customer network.

The security device can securely forward packets to the Internet or destination networks using the Layer 2 Tunneling Protocol (L2TP) for IPsec virtual private network (VPN) tunnels. (Note, however, that SRX Series devices do not support full L2TP.)

Figure 180 on page 998 illustrates the implementation of a security device to protect a PLMN on the Gi interface.

Figure 180: Gi Interface

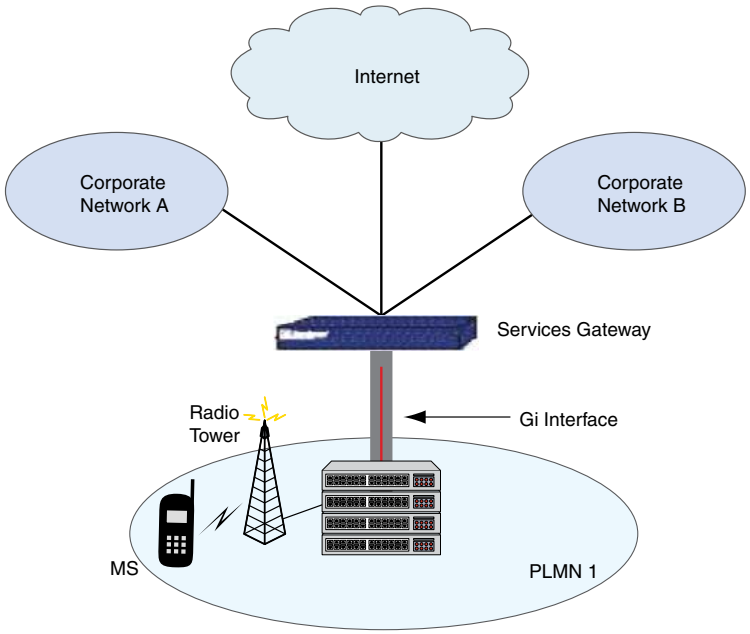

Operational Modes

JUNOS Software supports two interface operational modes with GTP: transparent mode and route mode. If you want the security device to participate in the routing infrastructure of your network, you can run it in route mode. This requires a certain amount of network redesign. Alternatively, you can implement the security device into your existing network in transparent mode without having to reconfigure the entire network. In transparent mode, the security device functions as a Layer 2 switch or bridge, and the IP addresses of interfaces are set at 0.0.0.0, making the presence of the security device invisible, or *transparent*, to users.

JUNOS Software supports Network Address Translation (NAT) on interfaces and policies that do not have GTP inspection enabled.

Currently in JUNOS Software, route mode supports active/passive, and active/active chassis cluster. Transparent mode supports active/passive only.

For more information about operational modes and high availability, see the *JUNOS Software Interfaces and Routing Configuration Guide* and “Chassis Cluster” on page 867, respectively.

Understanding Policy-Based GPRS Tunneling Protocol

By default, the PLMN that the SRX Series device protects is in the Trust zone. The SRX Series device protects the PLMN in the Trust zone against other PLMNs in other zones. You can place all the PLMNs against which you are protecting your PLMN in

the Untrust zone, or you can create user-defined zones for each PLMN. A PLMN can occupy one security zone or multiple security zones.

You must create policies to enable traffic to flow between zones and PLMNs. Policies contain rules that permit, deny, or tunnel traffic. A SRX Series device performs GTP policy filtering by checking every GTP packet against policies that regulate GTP traffic and by then forwarding, dropping, or tunneling the packet based on these policies.

By selecting the GTP service in a policy, you enable the SRX Series device to permit, deny, or tunnel GTP traffic. However, this does not enable the device to inspect GTP traffic. In order for the SRX Series device to inspect GTP traffic, you must apply a GTP configuration, also referred to as a *GTP inspection object*, to a policy.

Before you can apply a GTP configuration to a policy, you must first create a GTP inspection object (see "Understanding GPRS Tunneling Protocol Inspection Object" on page 1000). You can apply only one GTP inspection object per policy, but you can apply a GTP inspection object to multiple policies. Using policies, you can permit or deny the establishment of GTP tunnels from certain peers such as an SGSN.

You can configure policies that specify "Any" as the source or destination zone (thereby including all hosts in the zone), and you can configure policies that specify multiple source and destination addresses.

In policies, you can enable traffic logging.

Configuring Policies to Enable GTP Inspection

In this example, you configure interfaces and create addresses and two policies to allow bidirectional traffic between two networks within the same PLMN. You also apply a GTP inspection object to the policies. Enter the following commands in edit mode.

CLI Configuration

1. **Enable GTP**
   ```
   user@host# set security gprs gtp enable
   user@host# commit
   user@host# exit
   user@host# request system reboot
   ```

NOTE: GTP is disabled by default. The device must be restarted after enabling GTP.

2. **GTP Inspection Object**
   ```
   user@host# set security gprs gtp profile gtp1
   ```

3. **Interfaces**
   ```
   user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.0.0.254/8
   user@host# set interfaces ge-0/0/2 unit 0 family inet address 3.0.0.254/8
   ```

4. **Security Zones**
   ```
   user@host# set security zones security-zone sgsn interfaces ge-0/0/1.0
   ```

```
user@host# set security zones security-zone sgsn host-inbound-traffic
  system-services all
user@host# set security zones security-zone sgsn host-inbound-traffic protocols
  all
user@host# set security zones security-zone ggsn interfaces ge-0/0/2.0
user@host# set security zones security-zone ggsn host-inbound-traffic
  system-services all
user@host# set security zones security-zone ggsn host-inbound-traffic protocols
  all
```

5. **Addresses**

```
user@host# set security zones security-zone sgsn address-book address local-sgsn
  2.0.0.5/32
user@host# set security zones security-zone ggsn address-book address
  remote–ggsn 3.0.0.6/32
```

6. **Policies**

```
user@host# set security policies from-zone sgsn to-zone ggsn policy sgsn_to_ggsn
  match source-address local-sgsn destination-address remote–ggsn application
  junos-gprs-gtp
user@host# set security policies from-zone sgsn to-zone ggsn policy sgsn_to_ggsn
  then permit application-services gprs-gtp-profile gtp1
user@host# set security policies from-zone ggsn to-zone sgsn policy ggsn_to_sgsn
  match source-address remote–ggsn destination-address local–sgsn application
  junos-gprs-gtp
user@host# set security policies from-zone ggsn to-zone sgsn policy sgsn_to_ggsn
  then permit application-services gprs-gtp-profile gtp1
```

Understanding GPRS Tunneling Protocol Inspection Object

To enable the SRX Series device to perform the inspection of GPRS Tunneling Protocol (GTP) traffic, you must create a GTP inspection object and then apply it to a policy. GTP inspection objects provide more flexibility in that they allow you to configure multiple policies that enforce different GTP configurations. You can configure the SRX Series device to control GTP traffic differently based on source and destination zones and addresses, action, and so on.

To configure GTP features, you must enter the context of a GTP configuration. To save your settings in the CLI, you must first exit the GTP configuration, then enter the **commit** command.

Creating a GTP Inspection Object

In this example, you create a GTP inspection object named LA-NY. You preserve most of the default values, but you enable the Sequence Number Validation feature.

CLI Configuration

```
user@host# set security gprs gtp profile la-ny
user@host# set security gprs gtp profile la-ny seq-number-validated
user@host# commit
```

Understanding GTP Message Filtering

When a SRX Series device receives a GTP packet, it checks the packet against policies configured on the device. If the packet matches a policy, the device inspects the packet according to the GTP configuration applied to the policy. If the packet fails to meet any of the GTP configuration parameters, the device will pass or drop the packets based on the configuration of the GTP inspection object.

This section describes features that constitute a GTP configuration, which the SRX Series device uses to perform GTP traffic inspection.

This topic includes:

- Message-Length Filtering on page 1001
- Setting GTP Message Lengths on page 1001
- CLI Configuration on page 1001
- Message-Type Filtering on page 1002
- Permitting and Denying Message Types on page 1002
- CLI Configuration on page 1002
- Supported Message Types on page 1002
- Message-Rate Limiting on page 1004
- Setting a Rate Limit on page 1005
- CLI Configuration on page 1005
- Sequence Number Validation on page 1005
- Enabling Sequence Number Validation on page 1005
- IP Fragmentation on page 1005

Message-Length Filtering

You can configure the SRX Series device to drop packets that do not meet your specified minimum or maximum message lengths. In the GTP header, the message length field indicates the length, in octets, of the GTP payload. It does not include the length of the GTP header itself, the UDP header, or the IP header. The default minimum and maximum GTP message lengths are 0 and 65535 bytes, respectively.

Setting GTP Message Lengths

In this example, you configure the minimum GTP message length to be 8 octets and the maximum GTP message length to be 1200 octets for the GPRS GTP inspection object.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 min-message-length 8
```

```
user@host# set security gprs gtp profile gtp1 max-message-length 1200
user@host# commit
```

Message-Type Filtering

You can configure the SRX Series device to filter GTP packets and permit or deny them based on their message type. By default, the SRX Series device permits all GTP message types.

A GTP message type includes one or many messages. When you permit or deny a message type, you automatically permit or deny all messages of the specified type. For example, if you select to drop the **sgsn-context** message type, you thereby drop **sgsn context request**, **sgsn context response**, and **sgsn context acknowledge** messages. For more information about message types, see “Supported Message Types” on page 1002.

You permit and deny message types based on the GTP version number. For example, you can deny message types for one version while you permit them for the other version.

Permitting and Denying Message Types

In this example, for the gtp1 profile, you configure the SRX Series device to drop the error-indication and failure-report message types for version 1.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 drop error-indication 1
user@host# set security gprs gtp profile gtp1 drop failure-report 1
user@host# commit
```

Supported Message Types

Table 117 on page 1002 lists the GPRS Tunneling Protocol (GTP) messages supported in GTP Releases 1997 and 1999 (including charging messages for GTP’) and the message types that you can use to configure GTP message-type filtering.

Table 117: GPRS Tunneling Protocol (GTP) Messages

Message	Message Type	Version 0	Version 1
create AA pdp context request	create-aa-pdp	b	
create AA pdp context response	create-aa-pdp	b	
create pdp context request	create-pdp	b	b
create pdp context response	create-pdp	b	b
Data record request	data-record	b	b

Table 117: GPRS Tunneling Protocol (GTP) Messages *(continued)*

Message	Message Type	Version 0	Version 1
Data record response	data-record	b	b
delete AA pdp context request	delete-aa-pdp	b	
delete AA pdp context response	delete-aa-pdp	b	
delete pdp context request	delete-pdp	b	b
delete pdp context response	delete-pdp	b	b
echo request	echo	b	b
echo response	echo	b	b
error indication	error-indication	b	b
failure report request	failure-report	b	b
failure report response	failure-report	b	b
forward relocation request	fwd-relocation	b	b
forward relocation response	fwd-relocation	b	b
forward relocation complete	fwd-relocation	b	b
forward relocation complete acknowledge	fwd-relocation	b	b
forward SRNS context	fwd-srns-context	b	b
forward SRNS context acknowledge	fwd-srns-context	b	b
identification request	identification	b	b
identification response	identification	b	b
node alive request	node-alive	b	b
node alive response	node-alive	b	b
note MS GPRS present request	note-ms-present	b	b
note MS GPRS present response	note-ms-present	b	b
pdu notification request	pdu-notification	b	b
pdu notification response	pdu-notification	b	b
pdu notification reject request	pdu-notification	b	b
pdu notification reject response	pdu-notification	b	b
RAN info relay	ran-info	b	b

Table 117: GPRS Tunneling Protocol (GTP) Messages *(continued)*

Message	Message Type	Version 0	Version 1
redirection request	redirection	b	b
redirection response	redirection	b	b
relocation cancel request	relocation-cancel	b	b
relocation cancel response	relocation-cancel	b	b
send route info request	send-route	b	b
send route info response	send-route	b	b
sgsn context request	sgsn-context	b	b
sgsn context response	sgsn-context	b	b
sgsn context acknowledge	sgsn-context	b	b
supported extension headers notification	supported-extension	b	b
g-pdu	gtp-pdu	b	b
update pdp context request	update-pdp	b	b
updated pdp context response	update-pdp	b	b
version not supported	version-not-supported	b	b

Message-Rate Limiting

You can configure the SRX Series device to limit the rate of network traffic going to a GSN. You can set separate thresholds, in packets per second, for GTP-Control (GTP-C) messages. Because GTP-C messages require processing and replies, they can potentially overwhelm a GSN. By setting a rate limit on GTP-C messages, you can protect your GSNs from possible denial of service (DoS) attacks such as the following:

- **Border gateway bandwidth saturation:** A malicious operator connected to the same GRX as your PLMN can direct so much network traffic at your Border Gateway that legitimate traffic is starved for bandwidth in or out of your PLMN, thus denying roaming access to or from your network.
- **GTP flood:** GTP traffic can flood a GSN, forcing it to spend its CPU cycles processing illegitimate data. This can prevent subscribers from roaming, forwarding data to external networks, and can prevent a GPRS from attaching to the network.

This feature limits the rate of traffic sent to each GSN from the Juniper Networks SRX Series device. The default rate is unlimited.

Setting a Rate Limit

In the following example, you limit the rate of incoming GTP-C messages to 300 packets per second.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 rate-limit 300
user@host# commit
```

Sequence Number Validation

You can configure a SRX Series device to perform sequence-number validation.

The header of a GTP packet contains a Sequence Number field. This number indicates to the GGSN receiving the GTP packets the order of the packets. During the Packet Data Protocol (PDP) context-activation stage, a sending GGSN uses zero (0) as the sequence-number value for the first G-PDU it sends through a tunnel to another GGSN. The sending GGSN increments the sequence number value for each following G-PDU it sends. The value resets to zero when it reaches 65535.

During the PDP context-activation stage, the receiving GGSN sets its counter to zero. Subsequently, whenever the receiving GGSN receives a valid G-PDU, the GGSN increments its counter by one. The counter resets to zero when it reaches 65535.

Normally, the receiving GGSN compares the sequence number in the packets it received with the sequence number from its counter. If the numbers correspond, the GGSN forwards the packet. If they differ, the GGSN drops the packet. By implementing a SRX Series device between the GGSNs, the device can perform this validation for the GGSN and drop packets that arrive out of sequence. This feature helps conserve GGSN resources by preventing the unnecessary processing of invalid packets.

Enabling Sequence Number Validation

In this example, you enable the Sequence Number Validation feature.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 seq-number-validated
user@host# commit
```

IP Fragmentation

A GTP packet consists of the message body and three headers: GTP, UDP and IP. If the resulting IP packet is larger than the message transmission unit (MTU) on the transferring link, the sending SGSN or GGSN performs an IP fragmentation.

By default, a SRX Series device buffers IP fragments until it receives a complete GTP message, and then inspects the GTP message.

Understanding GTP Information Elements

Information Elements (IEs) are included in all GTP control message packets. IEs provide information about GTP tunnels, such as creation, modification, deletion, and status. JUNOS Software supports IEs consistent with 3GPP Release 6. If you have contractual agreements with operators running earlier releases of 3GPP, you can reduce network overhead by restricting control messages containing unsupported IEs.

This section describes IEs contained in control messages you can configure the SRX Series device to screen based on IEs.

This topic includes:

- Access Point Name Filtering on page 1006
- IMSI Prefix Filtering on page 1007
- Supported R6 Information Elements on page 1008
- 3GPP R6 IE Removal on page 1011

Access Point Name Filtering

An Access Point Name (APN) is an IE included in the header of a GTP packet that provides information about how to reach a network. An APN comprises two elements:

- **Network ID**—Identifies the name of an external network such as “mobiphone.com”
- **Operator ID**—Which uniquely identifies the operators’ PLMN such as “mnc123.mcc456”

By default, the SRX Series device permits all APNs. However, you can configure the device to perform APN filtering to restrict roaming subscribers’ access to external networks.

To enable APN filtering, you must specify one or more APNs. To specify an APN, you need to know the domain name of the network (for example, mobiphone.com) and, optionally, the operator ID. Because the domain name (network ID) portion of an APN can potentially be very long and contain many characters, you can use the wildcard “*” as the first character of the APN. The wildcard indicates that the APN is not limited only to mobiphone.com but also includes all the characters that might precede it.

You may also set a *selection mode* for the APN. The selection mode indicates the origin of the APN and whether or not the Home Location Register (HLR) has verified the user subscription. You set the selection mode according to the security needs of your network. Possible selection modes include the following:

- **Mobile Station —**Mobile station-provided APN, subscription not verified

This selection mode indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.

- **Network**—Network-provided APN, subscription not verified

This selection mode indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.

- **Verified**—MS or network-provided APN, subscription verified

This selection mode indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.

APN filtering applies only to **create pdp request** messages. When performing APN filtering, the SRX Series device inspects GTP packets to—look for APNs that match APNs that you set. If the APN of a GTP packet matches an APN that you specified, the device then verifies the selection mode and only forwards the GTP packet if both the APN and the selection mode match the APN and the selection mode that you specified. Because APN filtering is based on perfect matches, using the wildcard "*" when setting an APN suffix can prevent the inadvertent exclusion of APNs that you would otherwise authorize. The device automatically denies all other APNs that do not match.

Additionally, an SRX Series device can filter GTP packets based on the combination of an International Mobile Subscriber Identity (IMSI) prefix and an APN.

Setting an APN and a Selection Mode

In this example, you set **mobiphone.com.mnc123.mcc456.gprs** as an APN and use the wildcard "*". You also set **Network** as the selection mode.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 apn
  *mobiphone.com.mnc123.mcc456.gprs mcc-mnc * action selection net
user@host# commit
```

IMSI Prefix Filtering

A GPRS Support Node (GSN) identifies a mobile station (MS) by its International Mobile Station Identity (IMSI). An IMSI comprises three elements: the Mobile Country Code (MCC), the Mobile Network Code (MNC), and the Mobile Subscriber Identification Number (MSIN). The MCC and MNC combined constitute the IMSI prefix and identify the mobile subscriber's home network, or Public Land Mobile Network (PLMN).

By setting IMSI prefixes, you can configure the SRX Series device to deny GTP traffic coming from nonroaming partners. By default, a device does not perform IMSI prefix filtering on GTP packets. By setting IMSI prefixes, you configure the device to filter

create pdp request messages and permit only GTP packets with IMSI prefixes that match the ones you set. The device allows GTP packets with IMSI prefixes that do not match any of the IMSI prefixes that you set. To block GTP packets with IMSI prefixes that do not match any of the IMSI prefixes set, use an explicit wildcard for the IMSI filter, and the action: **drop** should be the last IMSI prefix filtering policy.

When you filter GTP packets based on an IMSI prefix, you must also specify an APN. See "Setting a Combined IMSI Prefix and APN Filter" on page 1008

Setting a Combined IMSI Prefix and APN Filter

In this example, you set **mobiphone.com.mnc123.mcc456.gprs** as an APN and use the wildcard "*". You permit all selection modes for this APN. You also set the IMSI prefix for a known PLMN, which is 246565. The MCC-MNC pair can be five or six digits.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 apn
   *mobiphone.com.mnc123.mcc456.gprs mcc-mnc 246565 action pass
user@host# commit
```

Supported R6 Information Elements

JUNOS Software supports all 3GPP R6 IEs for GTP, as listed in Table 118 on page 1008.

Table 118: Supported Information Elements

IE Type Value	Information Element
1	Cause
2	International Mobile Subscriber Identity (IMSI)
3	Routing Area Identity (REI)
4	Temporary Logical Link Identity (TLLI)
5	Packet TMSI (P-TMSI)
8	Reordering Required
9	Authentication Triplet
11	MAP Cause
12	P-TMSI Signature
13	MS Validated
14	Recovery
15	Selection Mode

Table 118: Supported Information Elements *(continued)*

IE Type Value	Information Element
16	Tunnel Endpoint Identifier Data I
17	Tunnel Endpoint Identifier Control Plane
18	Tunnel Endpoint Identifier Data II
19	Teardown ID
20	NSAPI
21	RANAP Cause
22	RAB Context
23	Radio Priority SMS
24	Radio Priority
25	Packet Flow ID
26	Charging Characteristics
27	Trace Reference
28	Trace Type
29	MS Not Reachable Reason
127	Charging ID
128	End User Address
129	MM Context
130	PDP Context
131	Access Point Name
132	Protocol Configuration Options
133	GSN Address
134	MS International PSTN/ISDN Number (MSISDN)
135	Quality of Service Profile
136	Authentication Quintuplet
137	Traffic Flow Template
138	Target Identification
139	UTRAN Transparent Container

Table 118: Supported Information Elements *(continued)*

IE Type Value	Information Element
140	RAB Setup Information
141	Extension Header Type List
142	Trigger Id
143	OMC Identity
144	RAN Transparent Container
145	PDP Context Prioritization
146	Additional RAB Setup Information
147	SGSN Number
148	Common Flags
149	APN Restriction
150	Radio Priority LCS
151	RAT Type
152	User Location Information
153	MS Time Zone
154	IMEI-SV
155	CAMEL Charging Information Container
156	MBMS UE Context
157	Temporary Mobile Group Identity (TMGI)
158	RIM Routing Address
159	MBMS Protocol Configuration Options
160	MBMS Service Area
161	Source TNC PDCP context Information
162	Additional Trace Information
163	Hop Counter
164	Selected PLMN ID
165	MBMS Session Identifier
166	MBMS2G/3G Indicator

Table 118: Supported Information Elements *(continued)*

IE Type Value	Information Element
167	Enhanced NSAPI
168	MBMS Session Duration
169	Additional MBMS Trace Information
251	Charging Gateway Address
255	Private Extension

3GPP R6 IE Removal

The 3GPP R6 IE Removal feature allows you to retain interoperability in roaming between 2GPP and 3GPP networks. You can configure the GTP-aware SRX Series device, residing on the border of a PLMN and a GRX and acting as a Gp firewall, to remove 3GPP-specific attributes from the GTP packet header when the packet passes into a 2GPP network. You can configure the device to remove the RAT, RAI, ULI, IMEI-SV, and APN Restriction IEs from GTP messages prior to forwarding these messages to the GGSN.

R6 Removal

In this example, you configure the Gp interface of the security device to remove newly added R6 IEs (RAT, ULI, IMEI-SV and APN Restriction) from the GTP message.

CLI Configuration

```
user@host# set security gprs gtp profile gtp1
user@host# set security gprs gtp profile gtp1 remove-r6
user@host# commit
save
```

Understanding GGSN Redirection

SRX Series devices support GTP traffic GGSN redirection.

GGSN Redirection—A GGSN (X) can send create-pdp-context responses in which it can specify different GGSN IP addresses (GGSN Y and GGSN Z) for subsequent GTP-C and GTP-U messages. Consequently, the SGSN sends the subsequent GTP-C and GTP-U messages to GGSNs Y and Z, instead of X.

Part 13
Index

■ Index on page 1015

Index

Symbols

#, comments in configuration statements.....xlv
( ), in syntax descriptions.....xlv
3DES.....359
< >, in syntax descriptions.....xlv
[ ], in configuration statements.....xlv
{ }, in configuration statements.....xlv
| (pipe), in syntax descriptions.....xlv

.....188, 763, 769, 773, 779, 780, 782, 815
<Emphasis> *See* <Default Para Font> DoS
attributes.....769
loose source route.....782
network.....823
OS-specific.....847
SDP.....188

A

AAA.....306
Access Manager
adding firewall connection to.....439
auto-upgrading.....438
client-side files.....443
downloading to user's computers.....435
error messages.....446
launching from client.....439
logging.....446
overview.....433
system requirements.....442
Windows registry changes.....445
Access Point Name *See* APN
access profile
configuring
dynamic VPN.....449, 456, 473
NetScreen-Remote client.....485
accommodating end-to-end TCP communication
end-to-end TCP communication.....39
address sweep.....764
Advanced Encryption Standard (AES).....359
AES.....359
agentless access *See* UAC, Infranet Agent
agents, zombie.....807
aggressive mode.....361
AH (authentication header) protocol
overview.....355
ALGs
MS RPC.....299
SIP.....187
SIP NAT.....204
Sun RPC.....297
antispam filtering.....619
local list.....626
message handling.....636
server-based.....619
verifying.....637
antivirus
verifying.....679
antivirus, express.....685
EICAR file.....686
limitations.....686
testing.....686
updating antivirus patterns.....687
antivirus, full.....639
application protocol scanning.....647
content size limits.....642
decompression layer limit.....643
file extension scanning.....641
intelligent prescreening.....642
notification options.....676
scan session throttling.....645
scanning timeout.....643
signature database support.....642
updating antivirus patterns.....645
APN
filtering.....1006
selection mode.....1006
AppDDoS
application-level DDoS protection overview.....808
AppDDoS Protection
enabling example.....812
application binding.....545, 591
support on different device types.....589
application identification.....589
application binding.....591
configuring policies.....593
disable.....595
memory limit.....596
overview.....590

service binding....591
session limit....596
support on different device types....589
system cache....593
verifying cache statistics....598
verifying counters....599
See also IDP
application sets
IDP, configuring....535
overview....533
application system cache....593
overview....593
support on different device types....589
application-level DDoS....807
application-level DDoS protection overview....808
application-level DDoS protection configuration....812
applications
IDP, configuring....533
associating policy to schedulers....110
attack detection
overview....761
attack object groups....575
predefined....575
attack objects
custom....544
predefined....575
attacks
DOS....807, 815, 853
ICMP
floods....839, 840
fragments....792
IP packet fragments....802
Land....844, 845
large ICMP packets....794
Ping of Death....847
replay....363
session table floods....781, 815
SYN floods....823, 835
SYN fragments....804
Teardrop....849, 851
UDP floods....842, 843
unknown protocols....799
WinNuke....852, 853
auth users
groups....319
pass-through authentication....307
authentication
administrative....306
algorithms....359
client groups....319
configuring
external authentication servers....322, 327
SecurID server....327
pass-through....307
configuring....310
Quick Configuration....332
table....329
Web....308
configuring....315
authentication tables *See* UAC, authentication tables
authentication, authorization, and accounting servers....306
AutoKey IKE VPN....358
management....357

B

banners....330
braces, in configuration statements....xlv
brackets
angle, in syntax descriptions....xlv
square, in configuration statements....xlv

C

CA certificates....405
certificates....358
CA....405
loading....420
local....410
revocation....405
self-signed....408
UAC deployments *See* UAC, device authentication
changing session characteristics....7, 37
chassis cluster
ISSU upgrading....939
chassis clusters
about....867
control interfaces....903
creating a J Series cluster....906
creating an SRX Series cluster....904
deployment scenarios....942
disabling....942
enabling....895
fabric interfaces....903
formation....869
hardware setup for J Series devices....894
hardware setup for SRX Series devices....890
management interfaces on J Series devices....903
management interfaces on SRX Series devices....903
node interfaces on J Series devices....901
node interfaces on SRX Series devices....896
redundancy groups....870
setting node and cluster IDs....910
upgrading....939
verifying....932
verifying configuration....931
verifying interfaces....932
verifying redundancy group status....937

verifying statistics....933
verifying status....936
client groups for firewall authentication....319
cold sync
monitoring....878
comments, in configuration statements....xlv
compiling IDP policy....585
completing NetScreen-Remote client installation....480
compound attack sample....560
conditional route advertising configuration....929
Configuration
anto tunnel....375
IPSec manual tunnel VPN....391
VPN global settings....372
configuring
anomaly attack objects....565
application identification, memory limit....596
application identification, session limit....596
applications and application sets....136
chassis cluster information....919
conditional route advertising....929
control link recovery....924
control ports....909
dampening time between back-to-back redundancy group failovers....926
destination NAT....973
DSCP in IDP policy....568
dynamic VPN client configurations....472
dynamic VPN global settings....475
exempt rulebase....519, 528
external authentication servers....322, 327
fabric....920
firewall authentication....332
firewall on router....483
Gigabit Ethernet interface....63
host inbound traffic....53
protocols....58
system services....54
IDP application sets....535
IDP applications....533
IDP in security policy....536
IDP policy....522
IDP policy, application identification....593
IDP services....533
IKE gateway and peer authentication....385
IKE policy, authentication, and proposal....384
interface monitoring....925
interface source NAT for incoming SIP calls....216
interface source NAT pool for incoming SIP calls....218
interfaces....62
IPS rulebase....515, 525
IPsec AutoKey....390
options....468
Quick Configuration....465
IPsec policy....388
IPsec tunnel overview....369
log suppression....610
management interfaces....911
MGCP ALG....272
pass-through authentication....310
phase 2 proposals....387
policies....80
redundancy groups....921
redundant Ethernet interfaces....923
SCCP DoS attack protection....260
SecurID....327
security zones....51
signature attack objects....562
signature database automatic download....582
signature database manual download....578
signature database, Quick Configuration....580
SIP ALG....193
SIP ALG options....195
SIP DoS attack protection....199
SIP proxy
private zone....228
public zone....230
source NAT....982
static NAT for incoming SIP calls....223
TCP-reset parameter....60
terminal rules....541
three-zone SIP scenario....235
VPN global settings....370
Web authentication....315
Content Filtering....707
filter types....707
protocol support....708
verifying....722
context
router....44
secure....43
control link....882
failure and recovery....884
control link recovery
configuring....924
control plane
overview....881
control ports
configuring....909
controlling session termination....7, 38
conventions
notice icons....xliv
text and syntax....xliv
cookies, SYN....836
CoS features....6, 37
counters, verifying
for application identification....599
creating a J Series chassis cluster....906
creating a new connection for NetScreen-Remote client....487
creating an SRX Series chassis cluster....904

curly braces, in configuration statements.....xlv
custom attacks
application binding.....545
compound.....558
configuring.....562, 565
name.....545
overview.....544
protocol anomaly.....557
protocol binding.....549
service binding.....545
severity.....545
signature.....551
time binding.....550
customer support.....xlvi
contacting JTAC.....xlvi

D

data
fabric.....886
forwarding.....887
plane.....885
Data Encryption Standard (DES).....359
data path.....39
fast-path processing.....42
forward processing.....40
session-based processing.....40
data processing, stateful and stateless.....3, 33
DDoS.....807
application-level.....807
defining
exempt rulebase.....519
IPS rulebase.....515
defining IPsec protocols for NetScreen-Remote client.....491
DES.....359
Diffie-Hellman.....358
Diffserv
configuring in IDP policy.....568
digital signature.....403
disabling
chassis clusters.....942
disabling TCP packet security checks.....8, 38
displaying authentication table.....329
documentation
comments on.....xlvi
DoS
firewall.....821
session table floods.....781, 815
DoS attacks.....807
download
Access Manager.....435
client configuration, dynamic VPN.....435
signature database automatic.....582
signature database manually.....578
signature database overview.....577
dual control links
about.....882
connecting.....893
dynamic auth table provisioning *See* UAC, dynamic auth table provisioning
dynamic packet filtering.....761
dynamic VPNs
client configurations.....472
Quick Configuration.....469
configuration overview.....449
downloading client configurations.....435, 438
global settings.....475
Quick Configuration.....473
IKE gateways
Quick Configuration.....456
IKE policies
Quick Configuration.....454
IPsec AutoKeys
Quick Configuration.....465
IPsec policies
Quick Configuration.....463
overview.....433
Phase 1 proposals
Quick Configuration.....451
Phase 2 proposals
Quick Configuration.....460

E

enabling chassis clusters.....895
encryption algorithms.....359
encryption and hash algorithms.....495
ESP.....359
ESP (Encapsulating Security Payload) protocol
overview.....355
exempt rulebase
configuring.....519

F

fabric configuration.....920
fabric data link.....886
fabric data-link failure.....887
fabric interfaces.....903
fast-path processing.....42
filters, stateless firewall.....6, 36
FIN scans.....780
FIN without ACK flag attack detection
overview.....774
firewall screen options.....855
defined.....856
firewall users, pass-through
auth process.....308
floods
ICMP.....839, 840
session table.....815

SYN....823, 835, 836
UDP....842, 843
flow-based packet processing
defined....4, 34
flowd
monitoring....878
font conventions....xliv
forward processing....40
forwarding features....42

G

gatekeeper devices....163
Gi interface....996
glossary
IDP policy....502
Gp interface....996
gprs
about....995
tunneling protocol SRX Series devices....996
GTP
Access Point Name (APN) filtering....1006
IMSI prefix filtering....1007
inspection objects....998
IP fragmentation....1005
policy-based....998
GTP messages....1004
length, filtering by....1001
rate, limiting by....1004
type, filtering by....1002
types....1002
versions 0 and 1....1004

H

hardware
supported platforms....xlii
hardware setup, chassis cluster....890, 894
hash-based message authentication code....359
heartbeats....883
high availability....998
HMAC....359
Host Checker *See* UAC, Host Checker policy enforcement
hub-and-spoke....393

I

ICMP
floods....839, 840
fragments....792
large packets....794
ICMP header flags....556
IDP
application and services....533
application identification....589
application sets....533
application sets, configuring....535
custom attacks, overview....544
custom attacks, properties....551, 557, 558
deactivating rules....532
defining exempt rulebase....519
defining IPS rulebase....515
detector....560
DSCP....568
enabling IDP....536
exempt rulebase, Quick Configuration....528
inserting rule....530
IPS rulebase, Quick Configuration....525
log suppression....609
logging, overview....609
maximize-idp-sessions....607
performance and capacity tuning....607
policy....501
policy, manage....503
policy, overview....502
policy, Quick Configuration....522
protocol decoder....560
rulebase, application-level DDoS....504
rulebase, DDoS....504
rulebase, exempt....505
rulebase, IPS....505
rulebase, overview....503
rules, actions....510
rules, IP actions....512
rules, match conditions....507
rules, objects....508
rules, overview....506
send attack logs to the IC....611
setting terminal rules....541
signature database....571
signature database, Quick Configuration....580
terminal rules, overview....540
verify load status....585
verify policy compilation....585
verify signature database version....586
IDP policy
application identification....593
overview....502
rulebase, exempt....505
IKE....357
gateway
Quick Configuration....456
gateway and peer authentication....385
phase 1 proposals
predefined....360
Phase 1 proposals
configuration options....453
Quick Configuration....451
phase 2 proposals
configuring....387
predefined....362

Phase 2 proposals
options....462
Quick Configuration....460
policies
Quick Configuration....454
policy, authentication, and proposal
options....456
proxy IDs....362
IKE gateway configuration for NetScreen-Remote client....485
IMSI prefix filtering....1007
in-service upgrade
chassis cluster....939
Infranet Agent *See* UAC, Infranet Agent
Infranet Controller *See* UAC, Infranet Controller
Infranet Enforcer *See* UAC, JUNOS Enforcer
initiating manual redundancy group failover....927
inspections....761
installing Net-Screen-Remote client
from a network share drive....479
from CD-ROM....479
from Web site....479
interface monitoring configuration....925
interfaces....61
configuring....62, 63
control....903
fabric....903
Gigabit Ethernet interfaces, configuring....63
interfaces on J Series devices
management....903
node....901
interfaces on SRX Series devices
management....903
node....896
intrusion detection and prevention *See* IDP
IP options
incorrectly formatted....797
loose source route....770
record route....770, 771
security....770, 771
source route....782
stream ID....770, 771
strict source route....770
timestamp....770, 771
IP packet fragments....802
IP protocol header....554
IP spoofing....782, 787
IPS rulebase
configuring....515
IPsec
digital signature....403
manual key VPN
Configuration....391
overview....355
SAs....355, 360, 362
security protocols
Authentication Header (AH)....359
Encapsulating Security Protocol (ESP)....359
tunnel....355
creating through dynamic VPN feature....433
tunnel mode....364
tunnel negotiation....360
UAC support....337
IPsec policy
Quick Configuration....463
options....465

J

J-Web Cofiguration....320
J-Web Configuration....315, 321
JUEP *See* UAC, device authentication
JUNOS Enforcer *See* UAC, JUNOS Enforcer

L

L2TP....997
land attack detection
configuration....845
overview....844
local certificate....410
log suppression....609
configuring....610
logging
IDP, overview....609
logging in to NetScreen-Remote client....496
logging, traffic....998
loose source route IP detection
configuration....770

M

main mode....361
management interfaces....903
configuring....911
manual key management
overview....357
manuals
comments on....xlvi
MD5....359
Message Digest version 5 (MD5)....359
MGCP ALG....266
commands....269
entities....267
Quick Configuration....272
security....267
MGCP timeouts
inactivity....274
media inactivity....277
Mobile Station (MS) mode....1006

modes
aggressive....361
main....361
tunnel....364
modes, operational
NAT....998
route....998
transparent....998
modes, selection
APN....1006
Mobile Station (MS)....1006
network....1006
verified....1006
modular architecture....33
modulus....358
MPLS
context....43
MS RPC ALG, defined....299
multimedia sessions, SIP....187

N

NAT (Network Address Translation)....963
proxy ARP....991
rule-based....970
NAT mode....998
NetScreen-Remote client
creating a new connection....487
creating the preshared key....490
defining IPsec protocols....491
encryption and hash algorithms....495
login....496
system requirements....477
NetScreen-Remote client configuration
access profiles for XAuth....485
firewall on router....483
IKE gateway....485
PC or laptop....486
policies....486
security zone....483
tunnel interface....484
NetScreen-Remote client installation
completing....480
installing from CD-ROM....479
installing from network share drive....479
installing from Web site....479
PC or laptop....478
starting....478
network mode....1006
node interfaces on J Series devices....901
node interfaces on SRX Series devices....896
notice icons....xliv

O

Odyssey Access Client *See* UAC, Infranet Agent
operating system....33
operational modes
NAT....998
route....998
transparent....998

P

packet filtering....3, 33
packet processing....3, 33
stateful....3, 33
stateless....3, 33
packet-based processing....6, 36
parentheses, in syntax descriptions....xlv
pass-through authentication....307
Perfect Forward Secrecy *See* PFS
PFS....363
phase 1....360
proposals....360
proposals, predefined....360
phase 2....362
proposals....362
proposals, configuring....387
proposals, predefined....362
Phase 2
proposals, options....462
ping of death attack protection
configuration....848
overview....847
pinholes....190
PKI....405
using SCEP....411
policies....999
application services processing order....337
core section....781
quick configuration....76
schedulers
associating....110
quick configuration....106
shadowing....76
policies, configuring....998
policy
IDP *See* IDP
policy configuration for NetScreen-Remote client....486
policy templates
predefined....572
port scan attack protection
configuration....768
overview....766
predefined attack objects....575
predefined policy templates....572
overview....572
preshared key....358
preshared key creation for NetScreen-Remote client....490

probes
network....764
open ports....766
operating systems....773, 775
processing
data....3, 33
flow-based....4, 34
packet-based....6, 36
proposals
phase 1....360
phase 2....362
protocol anomaly....557
protocol anomaly attack....558
direction....557
expression (boolean expression)....559
member index....559
member index sample....559
order....558
reset....558
sample....557, 560
scope....558
test condition....557
protocol anomaly attack sample....557
protocol binding....549
sample format....550
proxy IDs....362
public/private key pair....407

Q

Quick Configuration
chassis cluster and redundancy groups....912
destination NAT....973
dymamic VPN client configurations....469
dymamic VPN global settings....473
exempt rulebase....528
firewall authentication....332
firewall screen options....855
IDP policy....522
IKE gateway....456
IKE polcies....454
IPS rulebase....525
IPsec AutoKey....465
IPsec Phase 1 proposal....451
IPsec Phase 2 proposal....460
IPsec policy....463
MGCP ALG....272
options....273
redundant Ethernet interfaces....915
SCCP ALG....255
options....256
security zones....51
signature database download....580
SIP ALG....193
source NAT....982
quick configuration
addresses and address sets....100
applications and application sets....136
policies....76
scheduler....106

R

rate limiting, GTP-C messages....1004
reconnaissance
address sweep....764
FIN scans....780
IP options....769
port scan....766
SYN and FIN flags set....773
TCP packet without flags....775
reconnaissance deterrence
IP address sweeps....764
blocking....764
overview....763
record route IP option....770, 771
redundancy group
configuring dampening time between back-to-back failovers....926
initiating manual failover....927
redundancy group configuration....921
redundancy groups
about....870
group 0....871
groups 1 through 128....871
interface monitoring....874
IP address monitoring....875
redundant Ethernet interfaces
configuring....923
Quick Configuration....915
understanding....879
registry changes, Access Manager....445
Remote Access Management Solution *See* dynamic VPNs
Remote Access Server
logging into for the first time....435
logging into subsequent sessions....438
overview....433
replay protection....363
resource access policies *See* UAC, resource access policies
RFCs
0792, Internet Control Message Protocol....148
1038, Revised IP Security Option....770
791, Internet Protocol....769, 770
793, Transmission Control Protocol....774
roles *See* UAC, user roles
route mode....998
RPC
Sun RPC....297

rulebase
exempt, attack objects....505
exempt, match condition....505
exempt, overview....505
IPS, action....505
IPS, attack objects....505
IPS, IP action....505
IPS, match condition....505
IPS, notification....505
IPS, overview....505
IPS, terminal flag....505
overview....503
rules....506
rules
actions....510
deactivating....532
inserting....530
IP actions....512
match conditions....507
objects....508
objects, address....508
objects, attack....509
objects, service....508
objects, zone....508
overview....506
terminal....540

S

SA parameters....360
SAs....362
SCCP
allowing unknown message types....258
configuring DoS attack protection....260
setting inactive media timeout....257
SCEP....411, 415
digital certificates....411
enrolling a local certificate....416
PKCS-10, PKCS-7....417
reenrolling certificates....421
RSA key....413
schedulers
configuration....106
screen
address sweep....764
bad IP options, drop....797
FIN with no ACK....785
FIN without ACK flag, drop....774
ICMP
fragments, block....792
ICMP floods....839, 840
IP options....769
IP packet fragments, block....802
IP spoofing....782
Land attacks....844, 845
large ICMP packets, block....794
loose source route IP option, detect....784
Ping of Death....847
port scan....766
source route IP option, deny....784
strict source route IP option, detect....784
SYN and FIN flags set....773
SYN floods....823, 835
SYN fragments, detect....804
SYN-ACK-ACK proxy floods....821
TCP packet without flags, detect....775
Teardrop....849, 851
UDP floods....842, 843
unknown protocols, drop....799
WinNuke attacks....852, 853
SCREEN
IP spoofing....787
secure and router contexts....42
Secure Hash Algorithm-1....359
SecurID....326
security checks, disabling TCP packet....8, 38
security IP option....770, 771
security policy
enabling IDP....536
security zone configuration for NetScreen-Remote client....483
security zones....47
creating....49
functional....48
Gigabit Ethernet interfaces....63
host inbound traffic....53
protocols....58
system services....54
interfaces....61
configuring....62
ports....62
options....53
Quick Configuration....51
TCP-reset parameter....60
selection modes
APN....1006
Mobile Station (MS)....1006
Network....1006
verified....1006
self-signed certificates
about....408
automatically generated....409, 427
manually generated....410, 428
sequence-number validation....1005
service binding....545, 591
support on different device types....589
services
IDP, configuring....533
timeout threshold....149
session
changing characteristics....7, 37
controlling termination....7, 38

session limits....816
destination-based....819
source-based....816, 818
session lookup....40
session table floods....781, 815
session-based processing....40
setting the node and cluster IDs....910
SHA-1....359
show security idp application-identification application-system-cache command....598
show security idp counters application-identification command....599
signature attack sample....556
signature custom attack....551
context....552
direction....553
ICMP header....556
IP protocol flags....554
pattern....553
protocol-specific parameters....553
sample....556
TCP header....555
UDP header....556
signature database....571
attack object groups....575
automatic update....582
manually update....578
overview....571
predefined attack objects....575
predefined policy templates....572
Quick Configuration....580
updating, overview....577
verify....584
verify load status....585
verify policy compilation....585
verify version....586
version, overview....584
See also IDP
SIP
connection information....189
defined....187
media announcements....189
messages....188
multimedia sessions....187
pinholes....188
request methods....191
response codes....212
RTCP....190
RTP....190
signaling....188
SIP ALG....191
call duration and timeouts....196
SIP NAT
call setup....204
defined....204
SIP timeouts
inactivity....196
media inactivity....199, 258
session inactivity....196
signaling inactivity....195, 197
SNMP failover traps....938
source IP route attack protection
overview....782
SPUs
monitoring....877
stateful....761
stateful and stateless data processing....3, 33
stateful inspection....761
stateful packet processing....3, 33
stateless firewall filters....6, 36
stateless packet processing....3, 33
statistics, verifying
for application identification....598
stream ID IP option....770, 771
strict source route IP option....770
Sun RPC ALG....297
call scenarios....297
defined....297
support, technical *See* technical support
SYN and FIN flags protection
overview....773
SYN checking....785
asymmetric routing....780
reconnaissance hole....781
session table floods....781
SYN cookies....836
SYN floods....823, 835
alarm threshold....827
attack threshold....826
destination threshold....828
source threshold....827
SYN cookies....836
threshold....824
timeout....828
SYN fragment protection
overview....804
SYN-ACK-ACK proxy floods....821
SYN-ACK-ACK-proxy flood protection
configuration....821
syntax conventions....xliv

T

TCP header flag attack protection
configuration....555
overview....775
teardrop attack protection
configuration....851
overview....849
technical support
contacting JTAC....xlvi

terminal rules
overview....540
setting....541
terminology
IDP policy....502
three-way handshakes....823
time binding....550
count....551
scope....551
timestamp IP option....770, 771
traffic
counting....998
logging....998
transparent mode....998
transport mode....364
Triple DES....359
tunnel interface configuration for NetScreen-Remote client....484
tunnel mode
overview....364

U

UAC
authentication tables
failover processing....338
overview....337
certificates *See* UAC, device authentication
clustering *See* UAC, failover processing
device authentication
configuring....342
overview....336
dynamic auth table provisioning....337
failover processing
configuring timeout actions....350
connecting to cluster....342
overview....338
Host Checker policy enforcement....339
Infranet Agent
agentless access....339
Odyssey Access Client....339
overview....335, 339
support information....339
Infranet Controller
communications with JUNOS Enforcer....336
configuring access to....342
overview....335
IPsec support....337
JUEP *See* UAC, device authentication
JUNOS Enforcer
communications with Infranet Controller....336
enabling....342
overview....335
logging....351
overview....335
policies
application services processing order....337
enforcement overview....337
resource access policies
failover processing....338
overview....337
show commands....337
test-only mode....351
timeout actions *See* UAC, failover processing
user roles....337
UDP header attack protection
configuration....556
Unified Access Control *See* UAC
Unified Threat Management
antispam filtering....619
antivirus protection, express....685
antivirus protection, full....639
content filtering....707
licensing....616
overview....615
platform support....616
web filtering....723
unknown protocol attack protection
overview....799
upgrading
chassis cluster ISSU....939
chassis clusters....939
user roles *See* UAC, user roles
UTM
WELF support for log files....756

V

verification
application system cache....598, 599
verified mode....1006
verifying
chassis cluster configuration....931
chassis cluster interfaces....932
chassis cluster redundancy group status....937
chassis cluster statistics....933
chassis cluster status....936
chassis clusters....932
IDP policy compilation....585
IDP policy load status....585
signature database....584
signature database version....586
version
application identification, support....589
signature database....584
VPNs
aggressive mode....361
AutoKey IKE....357
Diffie-Hellman exchange....358
Diffie-Hellman groups....358
dynamic VPN *See* dynamic VPNs

global settings.....370, 372
options.....372, 471, 474
main mode.....361
phase 1.....360
phase 2.....362
replay protection.....363

W

Web Filtering.....723
cache.....725
integrated.....724
local.....751
profiles.....726, 752
redirect.....739
verifying.....755
wildcards.....1006
Windows registry changes, Access Manager.....445
WinNuke attack protection
configuration.....853
overview.....852

Z

zombie agents.....807
zones
functional.....48
security.....47