**JUNOS® Software**

# Security Configuration Guide for J-series Services Routers and SRX-series Services Gateways

*Release 9.3*

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

Part Number: 530-025828-01, Revision 1

*JUNOS Software Security Configuration Guide*
Release 9.3


Revision History
October 2008—Revision 1

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

YEAR 2000 NOTICE

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

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 Web site 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

# Table of Contents

## Part 3 Intrusion Detection and Prevention Features

### Chapter 18 IDP Policies 639

# About This Guide

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



## Objectives

This guide describes how to use and configure key security features on J-series Services Routers running JUNOS software with enhanced services and SRX-series services gateways running JUNOS software. It provides conceptual information, suggested workflows, and examples where applicable.

**NOTE:** This manual documents Release 9.3 of JUNOS software. For additional information—either corrections to or information that might have been omitted from this manual—see the *JUNOS Software with Enhanced Services Release Notes* or *JUNOS Software for SRX-series Services Gateways Release Notes* at http://www.juniper.net.

## Audience

This manual is designed for anyone who installs, sets up, configures, monitors, or administers a J-series Services Router running JUNOS software with enhanced services 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

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

## Supported Routing Platforms

This manual describes features supported on J-series Services Routers running JUNOS software with enhanced services 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 with enhanced services for J-series Services Routers
- JUNOS software for SRX-series services gateways

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

For an annotated list of the documentation referred to in Table 1 on page xxxvi, see "List of Technical Publications" on page xl. All documents are available at http://www.juniper.net/techpubs/.

**Table 1: Tasks and Related Documentation**

| Task | Related Documentation |
|---|---|
| **Basic Device Installation and Setup** | |
| ■ Reviewing safety warnings and compliance statements<br>■ Installing hardware and establishing basic connectivity<br>■ Initially setting up a device | J-series Services Routers:<br>■ *JUNOS Software with Enhanced Services Quick Start*<br>■ *JUNOS Software with Enhanced Services Hardware Guide*<br>■ *JUNOS Software with Enhanced Services Release Notes*<br><br>SRX-series services gateways: the appropriate *Services Gateway Getting Started Guide* |
| **Migration from ScreenOS or JUNOS Software to JUNOS Software with Enhanced Services (if necessary)** | |
| ■ Migrating from JUNOS Release 8.3 or later to JUNOS software with enhanced services<br>■ Migrating from ScreenOS Release 5.4 or later JUNOS software with enhanced services | *JUNOS Software with Enhanced Services 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** | |

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

| Task | Related Documentation |
|---|---|
| Configuring device interfaces | ■ *JUNOS Software Interfaces and Routing Configuration Guide*<br>■ *JUNOS Software CLI Reference* |
| **Deployment Planning and Configuration** | |
| ■ Understanding and gathering information required to design network firewalls and IPsec VPNs<br>■ Implementing a JUNOS software with enhanced services firewall from a sample scenario<br>■ Implementing a policy-based IPsec VPN from a sample scenario | *JUNOS Software with Enhanced Services Design and Implementation Guide* (J-series Services Routers only) |
| **Security Configuration** | |
| Configuring and managing the following security services:<br>■ Stateful firewall policies<br>■ Zones and their interfaces and address books<br>■ IPsec VPNs<br>■ Firewall screens<br>■ Interface modes: Network Address Translation (NAT) mode and Router mode<br>■ Public Key Cryptography (PKI)<br>■ Application Layer Gateways (ALGs)<br>■ Chassis clusters<br>■ Intrusion Detection and Prevention (IDP) | ■ *JUNOS Software Security Configuration Guide*<br>■ *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<br>■ Configuring class-of-service (CoS) features, including traffic shaping and policing<br>■ Configuring packet-based stateless firewall filters (access control lists) to control access and limit traffic rates<br>■ Configuring MPLS to control network traffic patterns | ■ *JUNOS Software Interfaces and Routing Configuration Guide*<br>■ *JUNOS Software CLI Reference* |
| **WAN Acceleration Module Installation (Optional)** | |
| 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** | |

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

| Task | Related Documentation |
|---|---|
| ■ Administering user authentication and access<br>■ Monitoring the device, routing protocols, and routing operations<br>■ Configuring and monitoring system alarms and events, real-time performance (RPM) probes, and performance<br>■ Monitoring the firewall and other security-related services<br>■ Managing system log files<br>■ Upgrading software<br>■ Diagnosing common problems | *JUNOS Software Administration Guide* |
| **User Interfaces** | |
| ■ Understanding and using the J-Web interface<br>■ Understanding and using the CLI configuration editor | ■ *JUNOS Software with Enhanced Services Quick Start* (J-series Services Routers only)<br>■ *JUNOS Software Administration Guide* |

## Document Conventions

Table 2 on page xxxviii 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 xxxviii 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 configuration mode, type the **configure** command:<br><br>user@host> **configure** |

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

| Convention | Description | Examples |
|---|---|---|
| `Fixed-width text like this` | Represents output that appears on the terminal screen. | `user@host> show chassis alarms`<br>`No alarms currently active` |
| *Italic text like this* | ■ Introduces important new terms.<br>■ Identifies book names.<br>■ Identifies RFC and Internet draft titles. | ■ A policy *term* is a named structure that defines match conditions and actions.<br>■ *JUNOS System Basics Configuration Guide*<br>■ 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:<br>[edit]<br>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.<br>■ 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<br>(*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]<br>routing-options {<br>  static {<br>  route default {<br>  nexthop *address*;<br>  retain;<br>  }<br>  }<br>} |
| ; (semicolon) | Identifies a leaf statement at a configuration hierarchy level. | |
| **J-Web GUI Conventions** | | |
| **Bold text like this** | Represents J-Web graphical user interface (GUI) items you click or select. | ■ In the Logical Interfaces box, select **All Interfaces**.<br>■ To cancel the configuration, click **Cancel**. |

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

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

## List of Technical Publications

The following sections list hardware and software guides and release notes for SRX-series services gateways and J-series Services Routers running JUNOS software.

All documents are available at http://www.juniper.net/techpubs/.

**Hardware Guides**

- *SRX 5600 Services Gateway Hardware Guide*—Describes hardware components, installation, basic configuration, and basic troubleshooting procedures for the SRX 5600 services gateway. This guide explains how to prepare a site, unpack and install the device, replace device hardware, establish basic connectivity, and perform routine maintenance.
- *SRX 5800 Services Gateway Hardware Guide*—Describes hardware components, installation, basic configuration, and basic troubleshooting procedures for the SRX 5800 services gateway. This guide explains how to prepare a site, unpack and install the device, replace device hardware, establish basic connectivity, and perform routine maintenance.
- *JUNOS Software with Enhanced Services Quick Start*—Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity.
- *JUNOS Software with Enhanced Services Hardware Guide*—Provides an overview, basic instructions, and specifications for J-series Services Routers. This guide explains how to prepare a site, unpack and install the router, replace router hardware, and establish basic router connectivity. This guide contains hardware descriptions and specifications.

**Software Guides**

- *JUNOS Software Interfaces and Routing Configuration Guide*—Explains how to configure SRX-series and J-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification.
- *JUNOS Software Security Configuration Guide*—Explains how to configure and manage SRX-series and J-series security services such as stateful firewall policies, IPsec VPNs, firewall screens, Network Address Translation (NAT), Public Key Cryptography, chassis clusters, Application Layer Gateways (ALGs), and Intrusion Detection and Prevention (IDP).
- *JUNOS Software Administration Guide*—Shows how to monitor SRX-series and J-series devices and routing operations, firewall and security services, system alarms and events, and network performance. This guide also shows how to administer user authentication and access, upgrade software, and diagnose common problems.
- *JUNOS Software CLI Reference*—Provides the complete configuration hierarchy available on SRX-series and J-series devices. This guide also describes the

configuration statements and operational mode commands unique to these devices.

- *JUNOS Network Management Configuration Guide*—Describes enterprise-specific MIBs for JUNOS software. The information in this guide is applicable to M-series, T-series, EX-series, SRX-series, and J-series devices.
- *JUNOS System Log Messages Reference*—Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message. The information in this guide is applicable to M-series, T-series, EX-series, SRX-series, and J-series devices.
- *JUNOS Software with Enhanced Services Design and Implementation Guide*—Provides guidelines and examples for designing and implementing IPsec VPNs, firewalls, and routing on J-series Services Routers running JUNOS software with enhanced services.
- *JUNOS Software with Enhanced Services Migration Guide*—Provides instructions for migrating an SSG device running ScreenOS software or a J-series Services Router running the JUNOS software to JUNOS software with enhanced services.
- *WXC Integrated Services Module Installation and Configuration Guide*—Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration.

**Release Notes**

- *JUNOS Software for SRX-series Services Gateways Release Notes*—Summarizes new features and known problems for SRX-series services gateways and the JUNOS software running on those devices. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions.
- *JUNOS Software with Enhanced Services Release Notes*—Summarizes new features and known problems for J-series Services Routers and the JUNOS software running on those routers. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions.

## Documentation Feedback

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

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

## Requesting Technical Support

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

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

### Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

### Opening a Case with JTAC

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

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

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

# Part 1
# Support Overview for Security Features

Chapter 1

# Support for Security Features on SRX 5600 and SRX 5800 Services Gateways

The following tables list security features that are supported on SRX 5600 and SRX 5800 services gateways.

**Table 4: Support Information: Zones**

| Feature | More Information |
|---|---|
| Security zone | "Security Zone" on page 50 |
| Functional zone | "Functional Zone" on page 50 |
| For information about the interfaces that are supported on your device, see the *JUNOS Software Interfaces and Routing Configuration Guide*. | |

**Table 5: Support Information: Security Policy**

| Feature | More Information |
|---|---|
| Address books | "Configuring Address Books" on page 98 |
| Policy application sets | "Policy Application Sets Overview" on page 112 |
| Schedulers | "Configuring Schedulers" on page 103 |
| Policy applications | "Understanding Internet-Related Predefined Policy Applications" on page 117 |
| Internet Control Message Protocol (ICMP) predefined policy application | "Understanding the ICMP Predefined Policy Application" on page 113 |
| Internet-related predefined policy applications | "Understanding Internet-Related Predefined Policy Applications" on page 117 |
| Microsoft predefined policy applications | "Understanding Microsoft Predefined Policy Applications" on page 119 |
| Dynamic routing protocols predefined policy applications | "Understanding Dynamic Routing Protocols Predefined Policy Applications" on page 122 |

**Table 5: Support Information: Security Policy** *(continued)*

| Feature | More Information |
|---|---|
| Streaming video predefined policy applications | "Understanding Streaming Video Predefined Policy Applications" on page 123 |
| Sun remote procedure protocol (RPC) predefined policy applications | "Understanding Sun RPC Predefined Policy Applications" on page 124 |
| Security and tunnel predefined policy applications | "Understanding Security and Tunnel Predefined Policy Applications" on page 125 |
| IP-related predefined policy applications | "Understanding IP-Related Predefined Policy Applications" on page 126 |
| Instant messaging predefined policy applications | "Understanding Instant Messaging Predefined Policy Applications" on page 126 |
| Management predefined policy applications | "Understanding Management Predefined Policy Applications" on page 127 |
| Mail predefined policy applications | "Understanding Mail Predefined Policy Applications" on page 129 |
| UNIX predefined policy applications | "Understanding UNIX Predefined Policy Applications" on page 129 |
| Miscellaneous predefined policy applications | "Understanding Miscellaneous Predefined Policy Applications" on page 130 |
| Custom policy Applications | "Understanding Custom Policy Applications" on page 131 |
| Policy application timeouts | "Understanding Policy Application Timeouts" on page 141 |
| Policy verification | "Understanding Policy Ordering" on page 74 |

**Table 6: Support Information: Firewall Authentication**

| Feature | More Information |
|---|---|
| Web authentication | "Web Authentication" on page 150 |
| Pass-through authentication | "Pass-Through Authentication" on page 149 |
| Local authentication server | "Firewall User Authentication Overview" on page 147 |
| RADIUS authentication server | "Firewall User Authentication Overview" on page 147 |
| LDAP authentication server | "Firewall User Authentication Overview" on page 147 |
| SecurID authentication server | "Understanding SecurID User Authentication" on page 168 |

**Table 7: Support Information: Attack Detection and Prevention**

| Feature | More Information |
|---|---|
| Bad IP option | "Understanding Bad IP Option Protection" on page 214 |
| Block fragment traffic | "Blocking Fragmented ICMP Packets" on page 210 |
| FIN flag without ACK flag set protection | "Blocking Packets with FIN Flag/No ACK Flag Set" on page 195 |
| ICMP flood protection | "Understanding ICMP Flood Attacks" on page 249 |
| ICMP fragment protection | "Understanding ICMP Fragment Protection" on page 209 |
| Large size ICMP packet protection | "Understanding Large ICMP Packet Protection" on page 211 |
| Loose source route option | "Blocking Packets with Either a Loose or Strict Source Route Option Set" on page 206 |
| IP record route option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| IP security option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| IP address spoof | "Blocking IP Spoofing" on page 204 |
| IP stream option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| IP strict source route option | "Blocking Packets with Either a Loose or Strict Source Route Option Set" on page 206 |
| IP address sweep | "Understanding IP Address Sweeps" on page 182 |
| IP timestamp option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| Land attack protection | "Understanding Land Attacks" on page 254 |
| Ping of death attack protection | "Understanding Ping of Death Attacks" on page 257 |
| Port scan | "Understanding Port Scanning" on page 184 |
| Source IP based session limit | "Understanding Session Table Flood Attacks" on page 225 |
| SYN-ACK-ACK proxy protection | "Understanding SYN-ACK-ACK Proxy Flood Attacks" on page 230 |
| SYN and FIN flags set protection | "Blocking Packets with SYN and FIN Flags Set" on page 194 |
| SYN flood protection | "Understanding SYN Flood Attacks" on page 233 |
| SYN fragment protection | "Understanding SYN Fragment Protection" on page 221 |
| Teardrop attack protection | "Understanding Teardrop Attacks" on page 259 |
| TCP packet without flag set protection | "Blocking Packets with No Flags Set" on page 196 |

**Table 7: Support Information: Attack Detection and Prevention** *(continued)*

| Feature | More Information |
|---|---|
| Unknown protocol protection | "Understanding Unknown Protocol Protection" on page 216 |
| UDP flood protection | "Understanding UDP Flood Attacks" on page 252 |
| WinNuke attack protection | "Understanding WinNuke Attacks" on page 262 |

**Table 8: Support Information: Network Address Translation**

| Feature | More Information |
|---|---|
| Destination IP address translation | "Destination IP Address Translation Overview" on page 279 |
| Static Network Address Translation (NAT) | "Understanding Static NAT on SRX-series Services Gateways" on page 281 |
| Rule-based NAT | "Understanding Rule-Based Destination NAT on SRX-series Services Gateways" on page 285 |
| Source IP address translation | "Source IP Address Translation Overview" on page 291 |
| NAT interface source pools | "Understanding NAT Interface Source Pools" on page 292 |
| Configuring proxy Address Resolution Protocol (ARP) | "Configuring Proxy ARP on SRX-series Services Gateways" on page 301 |

**Table 9: Support Information: Chassis Cluster**

| Feature | More Information |
|---|---|
| Chassis cluster formation | "Understanding Chassis Cluster Formation" on page 316 |
| Active/passive chassis cluster (that is, no cross-box data forwarding over the fabric interface) | "Understanding Chassis Cluster Formation" on page 316 |
| Redundancy group 0 (backup for Routing Engine) | "Redundancy Group 0: Routing Engines" on page 318 |
| Redundancy groups 1 | "Redundancy Groups 1 Through 255" on page 319 |
| Redundant Ethernet interfaces | "Understanding Redundant Ethernet Interfaces" on page 322 |
| Control plane failover | "Understanding the Control Plane" on page 323 |
| Data plane failover | "Understanding the Data Plane" on page 326 |
| All JUNOS flow-based routing functionality (except for IPsec VPN) | *JUNOS Software Interfaces and Routing Configuration Guide* |

**Table 10: Support Information: IPsec**

| Feature | More Information |
|---|---|
| Policy-based and route-based VPNs | "Virtual Private Networks (VPNs)" on page 378 |
| Tunnel mode | "Understanding IPsec Operational Modes" on page 380 |
| Authentication Header (AH) protocol | "Understanding IPsec Security Protocols" on page 382 |
| Encapsulating Security Payload (ESP) protocol | "Understanding IPsec Security Protocols" on page 382 |
| IKE phase 1 | "Understanding IPsec Tunnel Negotiation" on page 391 |
| IKE phase 2 | "Understanding IPsec Tunnel Negotiation" on page 391 |
| Manual key management | "Understanding IPsec Key Management" on page 385 |
| Autokey management | "Understanding IPsec Key Management" on page 385 |
| Antireplay (packet replay attack prevention) | "Replay Protection" on page 395 |
| Dead peer detection (DPD) | "Configuring an IKE Gateway and Peer Authentication" on page 410 |
| XAuth extended authentication for remote access connections | "Configuring an Access Profile for XAuth" on page 623 |
| VPN monitoring | "Configuring VPN Global Settings" on page 395 |

**Table 11: Support Information: PKI**

| Feature | More Information |
|---|---|
| Internet Key Exchange (IKE) support | "Internet Key Exchange" on page 442 |
| Entrust, Microsoft, and Verisign certificate authorities (CAs) | "Understanding Certificates" on page 440 |
| Automated certificate enrollment using Simple Certificate Enrollment Protocol (SCEP) | "Using Digital Certificates" on page 448 |
| Automatic generation of self-signed certificates | "Understanding Self-Signed Certificates" on page 446 |
| Distinguished Encoding Rules (DER), Privacy-Enhanced Mail (PEM), Public-Key Cryptography Standard 7 (PKCS7), and X509 certificate encoding | "Manually Loading a CRL onto the Device" on page 461 |
| Manual installation of DER-encoded and PEM-encoded CRLs | "Manually Loading a CRL onto the Device" on page 461 |
| Online certificate revocation list (CRL) retrieval through LDAP and HTTP | "PKI Management and Implementation" on page 445 |
| CRL update at user-specified interval | "Understanding Certificate Revocation Lists" on page 443 |

**Table 12: Support Information: ALGs**

| Feature | More Information |
|---|---|
| FTP Application Layer Gateway (ALG) | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| Trivial File Transfer Protocol (TFTP) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| H.323 ALG | "Understanding the H.323 ALG" on page 476 |
| Media Gateway Control Protocol (MGCP) ALG | "Understanding the MGCP ALG" on page 578 |
| Point-to-Point Tunneling Protocol (PPTP) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| REAL ALG | Table 72 on page 474 |
| Remote procedure call (RPC) ALG | "Understanding the RPC ALG" on page 608 |
| Remote shell (RSH) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| Real-Time Streaming Protocol (RTSP) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| Skinny Call Control Protocol (SCCP) ALG | "Understanding the SCCP ALG" on page 561 |
| Session Initiation Protocol (SIP) ALG | "Understanding the SIP ALG" on page 499 |
| Structured Query Language (SQL) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| TALK ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |

**Table 13: Support Information: IDP Policy**

| Feature | More Information |
|---|---|
| Intrusion Detection and Prevention (IDP) Policy | "IDP Policies Overview" on page 640 |
| Intrusion prevention system (IPS) rulebase | "Defining Rules for an IPS Rulebase" on page 652 |
| Exempt rulebase | "Defining Rules for an Exempt Rulebase" on page 656 |
| Custom attacks | "Understanding Custom Attack Objects" on page 681 |
| Differentiated Services code point (DSCP) marking | "Configuring DSCP in an IDP Policy" on page 702 |

**Table 14: Support Information: IDP Signature Database**

| Feature | More Information |
|---|---|
| IDP signature database | "Understanding the IDP Signature Database" on page 705 |
| Predefined policy templates | "Using Predefined Policy Templates" on page 706 |
| Signature database—manual download | "Updating the Signature Database Manually" on page 712 |
| Signature database—automatic download | "Updating the Signature Database Automatically" on page 716 |
| Signature database version | "Understanding the Signature Database Version" on page 717 |

**Table 15: Support Information: IDP Application Identification**

| Feature | More Information |
|---|---|
| Application identification | "Understanding Application Identification" on page 721 |
| Service and application bindings | "Understanding Service and Application Bindings" on page 722 |
| Application system cache | "Understanding Application System Cache" on page 724 |

**Table 16: Support Information: IDP Logging**

| Feature | More Information |
|---|---|
| IDP logging | ■ Understanding IDP Logging on page 739<br>■ *JUNOS Software Administration Guide* |

**Table 17: Support Information: IDP SSL Inspectioon**

| Feature | More Information |
|---|---|
| IDP SSL Inspection | "IDP SSL Inspection" on page 733 |

Chapter 2

# Support for Security Features on J-series Services Routers

The following tables list security features that are supported on J-series Services Routers.

**Table 18: Support Information: Zones**

| Feature | More Information |
|---|---|
| Security zone | "Security Zone" on page 50 |
| Functional zone | "Functional Zone" on page 50 |
| For information about the interfaces that are supported on your device, see the *JUNOS Software Interfaces and Routing Configuration Guide*. | |

**Table 19: Support Information: Security Policy**

| Feature | More Information |
|---|---|
| Address books | "Configuring Address Books" on page 98 |
| Policy application sets | "Policy Application Sets Overview" on page 112 |
| Schedulers | "Configuring Schedulers" on page 103 |
| Policy applications | "Understanding Internet-Related Predefined Policy Applications" on page 117 |
| Internet Control Message Protocol (ICMP) predefined policy application | "Understanding the ICMP Predefined Policy Application" on page 113 |
| Internet-related predefined policy applications | "Understanding Internet-Related Predefined Policy Applications" on page 117 |
| Microsoft predefined policy applications | "Understanding Microsoft Predefined Policy Applications" on page 119 |
| Dynamic routing protocols predefined policy applications | "Understanding Dynamic Routing Protocols Predefined Policy Applications" on page 122 |

**Table 19: Support Information: Security Policy** *(continued)*

| Feature | More Information |
|---|---|
| Streaming video predefined policy applications | "Understanding Streaming Video Predefined Policy Applications" on page 123 |
| Sun remote procedure protocol (RPC) predefined policy applications | "Understanding Sun RPC Predefined Policy Applications" on page 124 |
| Security and tunnel predefined policy applications | "Understanding Security and Tunnel Predefined Policy Applications" on page 125 |
| IP-related predefined policy applications | "Understanding IP-Related Predefined Policy Applications" on page 126 |
| Instant messaging predefined policy applications | "Understanding Instant Messaging Predefined Policy Applications" on page 126 |
| Management predefined policy applications | "Understanding Management Predefined Policy Applications" on page 127 |
| Mail predefined policy applications | "Understanding Mail Predefined Policy Applications" on page 129 |
| UNIX predefined policy applications | "Understanding UNIX Predefined Policy Applications" on page 129 |
| Miscellaneous predefined policy applications | "Understanding Miscellaneous Predefined Policy Applications" on page 130 |
| Custom policy Applications | "Understanding Custom Policy Applications" on page 131 |
| Policy application timeouts | "Understanding Policy Application Timeouts" on page 141 |
| Policy verification | "Understanding Policy Ordering" on page 74 |

**Table 20: Support Information: Firewall Authentication**

| Feature | More Information |
|---|---|
| Web authentication | "Web Authentication" on page 150 |
| Pass-through authentication | "Pass-Through Authentication" on page 149 |
| Local authentication server | "Firewall User Authentication Overview" on page 147 |
| RADIUS authentication server | "Firewall User Authentication Overview" on page 147 |
| LDAP authentication server | "Firewall User Authentication Overview" on page 147 |
| SecurID authentication server | "Understanding SecurID User Authentication" on page 168 |

**Table 21: Support Information: Attack Detection and Prevention**

| Feature | More Information |
|---|---|
| Bad IP option | "Understanding Bad IP Option Protection" on page 214 |
| Block fragment traffic | "Blocking Fragmented ICMP Packets" on page 210 |
| FIN flag without ACK flag set protection | "Blocking Packets with FIN Flag/No ACK Flag Set" on page 195 |
| ICMP flood protection | "Understanding ICMP Flood Attacks" on page 249 |
| ICMP fragment protection | "Understanding ICMP Fragment Protection" on page 209 |
| Large size ICMP packet protection | "Understanding Large ICMP Packet Protection" on page 211 |
| Loose source route option | "Blocking Packets with Either a Loose or Strict Source Route Option Set" on page 206 |
| IP record route option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| IP security option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| IP address spoof | "Blocking IP Spoofing" on page 204 |
| IP stream option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| IP strict source route option | "Blocking Packets with Either a Loose or Strict Source Route Option Set" on page 206 |
| IP address sweep | "Understanding IP Address Sweeps" on page 182 |
| IP timestamp option | "SCREEN Options for Detecting IP Options Used For Reconnaissance" on page 189 |
| Land attack protection | "Understanding Land Attacks" on page 254 |
| Ping of death attack protection | "Understanding Ping of Death Attacks" on page 257 |
| Port scan | "Understanding Port Scanning" on page 184 |
| Source IP based session limit | "Understanding Session Table Flood Attacks" on page 225 |
| SYN-ACK-ACK proxy protection | "Understanding SYN-ACK-ACK Proxy Flood Attacks" on page 230 |
| SYN and FIN flags set protection | "Blocking Packets with SYN and FIN Flags Set" on page 194 |
| SYN flood protection | "Understanding SYN Flood Attacks" on page 233 |
| SYN fragment protection | "Understanding SYN Fragment Protection" on page 221 |
| Teardrop attack protection | "Understanding Teardrop Attacks" on page 259 |
| TCP packet without flag set protection | "Blocking Packets with No Flags Set" on page 196 |

**Table 21: Support Information: Attack Detection and Prevention** *(continued)*

| Feature | More Information |
|---|---|
| Unknown protocol protection | "Understanding Unknown Protocol Protection" on page 216 |
| UDP flood protection | "Understanding UDP Flood Attacks" on page 252 |
| WinNuke attack protection | "Understanding WinNuke Attacks" on page 262 |

**Table 22: Support Information: Network Address Translation**

| Feature | More Information |
|---|---|
| Destination IP address translation | "Destination IP Address Translation Overview" on page 279 |
| Static Network Address Translation (NAT) | "Example: Configuring Static NAT on SRX-series Services Gateways" on page 282 |
| Policy-based NAT | "Understanding NAT-Dst Policy-Based NAT on J-series Services Routers" on page 283 |
| Source IP address translation | "Source IP Address Translation Overview" on page 291 |
| NAT interface source pools | "Understanding NAT Interface Source Pools" on page 292 |

**Table 23: Support Information: Chassis Cluster**

| Feature | More Information |
|---|---|
| Chassis cluster formation | "Understanding Chassis Cluster Formation" on page 316 |
| Active/active chassis cluster (that is, cross-box data forwarding over the fabric interface) | "Understanding Chassis Cluster Formation" on page 316 |
| Redundancy group 0 (backup for Routing Engine) | "Redundancy Group 0: Routing Engines" on page 318 |
| Redundancy groups 1 through 255 | "Redundancy Groups 1 Through 255" on page 319 |
| Redundant Ethernet interfaces | "Understanding Redundant Ethernet Interfaces" on page 322 |
| Control plane failover | "Understanding the Control Plane" on page 323 |
| Data plane failover | "Understanding the Data Plane" on page 326 |
| All JUNOS flow-based routing functionality | *JUNOS Software Interfaces and Routing Configuration Guide* |

**Table 24: Support Information: IPsec**

| Feature | More Information |
|---|---|
| Policy-based and route-based VPNs | "Virtual Private Networks (VPNs)" on page 378 |

**Table 24: Support Information: IPsec** *(continued)*

| Feature | More Information |
|---|---|
| Tunnel mode | "Understanding IPsec Operational Modes" on page 380 |
| Authentication Header (AH) protocol | "Understanding IPsec Security Protocols" on page 382 |
| Encapsulating Security Payload (ESP) protocol | "Understanding IPsec Security Protocols" on page 382 |
| IKE phase 1 | "Understanding IPsec Tunnel Negotiation" on page 391 |
| IKE phase 2 | "Understanding IPsec Tunnel Negotiation" on page 391 |
| Manual key management | "Understanding IPsec Key Management" on page 385 |
| Autokey management | "Understanding IPsec Key Management" on page 385 |
| Antireplay (packet replay attack prevention) | "Replay Protection" on page 395 |
| Dead peer detection (DPD) | "Configuring an IKE Gateway and Peer Authentication" on page 410 |

**Table 25: Support Information: PKI**

| Feature | More Information |
|---|---|
| Internet Key Exchange (IKE) support | "Internet Key Exchange" on page 442 |
| Entrust, Microsoft, and Verisign certificate authorities (CAs) | "Understanding Certificates" on page 440 |
| Automated certificate enrollment using Simple Certificate Enrollment Protocol (SCEP) | "Using Digital Certificates" on page 448 |
| Automatic generation of self-signed certificates | "Understanding Self-Signed Certificates" on page 446 |
| Distinguished Encoding Rules (DER), Privacy-Enhanced Mail (PEM), Public-Key Cryptography Standard 7 (PKCS7), and X509 certificate encoding | "Manually Loading a CRL onto the Device" on page 461 |
| Manual installation of DER-encoded and PEM-encoded CRLs | "Manually Loading a CRL onto the Device" on page 461 |
| Online certificate revocation list (CRL) retrieval through LDAP and HTTP | "PKI Management and Implementation" on page 445 |
| CRL update at user-specified interval | "Understanding Certificate Revocation Lists" on page 443 |

**Table 26: Support Information: ALGs**

| Feature | More Information |
|---|---|
| FTP Application Layer Gateway (ALG) | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |

**Table 26: Support Information: ALGs** *(continued)*

| Feature | More Information |
|---|---|
| Trivial File Transfer Protocol (TFTP) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| H.323 ALG | "Understanding the H.323 ALG" on page 476 |
| Media Gateway Control Protocol (MGCP) ALG | "Understanding the MGCP ALG" on page 578 |
| Point-to-Point Tunneling Protocol (PPTP) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| REAL ALG | Table 72 on page 474 |
| Remote procedure call (RPC) ALG | "Understanding the RPC ALG" on page 608 |
| Remote shell (RSH) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| Real-Time Streaming Protocol (RTSP) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| Skinny Call Control Protocol (SCCP) ALG | "Understanding the SCCP ALG" on page 561 |
| Session Initiation Protocol (SIP) ALG | "Understanding the SIP ALG" on page 499 |
| Structured Query Language (SQL) ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |
| TALK ALG | "Configuring Application Layer Gateways—Quick Configuration" on page 473 |

**Table 27: Support Information: Netscreen Remote**

| Feature | More Information |
|---|---|
| Netscreen Remote VPN client | "NetScreen-Remote VPN Client" on page 615 |

## Part 2
# Security Features

Chapter 3

# 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 1 on page 20. 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.

This section includes:



## 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 1 on page 20 shows an architectural overview of traffic flow in a Services Router running JUNOS software. See Figure 3 on page 26 to follow the path of the traffic as it traverses through the Flow services module.

**Figure 1: Traffic Flow for Flow-Based Processing**

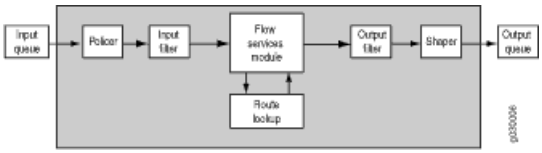


A flow is a stream of related packets that meet the same matching criteria and share the same characteristics. JUNOS software treats packets belonging to the same flow in the same manner.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, 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

## *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 2 on page 22 shows architectural overview of traffic flow in a standard JUNOS router.

**Figure 2: Traffic Flow for Packet-Based Processing**

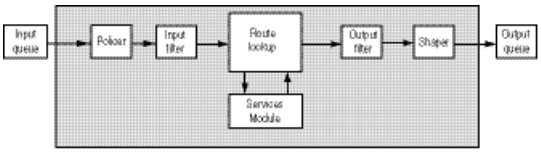


Filters and CoS features are typically associated with one or more interfaces to influence which packets are allowed to transit the system and to apply special actions to packets as necessary.

---

**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

The 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 the JUNOS software performs on the packet.

- Part 1—Forwarding Processing on page 26
- Part 2—Session-Based Processing on page 26
- Part 3—Forwarding Features on page 28

Figure 3 on page 26 shows the path of a data packet as it traverses through the Services Router. Refer to Figure 1 on page 20 to see how the flow module in Figure 3 on page 26 fits in with the architecture of the software.

**Figure 3: Data Packet Traversing the Flow Module on the Services Router**

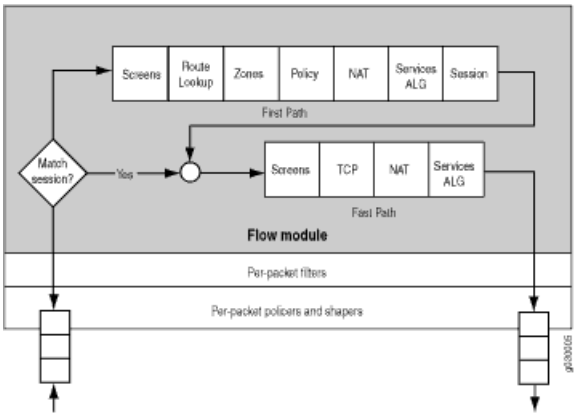


## *Part 1—Forwarding Processing*

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, the 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 27.)

If the packet matches a session, fast-path processing is performed. (See “Fast-Path Processing” on page 28.)

### 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 23, 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 179.

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 49.

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 275.

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

   For ALG details, see the "Application Layer Gateways (ALGs)" on page 471.

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 275.
4. Forwarding features are applied. (See the following section "Part 3—Forwarding Features" on page 28.)

### *Part 3—Forwarding Features*

After the packet has passed through session-based processing, the 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 28: 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 29 |
| Router context factory configuration | Yes | No | "Router Context" on page 30 |

## 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*.

Chapter 4

# 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 5600 and 5800 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. 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.

This section includes:



## Overview of SRX-series Services Gateways Running JUNOS Software

The SRX 5600 and 5800 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 34.) 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 4 on page 32 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 4: First-Packet Processing**

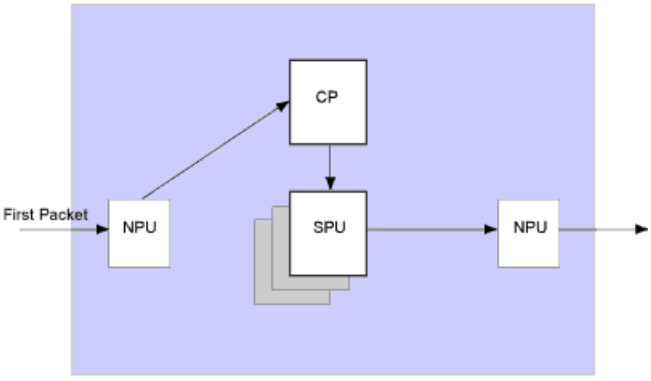


For details on session creation for the first packet in a flow, see “Understanding Session Creation: First-Packet Processing” on page 40.

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 5 on page 33 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 5: Fast-Path Processing**

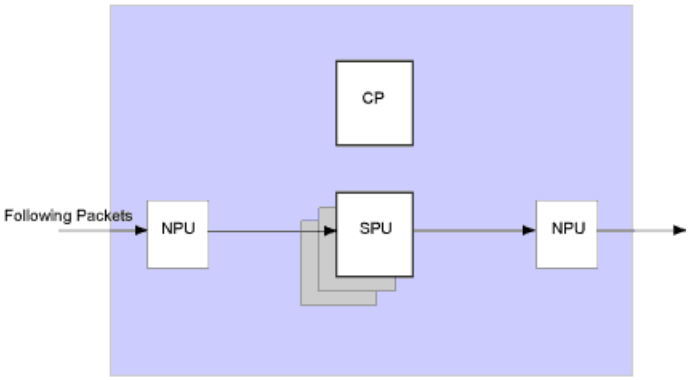


## Overview of Stateful and Stateless Data Processing

Traffic that enters and exits a 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 a services gateway 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.

- 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 6 on page 34 shows a conceptual view of how flow-based traffic processing occurs on an SPU of an SRX 5600 or SRX 5800 services gateway.

**Figure 6: Traffic Flow for Flow-Based Processing**

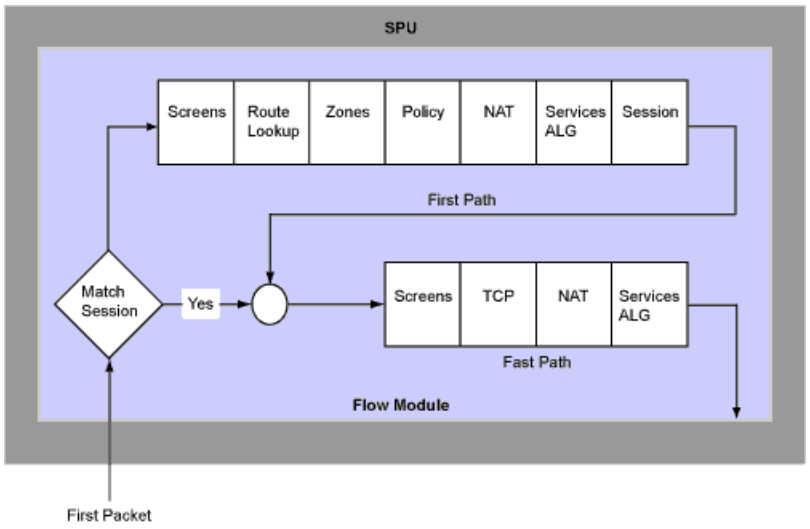


A flow is a stream of related packets that meet the same matching criteria and share the same characteristics. JUNOS software treats packets belonging to the same flow in the same manner.

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 services gateway, 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 services gateway 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
  ```

## Understanding Sessions

JUNOS software for SRX-series services gateways is a distributed parallel processing high throughput and high performance system. This section explains how a session is created and the process a packet undergoes as it transits the services gateway.

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 SRX 5600 and SRX 5800 services gateways:

- **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 SRX 5600 and SRX 5800 services gateways 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 services gateway 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 services gateway undertakes in establishing a session for packets belonging to a flow that transits the SRX-series services gateways.

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*

To determine if a packet belongs to an existing flow, the services gateway 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 a services gateway 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 43.)

**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 43.

Figure 7 on page 43 illustrates the first part of the process the first packet of a flow undergoes after it reaches the services gateway. 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 7: Session Creation: First-Packet Processing**

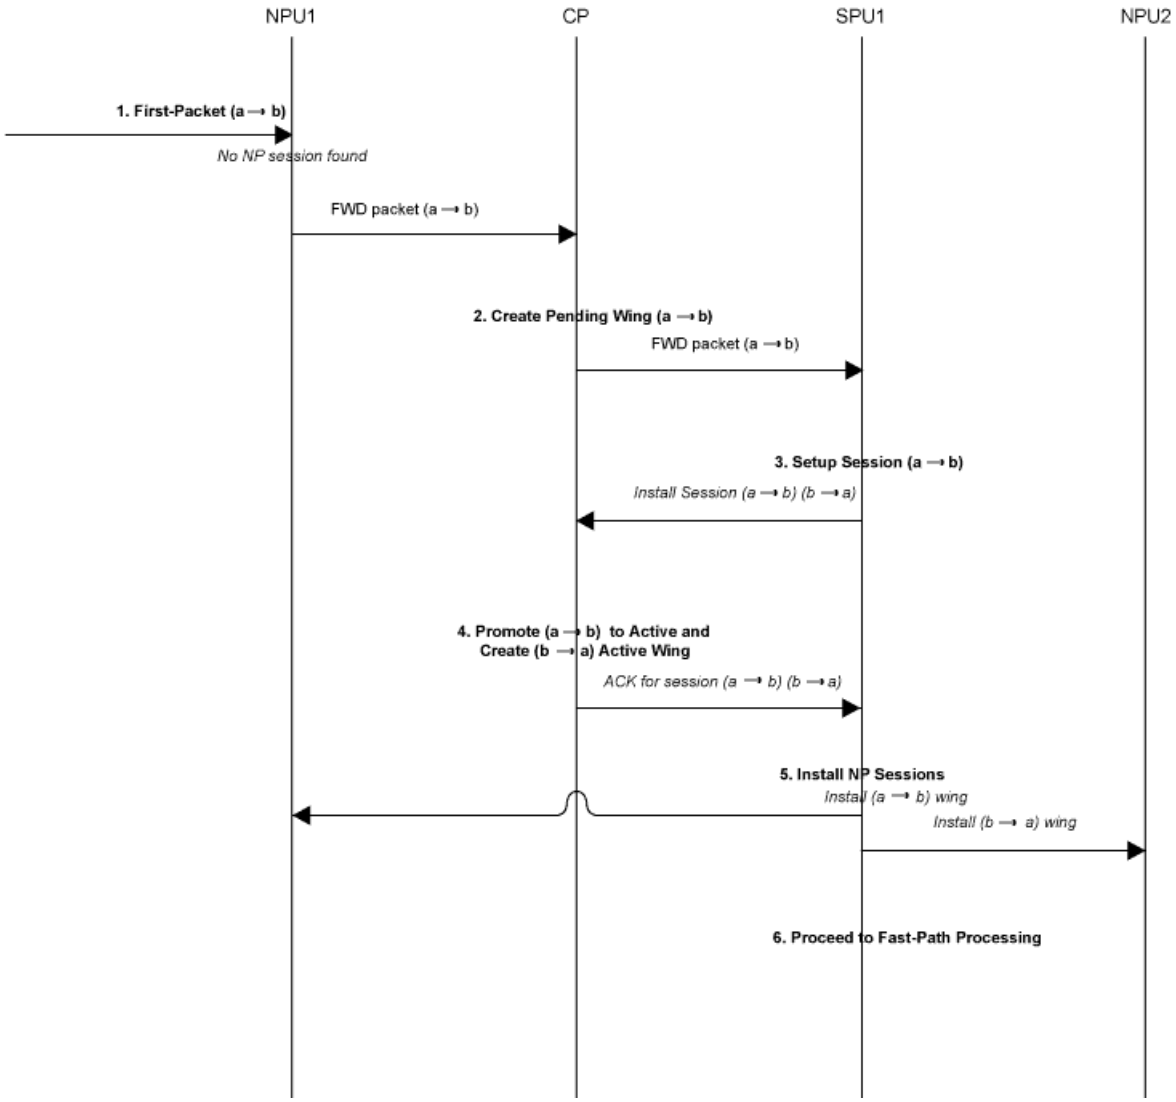


## *Understanding Fast-Path Processing*

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 services gateway's IOC.

1. The packet arrives at the services gateway'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 35.)

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 35.)

### 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 8 on page 46 illustrates the process a packet undergoes when it reaches the services gateway and a session exists for the flow that the packet belongs to.

**Figure 8: "Packet Walk" for Fast Path Processing**

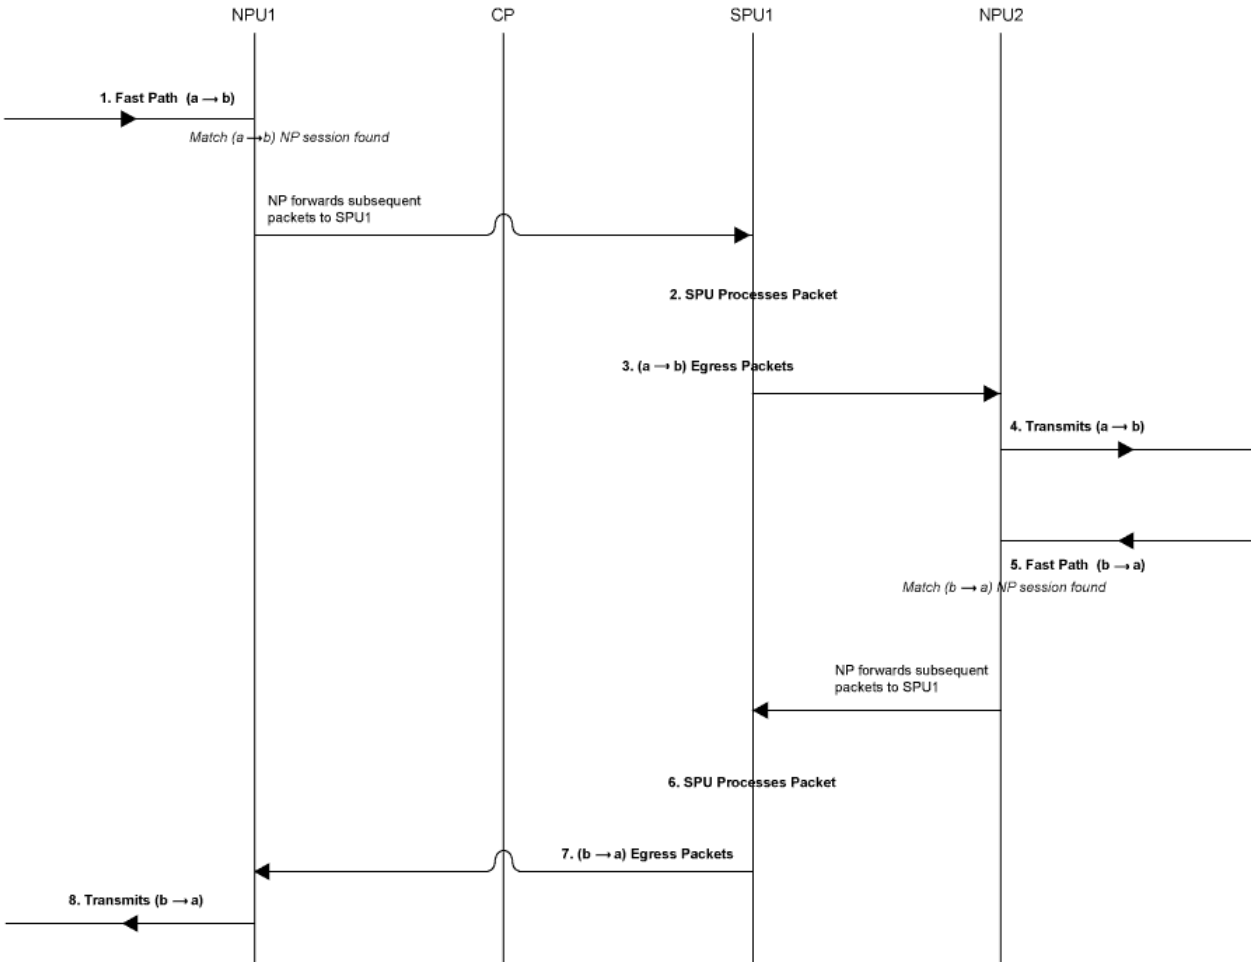


## *Obtaining Information About Sessions By Using the Configuration show Command*

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 36.
- syn-flood-protection-mode—Displays the SYN Proxy mode.

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

## *Obtaining Information About Sessions By Using the Operational show Command*

You can obtain information about the sessions and packet flows active on your device, including detailed information about specific sessions. (The services gateway 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.

### 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 services gateway 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 services gateway, 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 services gateway 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
```

### *Using the Operational clear Command to Terminate 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 5

# 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*.

This section includes:



## 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

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 181.

- 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 98.
- 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.

**NOTE:** JUNOS software supports only Layer 3 interfaces.

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 51
- Configuring Security Zones—Quick Configuration on page 53
- Configuring the TCP-Reset Parameter on page 62

## 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 49
- Configuring Host Inbound Traffic on page 55
- Configuring the TCP-Reset Parameter on page 62

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, use either J-Web or the CLI configuration editor.

This topic covers:



## J-Web Configuration

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

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 49
- Configuring Security Zones—Quick Configuration on page 53

## Configuring Security Zones—Quick Configuration

You can use J-Web Quick Configuration to quickly configure security zones. See Figure 9 on page 54.

**Before You Begin**

For background information, read:

- Understanding Security Zones on page 49
- Creating Security Zones on page 51

Figure 9 on page 54 shows the Quick Configuration Zones page.

**Figure 9: Quick Configuration Zones Page**

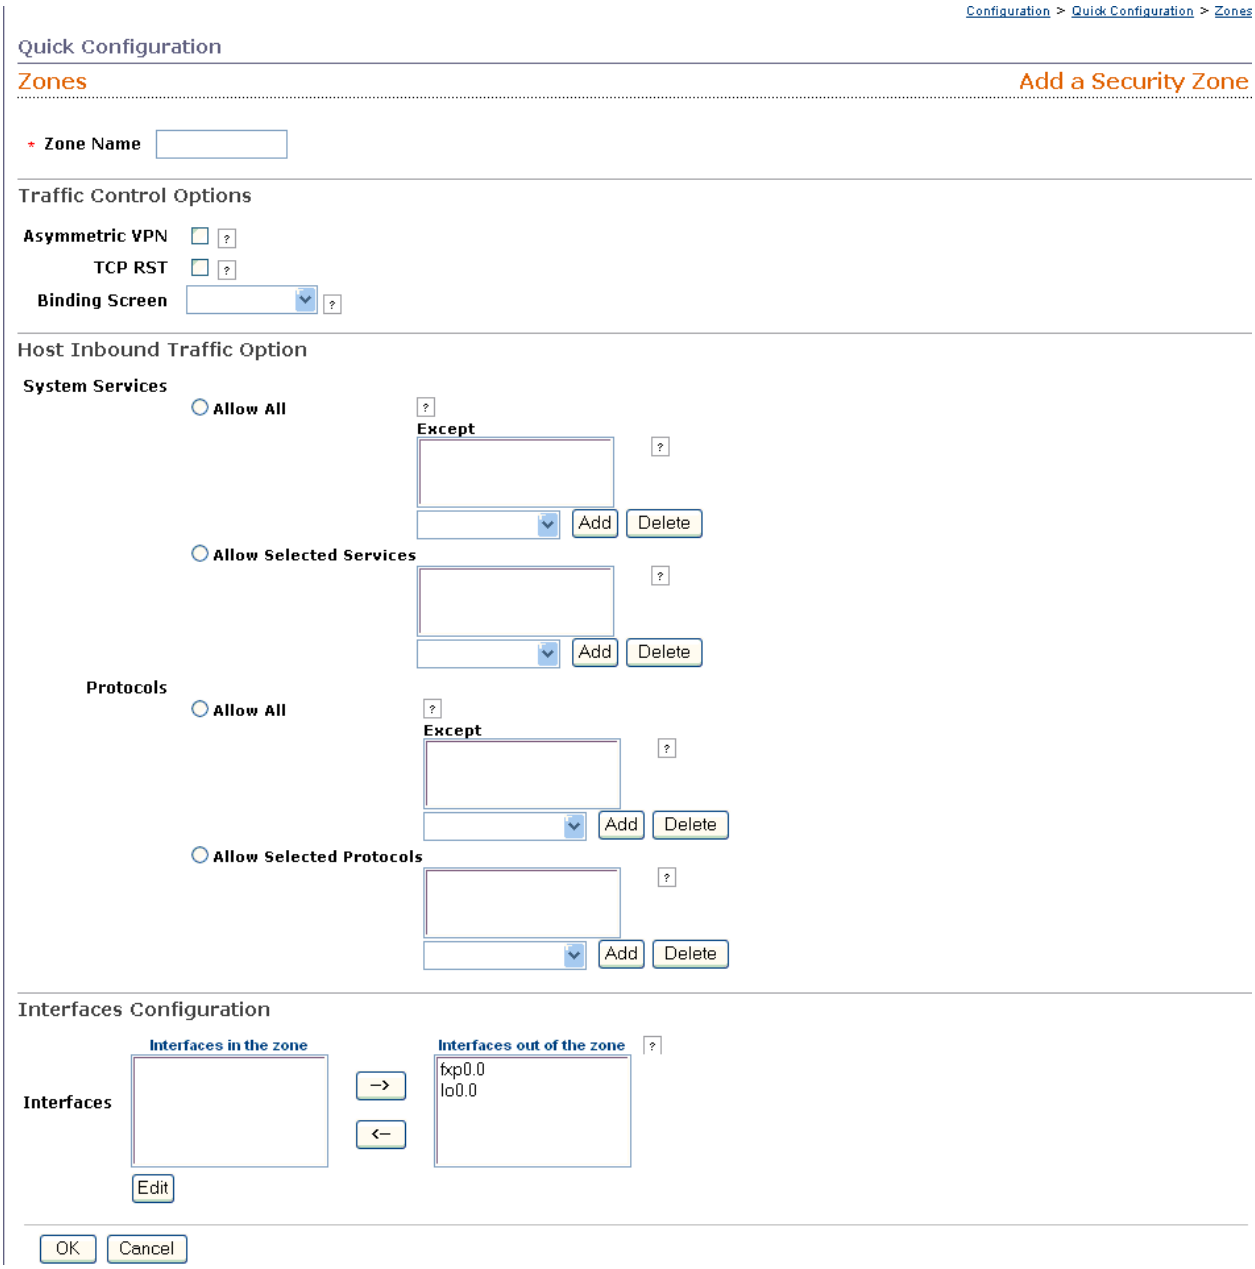


To configure security zones with Quick Configuration:

1. Select Configuration > Quick Configuration > Zones.
2. Click Add to create new zones.
3. Fill the form as shown in Table 29 on page 55.
4. 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 29: 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.<br><br>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.<br><br>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 49. |

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*

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 30 on page 56 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 30: 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 30 on page 56 (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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 all 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 49
- Configuring the TCP-Reset Parameter on page 62
- Creating Security Zones on page 51

## 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 31 on page 60 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 31: Supported Inbound System Protocols**

| Supported System Services | | | |
|---|---|---|---|
| all | igmp | pgm | sap |
| bfd | ldp | pim | vrrp |
| bgp | msdp | rip | |
| dlsw | 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

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 Configuration > View and Edit> Edit Configuration.

   The Configuration page appears.
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 and 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 49
- Creating Security Zones on page 51
- Configuring Host Inbound Traffic on page 55

## 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 49. |

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*

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

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 49
- Configuring Host Inbound Traffic on page 55

## 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 49. |

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 51
- Configuring Interfaces—Quick Configuration on page 64

## Configuring Interfaces—Quick Configuration

You can use J-Web Quick Configuration to quickly configure interfaces. See Figure 10 on page 65.

**Before You Begin**

For background information, read:

- Understanding Security Zones on page 49
- Creating Security Zones on page 51
- Understanding Security Zone Interfaces on page 63

A list of the network interfaces available on the routing platform appears, as shown in Figure 10 on page 65. The third column indicates whether the interface has been configured.

**Figure 10: Quick Configuration Interfaces Page**

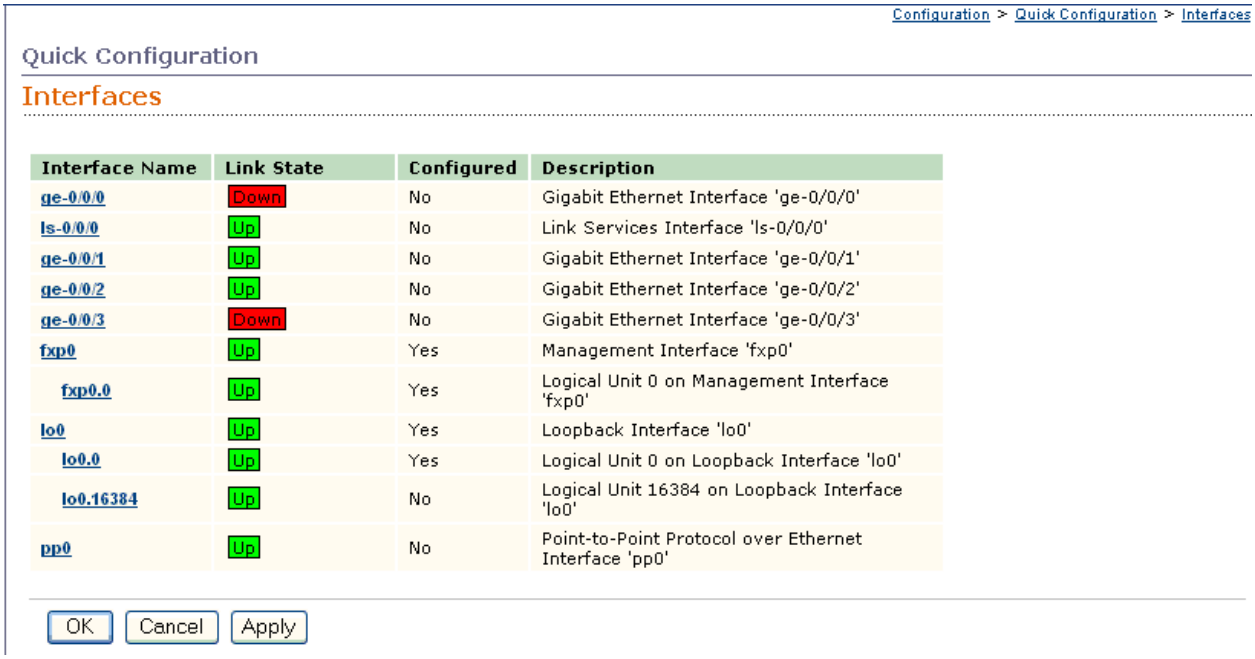


To configure a network interface with Quick Configuration:

1. Select Configuration > Quick Configuration > 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 65.

## 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 66 shows the Gigabit Ethernet Interface Quick Configuration page.

**Figure 11: Gigabit Ethernet Interface Quick Configuration**

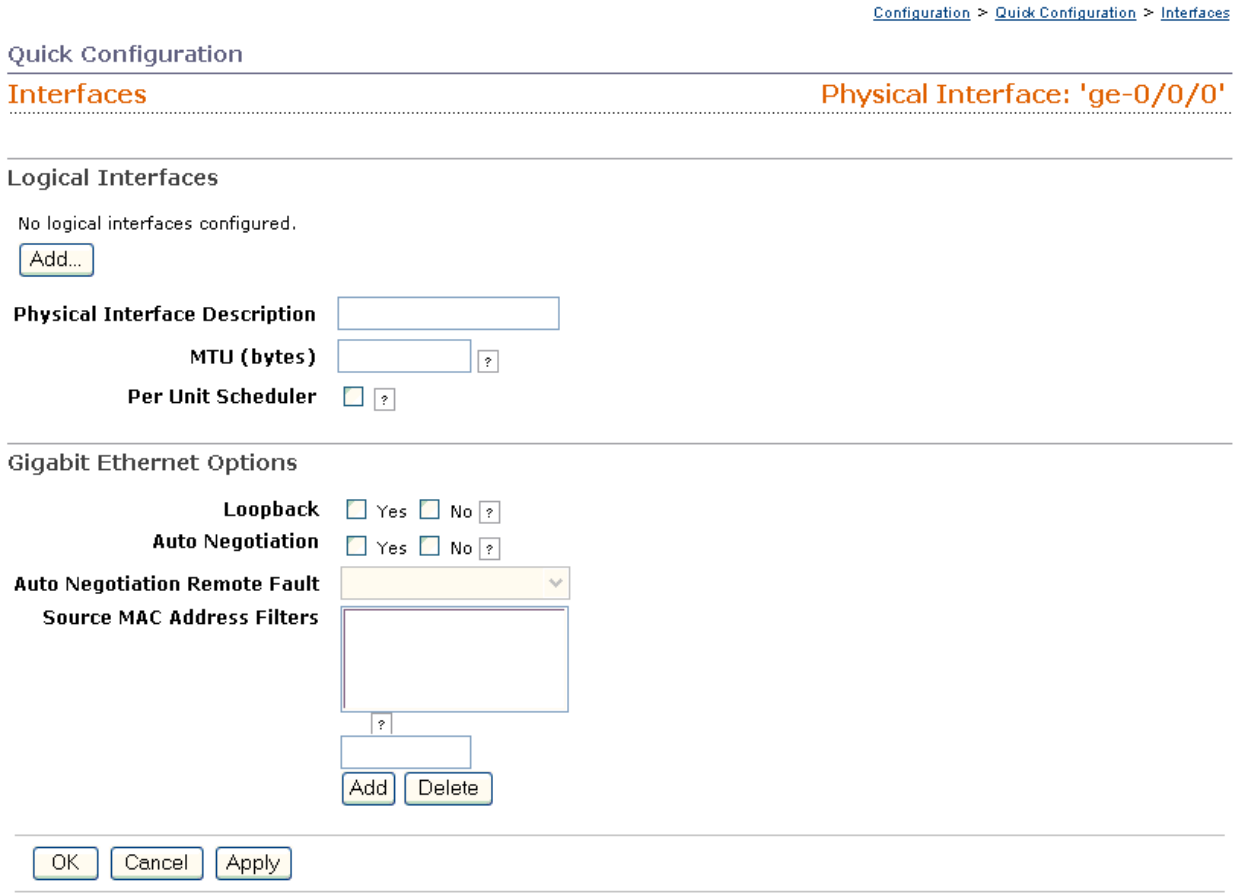


1. Select Configuration > Quick Configuration > Interfaces. The properties you can configure on a Gigabit Ethernet interface appear, as shown in Figure 11 on page 66.
2. Fill in the information as described in Table 32 on page 67.
3. Click one of the following buttons:
   - To apply the configuration 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 32: Gigabit Ethernet Quick Configuration Page Summary**

| Field | Function | Actions |
|---|---|---|
| **Logical Interfaces** | | |
| 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:<br>10.10.10.0/24.<br>2. Click **Add**.<br>3. Click **OK**.<br>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.<br><br>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.<br><br>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,<br>**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.<br>■ 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. |

**Table 32: Gigabit Ethernet Quick Configuration Page Summary** *(continued)*

| Field | Function | Actions |
|---|---|---|
| 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. |
| Per Unit Scheduler | Enables scheduling on logical interfaces.<br><br>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.<br>■ 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.<br><br>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.<br><br>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. |

Chapter 6

# 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.

This section includes:



## 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 70.

**Figure 12: Default Policy**

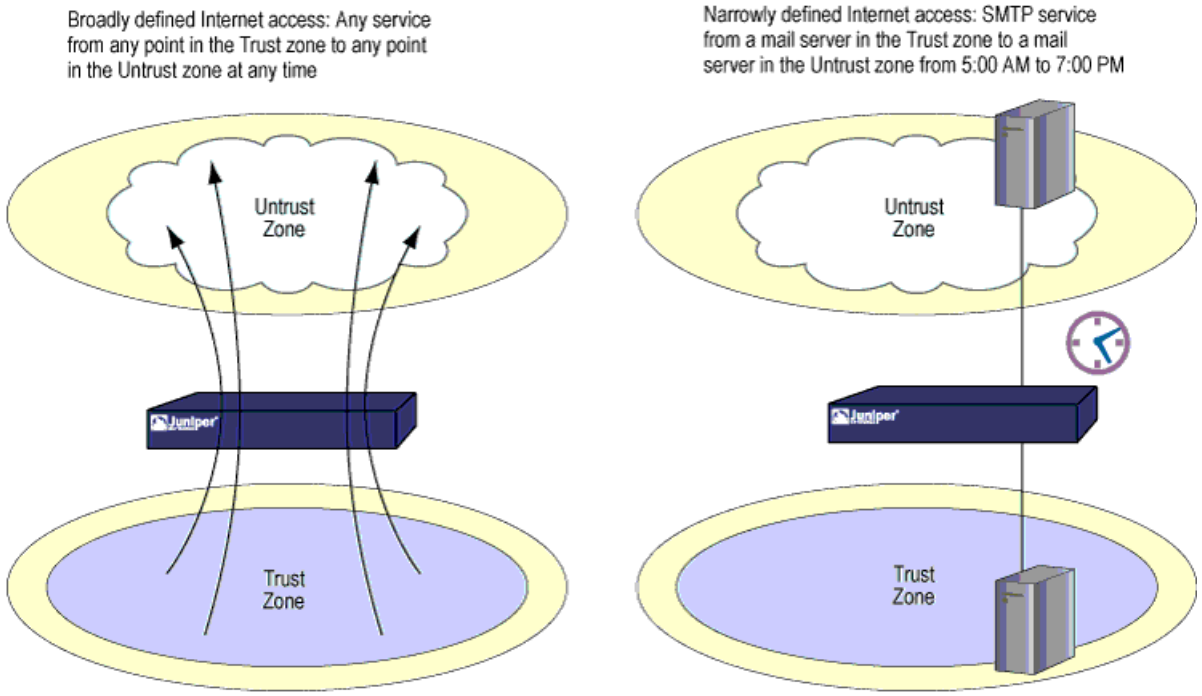

Every time a packet attempts to pass from one zone to another or between two interfaces bound to the same zone, the device checks for a policy that permits such traffic (see "Policy Application Sets Overview" on page 112). To allow traffic to pass from one security zone to another—for example, from zone 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*

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 79
- Security Policy Address Books and Address Sets on page 93
- Security Policy Schedulers on page 101
- Security Policy Applications on page 111

## 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 destination address any in policy 1 includes the more specific dst-A address in policy 2. When an HTTP packet arrives at JUNOS software from an address in the Internal zone bound for dst-A in the External zone, JUNOS software always first finds a match with policy 1.

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 84

## 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 53.
3. Configure address books. See “Configuring Address Books” on page 98.

To configure security policies with Quick Configuration:

1. Select Configuration > Quick Configuration > Security Policies > Policies.
2. Select the Default Policy Action, Deny All or Permit All.
3. Select the zone direction (from zone and to zone) as shown in Figure 13 on page 76. You must have preconfigured the security zones for which you want to set policies. For more information on zones, see “Configuring Security Zones—Quick Configuration” on page 53.

**Figure 13: Quick Configuration Policies Page for Security Policies**

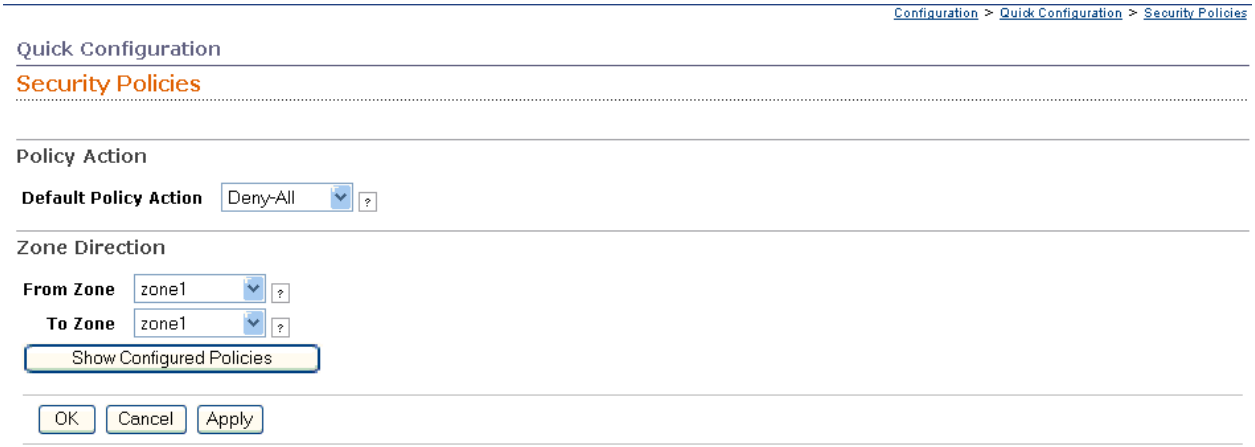


1. Click Show Configured Policies. The screen displays the message “ No policies have been defined for the selected zone direction.”
2. Click Add to configure a new policy.
3. Specify a policy name.
4. Specify a policy action. The form changes depending on the action specified (Selecting permit displays additional fields as shown in Figure 14 on page 77). See Table 33 on page 78 for the extended policy configuration fields.

**Figure 14: Security Policies Configuration**

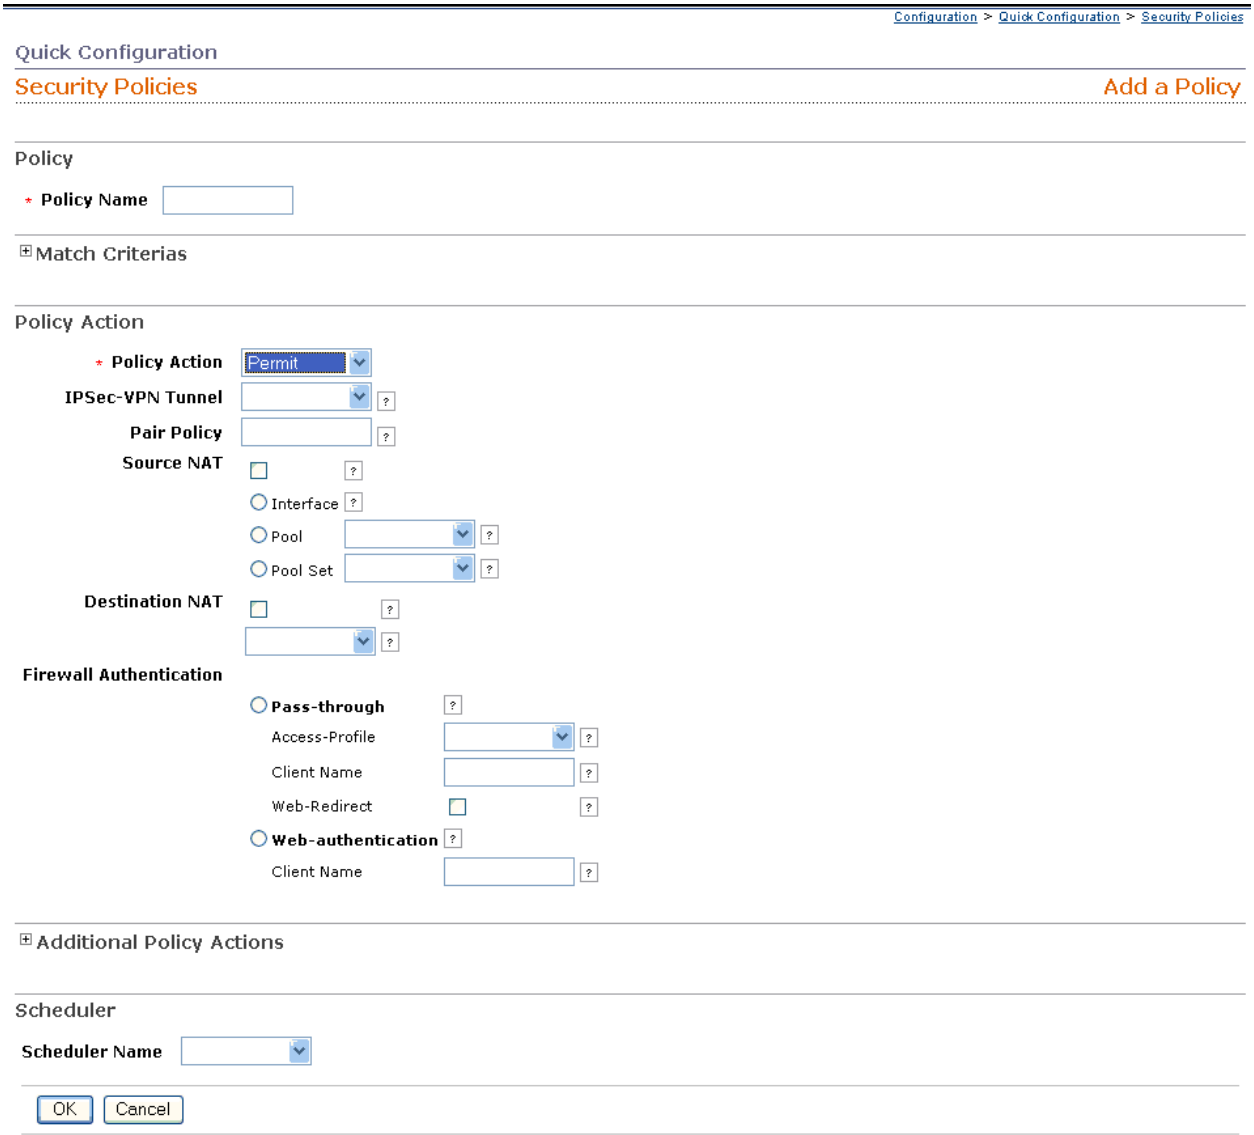


1. Optionally, you can select a scheduler name that you created earlier and whose schedule you want to associate with the policy. See “Configuring Schedulers” on page 103.
2. Click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
   - 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 33: Security Policies Configuration Options**

| Policy Action | Description |
|---|---|
| Match Criteria | Source Address—Name of the source address or address set as entered in the source zone's address book.<br><br>Destination Address—Name of the destination address or address set as entered in the destination zone's address book.<br><br>Application—Name of a preconfigured or custom application or application set. |
| Policy Action | Permit—Allows the packet to pass through the firewall.<br><br>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.<br><br>For TCP and UDP traffic, the firewall drops the packet and notifies the source host as action Deny.<br><br>Deny—Blocks and drops the packet from traversing the firewall, but does not send notification back to the source. |
| IPsec-VPN Tunnel | Name of the IPsec-VPN tunnel. |
| Pair Policy | Name of the policy with the same IPsec-VPN in the reverse direction to create a pair policy. |
| Source NAT | Enable source Network Address Translation (NAT-src) and permit address and port translation on the permitted traffic. |
| Destination NAT | Enable destination Network Address Translation (NAT-dst) and permit address and port translation on the permitted traffic. |
| Firewall Authentication | Authenticate the client before forwarding the traffic. Two types of firewall authentication:<br><br>Pass-through—Verifies traffic as it attempts to pass through the firewall.<br><br>Web authentication—Verifies client authentication.<br><br>For more information on authentication, see “Firewall User Authentication Overview” on page 147. |
| Additional Policy Actions | Count—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.<br><br>Log (session-init and session-close)—Logs session creation and session close events. |
| Scheduler | Optionally, name a scheduler whose schedule determines when the policy is active. |

## 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 51.
2. Configure the address book for the policy. (See “Configuring a Policy to Permit Traffic” on page 84.)
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 33 on page 78 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*

To configure policies using the J-Web configuration editor:

1. Select **Configuration > View and Edit > Edit Configuration.**

   The Configuration page appears.
2. Next to Interfaces, click Configure.
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 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 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 9 and click OK.

To configure interfaces:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
2. Next to Policy, select the check box and click Configure.
3. Next to Policy, click Add new entry.
4. In the From zone name box, type internal.
5. In the To zone name box, type dmz corp_net.
6. Next to Policy, click Add new entry.
7. In the Policy name box, type Pol1.
8. Select the Match check box.
9. Select the Then check box.
10. Next to Match, click Configure.
11. To match the policy to an application name, from the Source address choice list, select Source address.
12. Next to Source address, click Add new entry.
13. From the Value keyword list, select Enter specific value.
14. To specify the address of the address book, In the Address box, type corp_net and click OK.
15. To match the policy to a destination address, from the Destination address choice list, select Destination address.
16. Next to Destination address, click Add new entry.
17. In the Value keyword list, select Enter specific value.
18. To specify the destination address for the policy, in the Address box, type mail_svr and click OK.

19. To match the policy to an application set name, from the Application Choice list, select Application.
20. Next to Application, click Add new entry.
21. To specify the application set name to match the policy, from the Value keyword list, type MAIL-POP3 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.
25. To check the configuration, see "Verifying Policy Configuration" on page 83.

### *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 83.

### *Related Topics*

- Understanding Policies on page 70
- Configuring Address Books on page 98
- Policy Application Sets Overview on page 112
- Example: Configuring Applications and Application Sets on page 134
- Configuring Schedulers on page 103
- Verifying Policy Configuration on page 83

## 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 93

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 84.
2. Configure a policy to deny traffic. For more information, see "Configuring a Policy to Deny Traffic" on page 86.
3. Configure a policy with pass-through authentication. For more information, see "Configuring for Pass-Through Authentication" on page 152.
4. Configure a policy with WebAuth authentication. For more information, see "Configuring for Web Authentication" on page 157.
5. Reorder policies after they have been created. For more information, see "Reordering Policies After They Have Been Created" on page 88.

## 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 51.
3. Configure the address book for the policy. (See "Configuring a Policy to Permit Traffic" on page 84.)
4. For background information, read "Example: Configuring Security Policies—Detailed Configuration" on page 84.

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*

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

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 allowin.
7. Select the Match check box.
8. Select the Then check box.
9. Next to Match, click Configure.
10. 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. Next to Then, click Configure.
22. From the Action list, select Permit and click OK.
23. If you are finished configuring the device, commit the configuration.
24. To check the configuration, see “Verifying Policy Configuration” on page 83.

### 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 83.

### Related Topics

- Configuring a Policy to Deny Traffic on page 86
- Monitoring Policy Statistics on page 90

## 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 51.
3. Configure the address book for the policy. (See "Configuring a Policy to Permit Traffic" on page 84.)
4. Configure a policy to permit traffic. (See "Configuring a Policy to Permit Traffic" on page 84.)
5. For background information, read "Example: Configuring Security Policies—Detailed Configuration" on page 84.

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

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

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 listt, 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 83.

### *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 83.

### *Related Topics*

- Configuring for Web Authentication on page 157
- Verifying Scheduled Policies on page 108

## 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 51.)
3. Configure the address book for the policy. (See "Configuring a Policy to Permit Traffic" on page 84.)
4. Configure a policy to permit traffic. (See "Configuring a Policy to Permit Traffic" on page 84.)
5. Configure a policy to deny traffic. (See "Configuring a Policy to Deny Traffic" on page 86.)
6. Configure a policy with pass-through authentication. (See "Configuring for Pass-Through Authentication" on page 152.)
7. Configure a policy with WebAuth authentication. (See "Configuring for Web Authentication" on page 157.)
8. For background information, read "Example: Configuring Security Policies—Detailed Configuration" on page 84.

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 73

## Troubleshooting Policy Configuration

Most policy configuration failures occur during a commit or runtime.

This topic covers:



### *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.

Chapter 7

# 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.

This section includes:



## 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 15 on page 95.

**Figure 15: Address Sets**

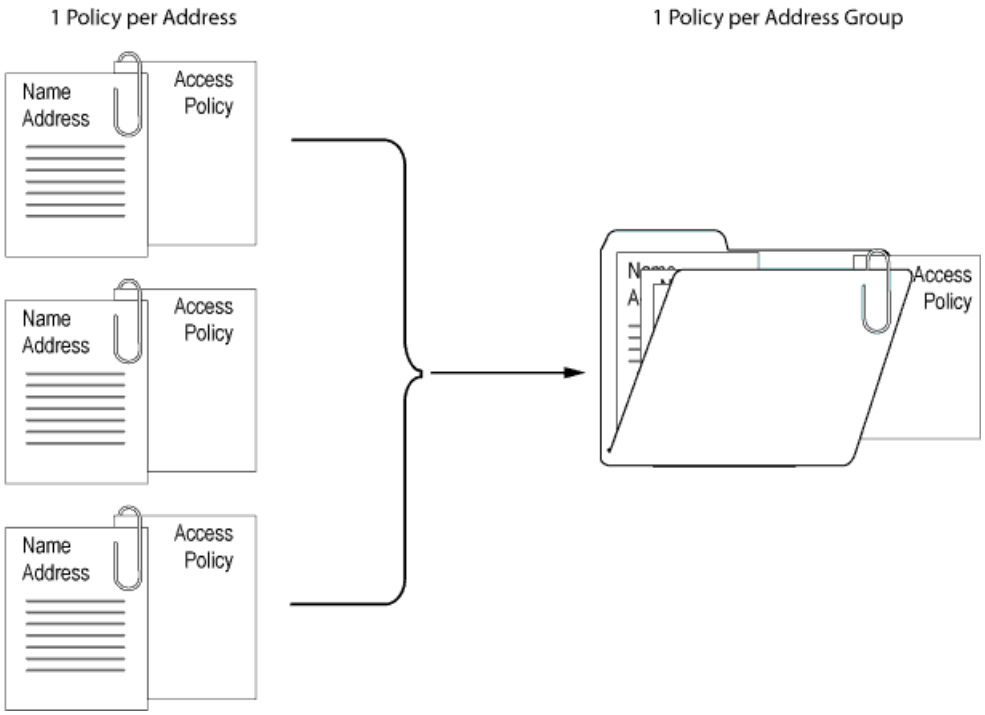


The address set option has the following features:

- 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 49 and "Configuring Address Books" on page 98.
2. Configure zones. For more information, see "Configuring Security Zones—Quick Configuration" on page 53.

To configure addresses and address sets with Quick Configuration:

1. Select **Configuration > Quick Configuration > Policy Elements > Address Books**.
2. Select the zone for which you want to configure address books. For more information on creating zones, see "Creating Security Zones" on page 51.
3. Click Add to add IP addresses to the address book. See Figure 16 on page 97.

**Figure 16: Quick Configuration Policy Elements Page for Address Books**

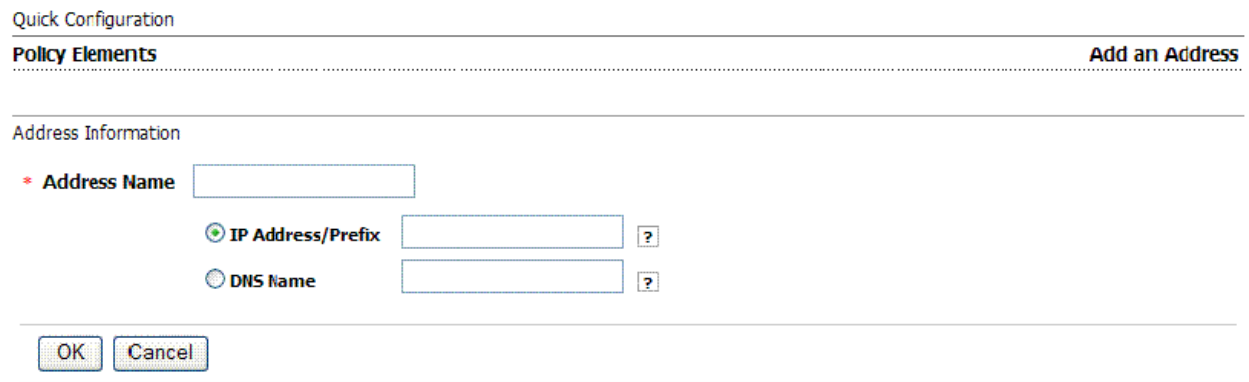


4. Specify an address name.
5. Specify either an IP address or a DNS name. The IP address must be an IPv4 address with the number of prefix bits. You can use domain names only if the system is configured to use DNS services.
6. Click one of the following buttons:
   - To apply the configuration and stay 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.
7. Continue adding addresses using Step 2 through Step 5. Once you have the addresses you need to create an address set, select the Address Sets tab.
8. Click Add. See Figure 17 on page 97.

**Figure 17: Configuring Address Sets**

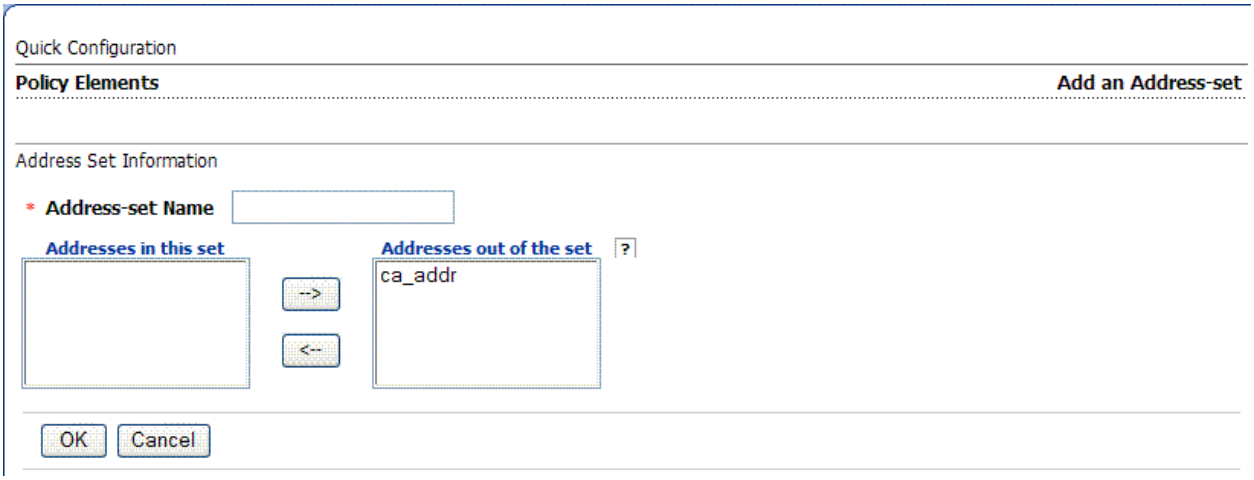

9. Specify an Address Set name.
10. To add the predefined addresses in the " Addresses in this set" column, use the left and right arrows.
11. Click OK to save the configuration and return to the main Configuration page.

## 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*

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

1. Select **Configuration > View and Edit > Edit Configuration**.

   The Configuration page appears.
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 IntranetGREEN.
6. Next to Address Book, **Configure**.
7. Next to Address, **Add new entry**.
8. In the Address name box, type G1 10.1.10.0/24 and **OK**.
9. Next to Address, **Add new entry**.

   You need to have more than one address in the Address book to configure an Address Set.
10. In the Address name box, type G2 10.1.10.0/32 and click **OK**.
11. Next to Address Set, click **Add new entry**.
12. In the Address set name box, type SerAll and click **OK.**

13. If you are finished configuring the device, commit the configuration.
14. To check the configuration, see "Verifying Address Book Configuration" on page 99.

### 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 99.

### Related Topics

- Understanding Policy Configuration on page 73
- Policy Application Sets Overview on page 112
- Example: Configuring Applications and Application Sets on page 134
- Configuring Policies on page 79
- Configuring Schedulers on page 103

## 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 8

# 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.

This section includes:



## Configuring a Scheduler—Quick Configuration

You can use J-Web Quick Configuration to quickly configure a schedule for the security policies. See Figure 18 on page 102.

**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 53.
3. Configure address books. See “Configuring Address Books” on page 98.
4. Configure security policies that are going to be associated with a schedule. See “Configuring Policies—Quick Configuration” on page 76.

**Figure 18: Quick Configuration Policies Page for Schedulers**

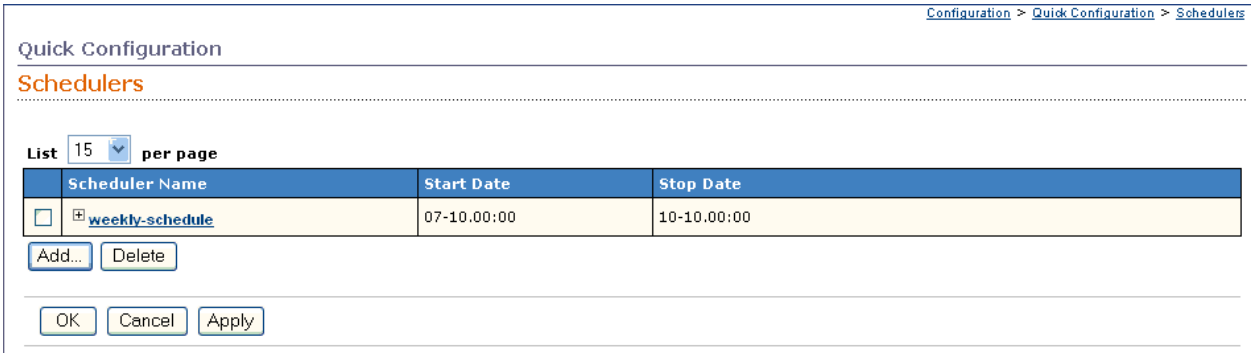


To configure a scheduler with Quick Configuration:

1. Select Configuration > Quick Configuration > Schedulers.
2. Click Add. See Figure 19 on page 102.

**Figure 19: Scheduler Configuration**

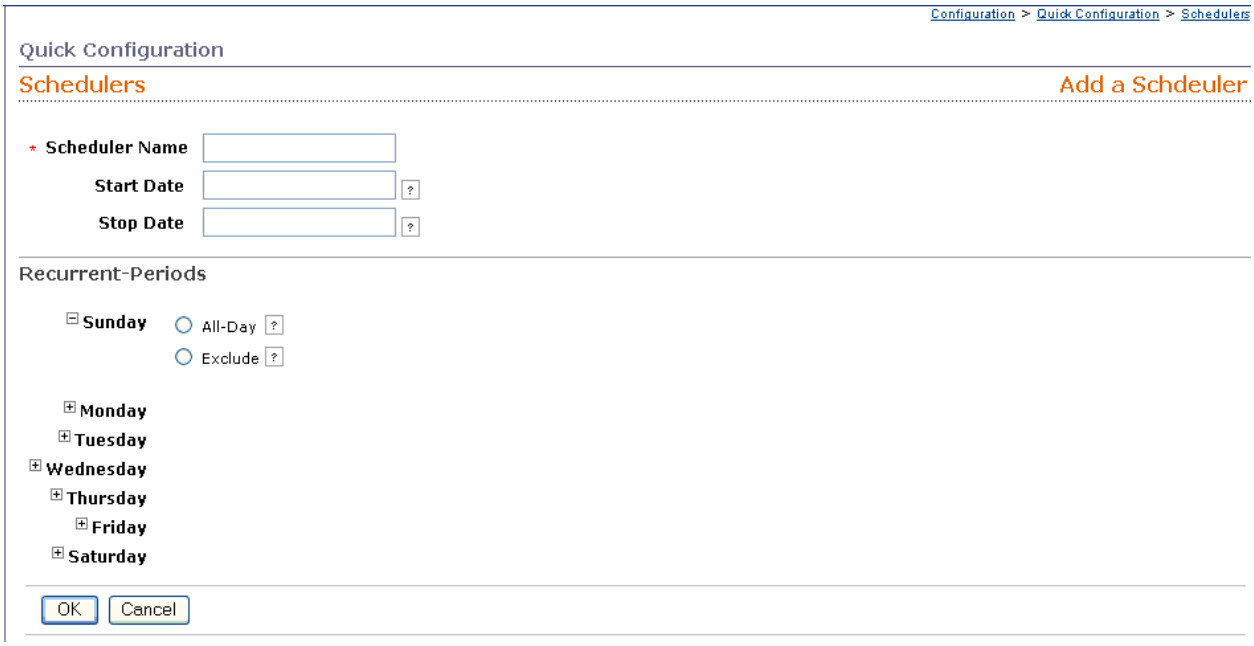


1. Specify a scheduler name.
2. Specify the start and stop date and times in the YYYY.MM.DD.hh.mm (year, month, date, hour, minutes) format.
3. Select the recurrent period, if needed. Click the plus sign by a specific day to select all day or exclude a day from the weekly schedule.
4. Click one of the 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.

## 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 103
- CLI Configuration on page 105
- Related Topics on page 106


### *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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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.

To associate the schedule with the policy, allowing access during regular work hours:

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
2. Next to Policies, select the check box and click Configure or Edit.
3. In the From zone name box, type green.
4. In the To zone name box, type red.
5. In the Policy name box, type abc.
6. 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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
2. In the Scheduler name box, type scheduler SatHrs.
3. Next to Saturday, click Configure.
4. From the Daily type list, select Start time.
5. Next to Start time, click Add new entry.
6. In the Start time value box, type 12:00:00.
7. In the Stop time box, type 18:00:00 and click OK.
8. 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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
2. Next to Policies, select the check box and click Configure or Edit.

3. In the From zone name box, type green.
4. In the To zone name box, type red.
5. In the Policy name box, type abc.
6. To bind these schedules to the policy during Saturdays, in the Scheduler name box, type SatHrs and click OK.
7. To bind these schedules to the policy during Sundays, in the Scheduler name box, type SunHrs and click OK.
8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see "Verifying Scheduled Policies" on page 108.

### *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 108.

### *Related Topics*

- Understanding Policy Configuration on page 73
- Associating a Policy to a Scheduler on page 106
- Configuring Address Books on page 98
- Example: Configuring Applications and Application Sets on page 134
- Configuring Policies on page 79

## 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 93

To associate a policy to a scheduler, use either J-Web or the CLI configuration editor.

This topic covers:

- J-Web Configuration on page 106
- CLI Configuration on page 107
- Related Topics on page 108

### *J-Web Configuration*

To associate a policy to a scheduler using the J-Web configuration editor:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
2. From the list of configured policies, click the policy you need to associate with a scheduler.
3. Under the Nested Configuration column, click Policy.
4. Under the Actions column, click Edit.
5. 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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
2. In the Scheduler name box, type scheduler SatHrs.
3. Next to Saturday, click Configure.
4. From the Daily type list, select Start time.
5. Next to Start time, click Add new entry.
6. In the Start time box, type 12:00:00.
7. In the Stop time box, type 18:00:00 and click OK.
8. 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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 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 108.

### *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 108.

### *Related Topics*

- Configuring Policies on page 79
- Configuring Schedulers on page 103
- Verifying Scheduled Policies on page 108

## 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 9

# 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 131.

This section includes:

## 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 set*s. 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*

## 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 111.

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 34 on page 113 lists ICMP message names, the corresponding code, type, and description.

**Table 34: ICMP Messages**

| ICMP Message Name | Type | Code | Description |
|---|---|---|---|
| ICMP-ANY | all | all | ICMP-ANY affects any protocol using ICMP.<br><br>Denying ICMP-ANY impairs any attempt to ping or monitor a network using ICMP.<br><br>Permitting ICMP-ANY allows all ICMP messages. |
| ICMP-ADDRESS-MASK<br>■ Request<br>■ Reply | 17<br><br>18 | 0<br><br>0 | ICMP address mask query is used for systems that need the local subnet mask from a bootstrap server.<br><br>Denying ICMP address mask request messages can adversely affect diskless systems.<br><br>Permitting ICMP address mask request messages might allow others to fingerprint the operating system of a host in your network. |

**Table 34: 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.<br><br>Codes 0, 1, 4, or 5 can be from a gateway. Codes 2 or 3 can be from a host (RFC 792).<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>We recommend denying these messages from the Internet to a trusted network.<br><br>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<br>■ Request<br>■ Reply | 15<br><br>16 | 0<br><br>0 | ICMP-INFO query messages allow diskless host systems to query the network and self-configure.<br><br>Denying ICMP address mask request messages can adversely affect diskless systems.<br><br>Permitting ICMP address mask request messages might allow others to broadcast information queries to a network segment to determine computer type. |

**Table 34: 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<br><br>We recommend denying these messages from the Internet to a trusted network.<br><br>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.<br><br>We recommend denying these messages from the Internet to a trusted network.<br><br>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.<br><br>We recommend denying these messages from the Internet to a trusted network.<br><br>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.<br><br>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 34: 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.<br><br>Denying these messages will not help or impair internal network performance.<br><br>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<br><br>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.<br><br>We recommend denying these messages from a trusted network out to the Internet. |
| ICMP-TIMESTAMP<br>■ Request<br>■ Reply | 13<br><br>14 | 0<br><br>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.<br><br>Denying ping functionality removes your ability to check to see if a host is active.<br><br>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.<br><br>We recommend denying these messages. |
| Traceroute<br>■ Forward<br>■ Discard | 30<br><br>30 | 0<br><br>1 | Traceroute is a utility to indicate the path to access a specific host.<br><br>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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 35 on page 117 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 35: 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<br>67 server | Dynamic Host Configuration Protocol allocates network addresses and delivers configuration parameters from server to hosts. |

**Table 35: Predefined Applications** *(continued)*

| Application Name | Port(s) | Application Description |
|---|---|---|
| DNS | 53 | Domain Name System translates domain names into IP addresses. |
| FTP<br>■ FTP-Get<br>■ FTP-Put | 20 data<br>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.<br><br>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.<br><br>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).<br><br>Denying HTTP application disables your users from viewing the Internet.<br><br>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.<br><br>Denying HTTPS disables your users from shopping on the Internet and from accessing certain online resources that require secure password exchange.<br><br>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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 36 on page 119 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 36: Predefined Microsoft Applications**

| Application | Parameter/UUID | Description |
| --- | --- | --- |
| MS-RPC-EPM | 135<br>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:<br>■ MS-AD-BR<br>■ MS-AD-DRSUAPI<br>■ MS-AD-DSROLE<br>■ MS-AD-DSSETUP |
| MS-EXCHANGE | 6 members | Microsoft Exchange application group includes:<br>■ MS-EXCHANGE-DATABASE<br>■ MS-EXCHANGE-DIRECTORY<br>■ MS-EXCHANGE-INFO-STORE<br>■ MS-EXCHANGE-MTA<br>■ MS-EXCHANGE-STORE<br>■ MS-EXCHANGE-SYSATD |

**Table 36: Predefined Microsoft Applications** *(continued)*

| Application | Parameter/UUID | Description |
|---|---|---|
| MS-IIS | 6 members | Microsoft IIS Server application group includes:<br>■ MS-IIS-COM<br>■ MS-IIS-IMAP4<br>■ MS-IIS-INETINFO<br>■ MS-IIS-NNTP<br>■ MS-IIS-POP3<br>■ MS-IIS-SMTP |
| MS-AD-BR | ecec0d70-a603-11d0-96b1-00a0c91ece30<br>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<br>f5cc5a7c-4264-101a-8c59-08002b2f8426<br>f5cc59b4-4264-101a-8c59-08002b2f8426 | Microsoft Exchange Directory application |
| MS-EXCHANGE-INFO-STORE | 0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde<br>1453c42c-0fa6-11d2-a910-00c04f990f3b<br>10f24e8e-0fa6-11d2-a910-00c04f990f3b<br>1544f5e0-613c-11d1-93df-00c04fd7bd09 | Microsoft Exchange Information Store application |
| MS-EXCHANGE-MTA | 9e8ee830-4459-11ce-979b-00aa005ffebe<br>38a94e72-a9bc-11d2-8faf-00c04fa378ff | Microsoft Exchange MTA application |
| MS-EXCHANGE-STORE | 99e66040-b032-11d0-97a4-00c04fd6551d<br>89742ace-a9ed-11cf-9c0c-08002be7ae86<br>a4f1db00-ca47-1067-b31e-00dd010662da<br>a4f1db00-ca47-1067-b31f-00dd010662da | Microsoft Exchange Store application |

**Table 36: Predefined Microsoft Applications** *(continued)*

| Application | Parameter/UUID | Description |
|---|---|---|
| MS-EXCHANGE-SYSATD | 67df7c70-0f04-11ce-b13f-00aa003bac6c<br>f930c514-1215-11d3-99a5-00a0c9b61b04<br>83d72bf0-0d89-11ce-b13f-00aa003bac6c<br>469d6ec0-0d87-11ce-b13f-00aa003bac6c<br>06ed1d30-d3d3-11cd-b80e-00aa004b9c30 | Microsoft Exchange System Attendant application |
| MS-FRS | f5cc59b4-4264-101a-8c59-08002b2f8426<br>d049b186-814f-11d1-9a3c-00c04fc9b232<br>a00c021c-2be2-11d2-b678-0000f87a8f8e | Microsoft File Replication application |
| MS-IIS-COM | 70b51430-b6ca-11d0-b9b9-00a0c922e750<br>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<br>130ceefb-e466-11d1-b78b-00c04fa32883 | Microsoft Inter-site Messaging application |
| MS-MESSENGER | 17fdd703-1827-4e34-79d4-24a55c53bb37<br>5a7b91f8-ff00-11d0-a9b2-00c04fb6e6fc | Microsoft Messenger application |
| MS-MQQM | fdb3a030-065f-11d1-bb9b-00a024ea5525<br>76d12b80-3467-11d3-91ff-0090272f9ea3<br>1088a980-eae5-11d0-8d9b-00a02453c33<br>5b5b3580-b0e0-11d1-b92d-0060081e87f0<br>41208ee0-e970-11d1-9b9e-00e02c064c39 | Microsoft Windows Message Queue Management application |
| MS-NETLOGON | 12345678-1234-abcd-ef00-01234567cffb | Microsoft Netlogon application |

**Table 36: Predefined Microsoft Applications** *(continued)*

| Application | Parameter/UUID | Description |
|---|---|---|
| MS-SCHEDULER | 1ff70682-0a51-30e8-076d-740be8cee98b<br>378e52b0-c0a9-11cf-822d-00aa0051e40f<br>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<br>811109bf-a4e1-11d1-ab54-00a0c91e9b45 | Microsoft WINS application |

### *Related Topics*

- Example: Adding a Custom Policy Application on page 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Depending on your network requirements, you can choose to permit or deny messages generated from and packets of these dynamic routing protocols. Table 37 on page 122 lists each supported dynamic routing protocol by name, port, and description.

**Table 37: 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 38 on page 123 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 38: Supported Streaming Video Applications**

| Application | Port | Description |
| --- | --- | --- |
| H.323 | TCP source 1-65535; TCP destination 1720, 1503, 389, 522, 1731<br><br>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<br><br>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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 39 on page 124 lists each Sun remote procedure call Application Layer Gateway (RPC ALG) application name, parameters, and full name.

**Table 39: 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<br>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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 40 on page 125 lists each supported application and gives the default port(s) and a description of each entry.

**Table 40: Supported Applications**

| Application | Port | Description |
| --- | --- | --- |
| IKE | UDP source 1-65535; UDP destination 500<br><br>4500 (used for NAT traversal) | Internet Key protocol (IKE) is a protocol to obtain authenticated keying material for use with ISAKMP.<br><br>When configuring auto IKE, you can choose from three predefined Phase 1 or Phase 2 proposals:<br>■ Standard: AES and 3DES<br>■ Basic: DES and two different types of authentication algorithms<br>■ Compatible: Four commonly used authentication and encryption algorithms |
| L2TP | 1723 | L2TP combines PPTP with Layer 2 Forwarding (L2F) for remote access. |
| PPTP | — | 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 41 on page 126 lists the predefined IP-related applications. Each entry includes the default port and a description of the application.

**Table 41: 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 42 on page 127 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 42: 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 43 on page 127 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 43: 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 43: 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 44 on page 129 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 44: 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 email. |

### *Related Topics*

- Example: Adding a Custom Policy Application on page 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 45 on page 129 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 45: Predefined UNIX Applications**

| Application | Port | Description |
| --- | --- | --- |
| FINGER | 79 | Finger is a UNIX program that provides information about the users. |

**Table 45: Predefined UNIX Applications** *(continued)*

| Application | Port | Description |
|---|---|---|
| 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

Table 46 on page 130 lists predefined miscellaneous applications. Each entry includes the application name, default or assigned port, and a description of the application.

**Table 46: 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;<br>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) |
| VNC | 5800 | Virtual Network Computing facilitates viewing and interacting with another computer or mobile Juniper Networks device connected to the Internet. |

**Table 46: Predefined Miscellaneous Applications** *(continued)*

| Application | Port | Description |
|---|---|---|
| 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 136
- Example: Modifying a Custom Policy Application on page 138

## 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 111. |

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 131
- Related Topics on page 132

### *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 136
- Example: Modifying a Custom Policy Application on page 138
- Example: Defining a Custom Internet Control Message Protocol Application on page 139

## 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.

**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 53.
3. Configure address books. For more information, see "Configuring Addresses and Address Sets—Quick Configuration" on page 96.

To configure applications and application sets with Quick Configuration:

1. Select Configuration > Quick Configuration > Policy Elements > Applications / Application Sets.
2. Select the Application Sets tab. See Figure 20 on page 133.

**Figure 20: Quick Configuration Page for Configuring Applications or Application Sets**

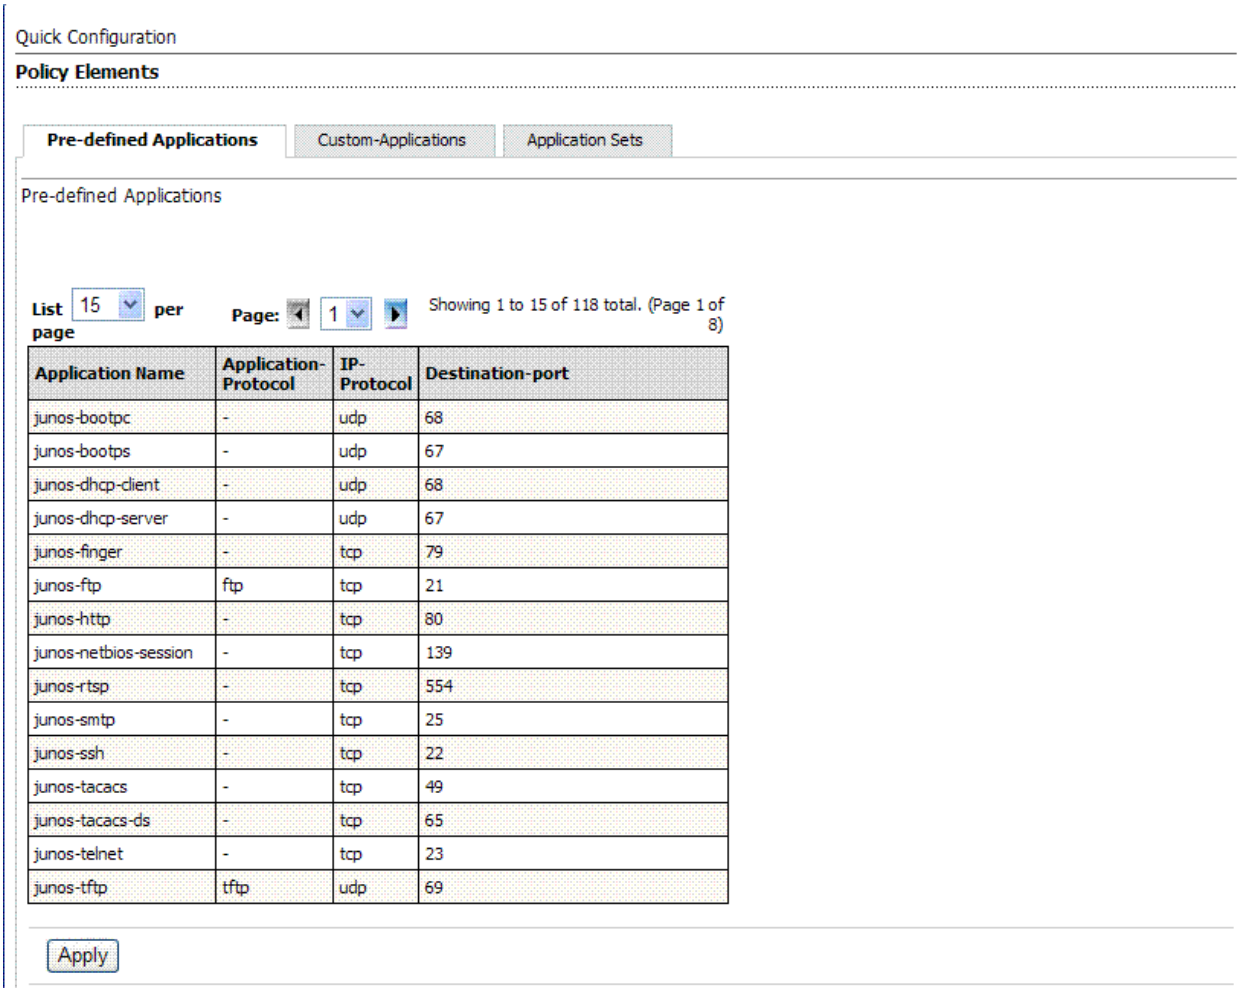


3. Click **Add**. The Add an Application Set window appears as shown in Figure 21 on page 134.

**Figure 21: Adding Application Sets**

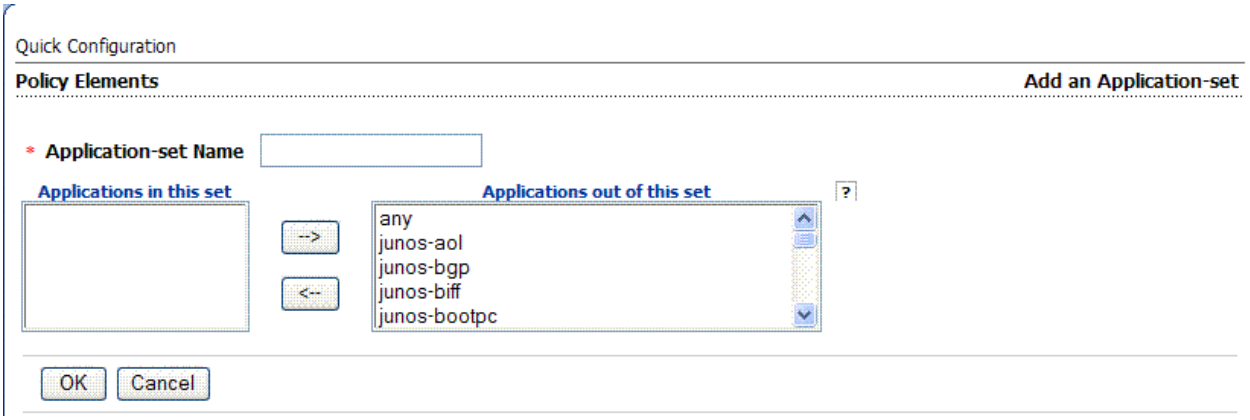


4. Enter an application set name (name for one or a group of preconfigured applications).
5. Use the left and right arrows to add or remove applications from the Applications in this set column.
6. Click one of the following buttons:
   - To apply the configuration and return to the main Configuration page, click OK.
   - To cancel your entries and return to the main page, click Cancel.

## 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*

To specify applications and application sets using the J-Web configuration editor:

1. Select **Configuration > View and Edit > Edit Configuration**.

   The Configuration page appears.
2. To create an application set for the three applications, next to Applications, click Configure or Edit.
3. Next to Application set, click Configure or Edit.
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 **Configuration > View and Edit > Edit Configuration**.

   The Configuration page appears.
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. Next to Application, click Add new entry.
6. In the Application name box, type smtp and click OK.
7. To specify another application name:
   a. Next to Application, click Add new entry.
   b. In the Application name box, type **http** and click OK.
8. To specify another application name:
   a. Next to Application, click Add new entry.
   b. In the Application name box, type **https** and click OK.
9. To specify another application name:
   a. Next to Application, click Add new entry.
   b. In the Application name box, type **POPS** and click OK.
10. 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 112

## 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 84.
3. Create zones. See "Creating Security Zones" on page 51.
4. For background information, read:
   - Understanding Custom Policy Applications on page 131
   - Security Policy Applications on page 111

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

To add a custom application to application book using the J-Web configuration editor:

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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: Defining a Custom Internet Control Message Protocol Application on page 139

## 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 84.
3. Create zones. See "Creating Security Zones" on page 51.
4. For background information, read:
   - Understanding Custom Policy Applications on page 131
   - Security Policy Applications on page 111

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 138
- CLI Configuration on page 139
- Related Topics on page 139

### *J-Web Configuration*

To modify the custom policy application using the J-Web configuration editor:

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 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 136
- Example: Defining a Custom Internet Control Message Protocol Application on page 139

## 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 113
   - Understanding Custom Policy Applications on page 131
   - Security Policy Applications on page 111

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 47 on page 140.

**Table 47: 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 140
- CLI Configuration on page 141
- Related Topics on page 141


## *J-Web Configuration*

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
2. Next to Applications, click Configure or Edit.
3. Next to Application, click Add new entry.
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. To specify the inactivity timeout, corresponding to the application name host-unreachable, click Edit.

10. In the Inactivity timeout box, type 4 and click OK.
11. 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 136
- Example: Modifying a Custom Policy Application on page 138
- Example: Defining a Custom Internet Control Message Protocol Application on page 139

## 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 111
- Understanding Policies on page 70

This topic covers:



### 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 48 on page 142.

**Table 48: 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 144
- Understanding Policies on page 70
- Policy Application Sets Overview on page 112

## 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 141
   - Address Books and Address Sets Overview on page 93

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 141

Chapter 10

# 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.

This section includes:



## 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

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 149
- Understanding Client Groups for Firewall Authentication on page 162
- Understanding Banner Customization on page 172

## 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 147.

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 149
- Web Authentication on page 150
- Related Topics on page 152

### *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 22 on page 150.

**Figure 22: Policy Lookup for a User**

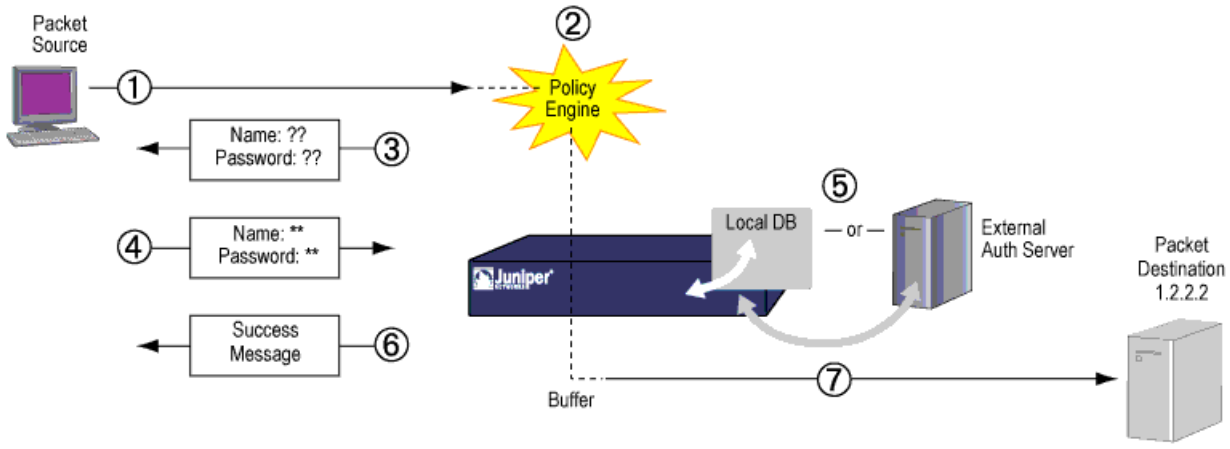


1. A client user sends an FTP, an HTTP, or a Telnet packet to 1.2.2.2.
2. The Juniper Networks device intercepts the packet, notes that its policy requires authentication from either the local database or an external authentication server, and buffers the packet.
3. The Juniper Networks device prompts 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 23 on page 151.

**Figure 23: Web Authentication Example**

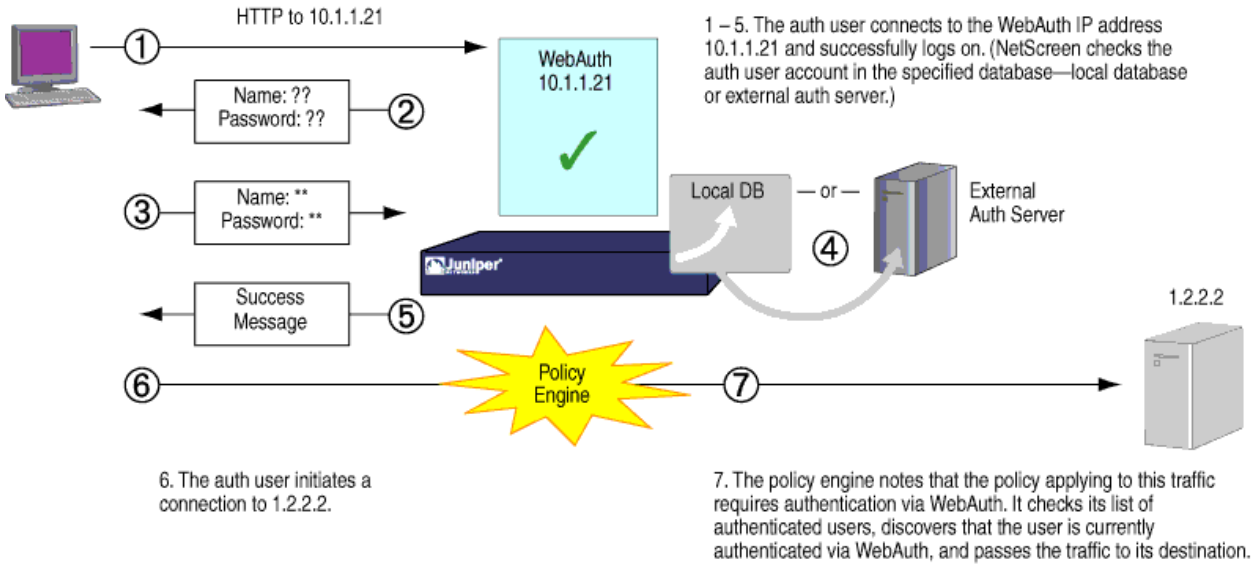


Follow these Web Authentication guidelines:

- 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 49.)
- 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 152
- Configuring for Web Authentication on page 157

## 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 24 on page 152.

**Before You Begin**

1. Establish basic connectivity. (See the Getting Started Guide for your device.)
2. Read "Understanding Authentication Schemes" on page 149.
3. Read "Firewall User Authentication Overview" on page 147.

**Figure 24: Configuring Pass-Through Firewall Authentication**

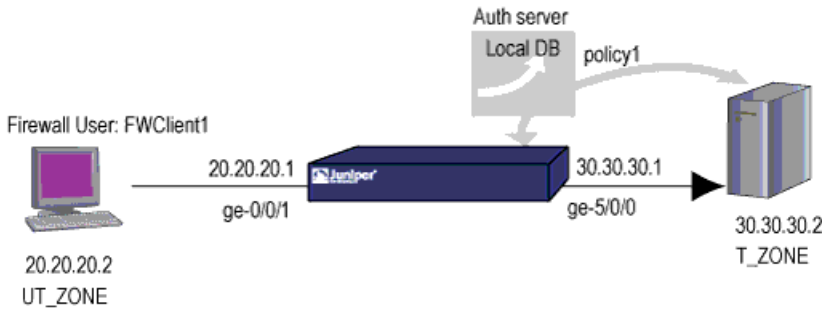


To configure pass-through firewall authentication, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 152
- CLI Configuration on page 155
- Related Topics on page 156

### *J-Web Configuration*

To create IP addresses for the interfaces on the device using the J-Web configuration editor:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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. Next to Firewall Authentication, click Configure or Edit.
2. Next to Pass through, click Configure or Edit.
3. 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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 177.

### *CLI Configuration*

To configure the device for pass-through firewall authentication as shown in Figure 24 on page 152, 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 177.

## *Related Topics*

- Firewall User Authentication Overview on page 147
- Understanding Client Groups for Firewall Authentication on page 162

## 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 149.
3. Read "Firewall User Authentication Overview" on page 147.

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 164.

To configure a server for web authentication in the scenario shown in Figure 23 on page 151, use either J-Web or the CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure the device for Web authentication using the J-Web configuration editor:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 succes 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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 Configuration>View and Edit >Edit Configuration.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 177

## *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 177.

   The following screen appears after the firewall user is authenticated.


## Related Topics

- Understanding Client Groups for Firewall Authentication on page 162

- Firewall User Authentication Overview on page 147

## 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 147.

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 162
- CLI Configuration on page 163
- J-Web Configuration on page 163
- CLI Configuration on page 164
- Related Topics on page 164

### *J-Web Configuration*

To configure a local user called client1 for groups G1, G2, and G3 using the J-Web configuration editor:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
2. Next to Access, click Configure or Edit.
3. Next to Profile, click Add new entry.
4. In the Profile name box, type Managers.
5. Next to Client, click Add new entry.
6. In the Name, type Client1.
7. Next to Firewall User, click Configure or Edit.
8. In the Password box, type test and click OK.
9. Next to Client group, click Add new entry.
10. In the Value box, type [g1 g2 g3] and click OK.
11. If you are finished configuring the device, commit the configuration.
12. To check the configuration, see "Verifying Firewall User Authentication" on page 177.

### *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 177.

### *J-Web Configuration*

To configure a default client group for all users using the J-Web configuration editor:

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 177.

### *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 177.

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 164

## 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 147.

To configure a server for external authentication, use either J-Web or the CLI configuration editor.

This topic covers:



## *J-Web Configuration*

To specify the RADIUS server for external authentication order using the J-Web configuration editor:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 Configure or Edit.
4. In the Value box, type alpha and click OK.
5. To specify another client group, in the Value box, type beta and click OK.
6. To specify another client group, in the Value box, type gamma 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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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. Next to Ldap options, click Configure or Edit.
2. In the Base distinguished name box, type CN=Users,DC=screenos,DC=spg,DC=juniper,DC=net
3. From the Search type list, select Search.
4. Next to Search, click Configure or Edit.
5. In the Search filter box, type sAMAccountName= and click OK.
6. Select the Admin search check box and click Configure or Edit.
7. In the Distinguished name box, type cn=administrator,cn=users,dc=screenos,dc=spg,dc=juniper,dc=net.
8. In the Password box, type pwd10 and click OK.
9. Next to Ldap server, click Add new entry.
10. In the Name box, type 3.3.3.3 and click OK.

To configure the IP addresses for the two RADIUS servers:

1. Next to Radius server, click Add new entry.
2. In the Address box, type 4.4.4.4 and click OK.
3. In the Secret box, type any unreadable data.
4. In the Retry box, type 10 and click OK.
5. Next to Radius server, click Add new entry.
6. In the Address box, type 5.5.5.5 and click OK.
7. In the Secret box, type any unreadable data.
8. If you are finished configuring the device, commit the configuration.
9. To check the configuration, see "Verifying Firewall User Authentication" on page 177

## *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 177

### *Related Topics*

- “Firewall User Authentication Overview” on page 147
- “Understanding Client Groups for Firewall Authentication” on page 162

## 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 147.

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 169
- Configuring for External Authentication Servers on page 164

## 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 168.

This example uses SecurID as the external authentication server.

The topic covers:

### *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 164
- Firewall User Authentication Overview on page 147

## 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*

To display the authentication table, select Monitor > Firewall Authentication as shown in Figure 25 on page 172.

**Figure 25: Firewall Authentication Table**

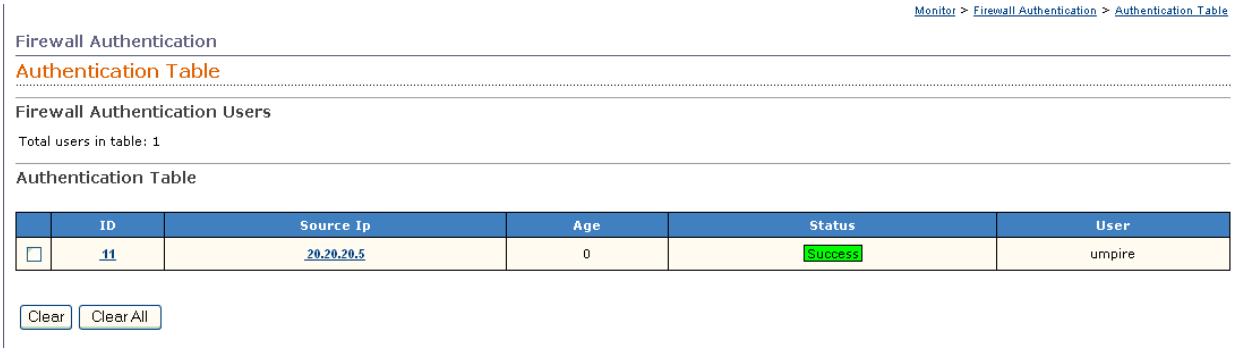


### *CLI Configuration*

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 147
- Understanding Authentication Schemes on page 149

## Understanding Banner Customization

A banner is a message that appears on a monitor in different places depending on the type of login:

**Figure 26: Banner Customization**


- At the top of a browser screen after a user has successfully logged onto a Web Authentication address as shown Figure 26 on page 172.

- 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 173

## 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 172. |

To customize a banner, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 173
- CLI Configuration on page 174
- Related Topics on page 175

### *J-Web Configuration*

To specify the banner text for failed pass-through authentication through FTP using the J-Web configuration editor:

1. Select Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 Configuration>View and Edit>Edit Configuration.

   The Configuration page appears.
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 177

## *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 177.

### *Related Topics*

- Understanding Banner Customization on page 172

## Configuring Firewall Authentication—Quick Configuration

You can use J-Web Quick Configuration to quickly configure banner pages for firewall authentication. Figure 27 on page 176 shows the Quick Configuration banner page for authentication.

**Before You Begin**

For background information, read

- Firewall User Authentication Overview on page 147
- Understanding Banner Customization on page 172
- Customizing a Banner on page 173

**Figure 27: Quick Configuration Page for Firewall Authentication Banners**

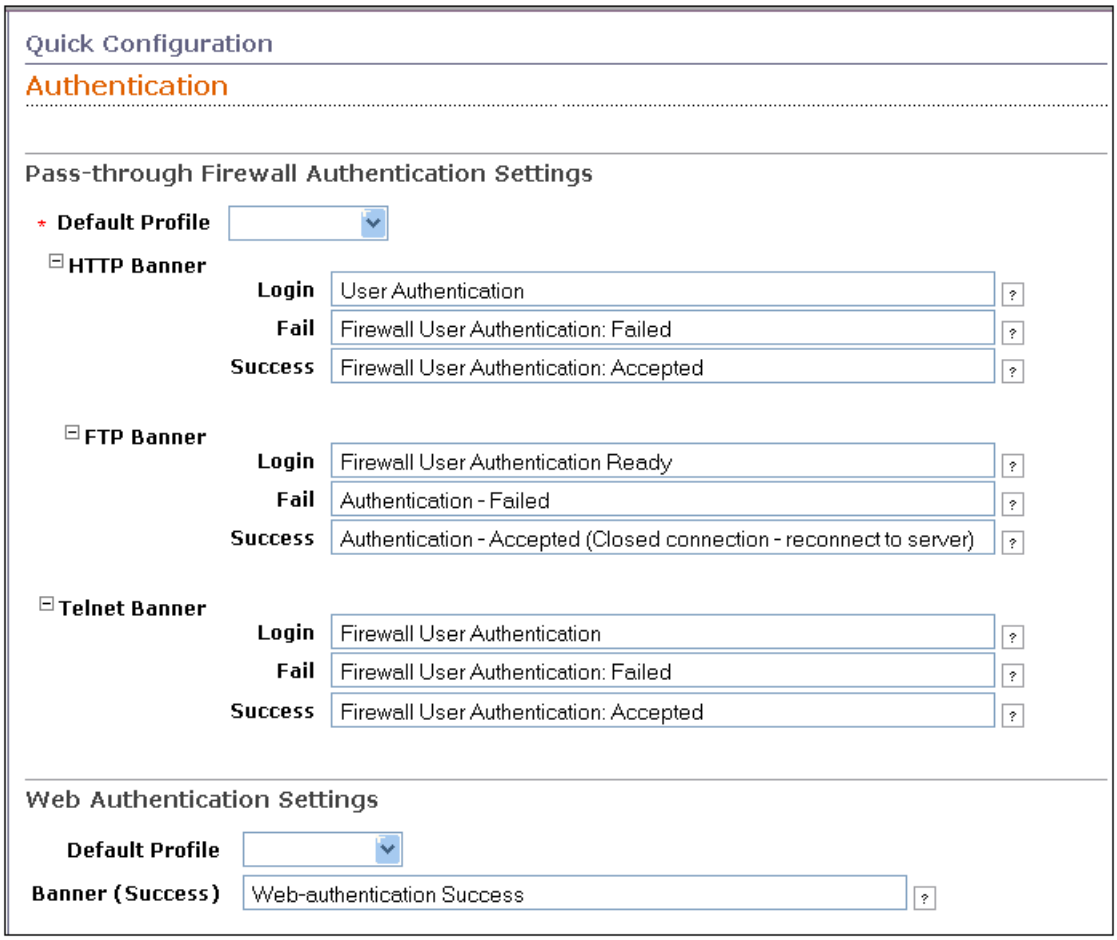


To configure banners for firewall authentication with Quick Configuration, follow these steps:

1. In the J-Web user interface, select Configuration > Quick Configuration > Authentication > Firewall Authentication.
2. Enter information into the Quick Configuration page for Firewall Authentication banners as described in Table 49 on page 177.
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 49: 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. |
| HTTP Banner | Customize banner text for users logging in using HTTP:<br><br>Login: This message displays the login prompt.<br><br>Fail: This message displays failed login.<br><br>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:<br><br>Login: This message displays the login prompt.<br><br>Fail: This message displays failed login.<br><br>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:<br><br>Login: This message displays the login prompt.<br><br>Fail: This message displays failed login.<br><br>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.<br><br>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 11

# 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 software software uses 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.

This section includes:

## 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 179.

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 182
- Understanding Port Scanning on page 184
- Understanding Network Reconnaissance Using IP Options on page 187
- Understanding Operating System Probes on page 191
- Understanding Attacker Evasion Techniques on page 197

## 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 179
- Reconnaissance Deterrence Overview on page 181

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 28 on page 183.

**Figure 28: Address Sweep**

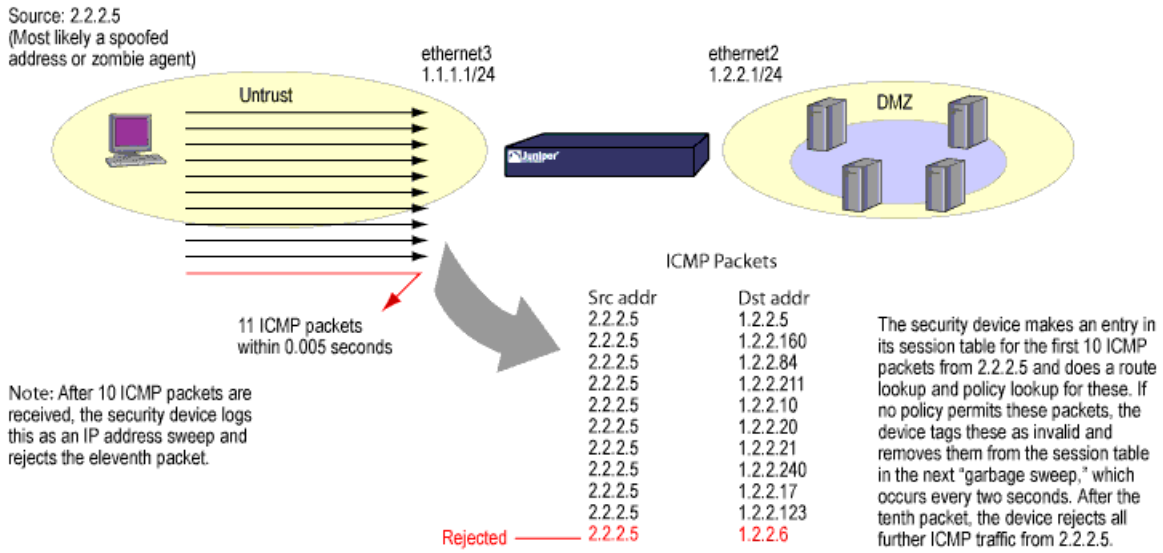


Consider enabling this SCREEN option for a security zone only if there is a policy permitting ICMP traffic from that zone. Otherwise, you do not need to enable it. The lack of such a policy denies all ICMP traffic from that zone, precluding an attacker from successfully performing an IP address sweep anyway.

### *Related Topics*

- Blocking IP Address Sweeps on page 183

## Blocking IP Address Sweeps

| **Before You Begin** |
| --- |
| For background information, read "Understanding IP Address Sweeps" on page 182. |

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 183
- CLI Configuration on page 184
- Related Topics on page 184

### *J-Web Configuration*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure.
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 **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 179
- Understanding Port Scanning on page 184

## 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 179
- Reconnaissance Deterrence Overview on page 181

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 29 on page 185.

**Figure 29: Port Scan**

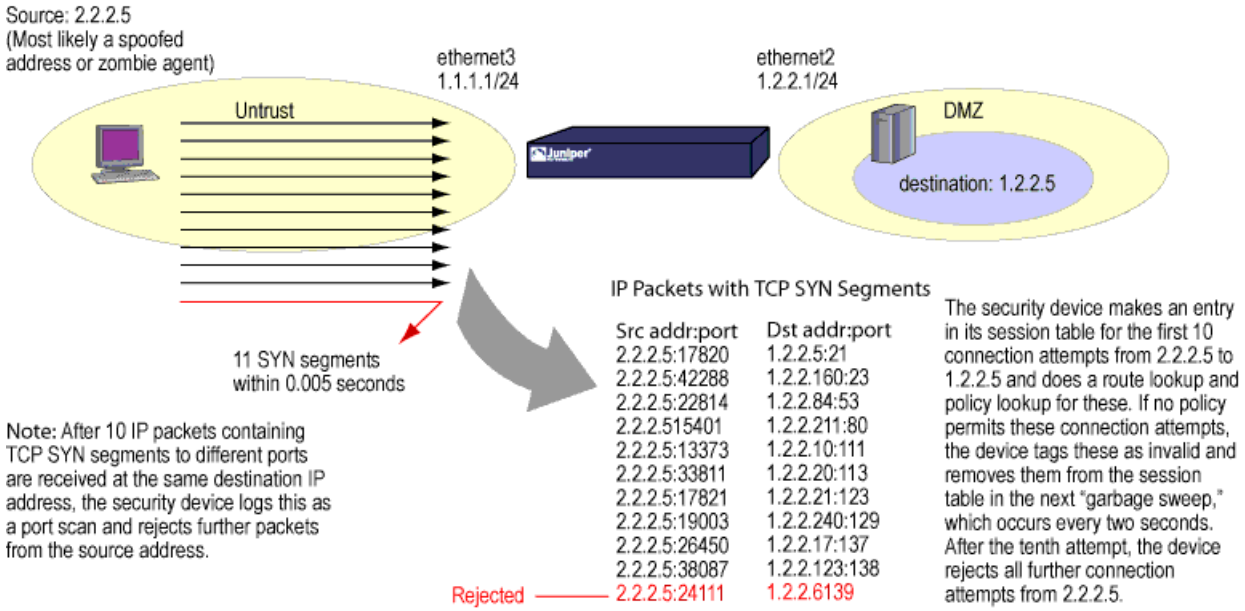


## *Related Topics*

- Blocking Port Scans on page 186
- Blocking IP Address Sweeps on page 183

## Blocking Port Scans

| Before You Begin |
| --- |
| For background information, read "Understanding Port Scanning" on page 184. |

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

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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

## 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 179
- Reconnaissance Deterrence Overview on page 181

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 30 on page 187. When they do appear, they are frequently being put to some illegitimate use.

**Figure 30: Routing Options**

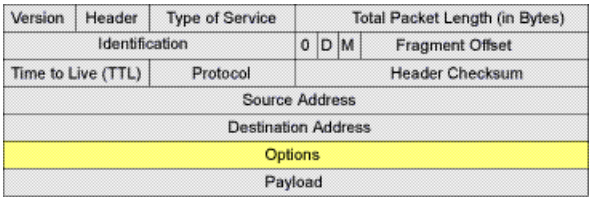


This topic covers:

- Uses for IP Packet Header Options on page 187
- SCREEN Options for Detecting IP Options Used For Reconnaissance on page 189
- Related Topics on page 189

### *Uses for IP Packet Header Options*

Table 50 on page 187 lists the IP options and their accompanying attributes.

**Table 50: 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 50: 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 206.) |
| 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 206.) |
| 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.<br><br>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 50: 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 189

## 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*

To configure screen and assign an Internet Protocol (IP) to it:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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-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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 179
- Reconnaissance Deterrence Overview on page 181
- Understanding Network Reconnaissance Using IP Options on page 187

This topic covers:



### *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 31 on page 192.

**Figure 31: TCP Header with SYN and FIN Flags Set**

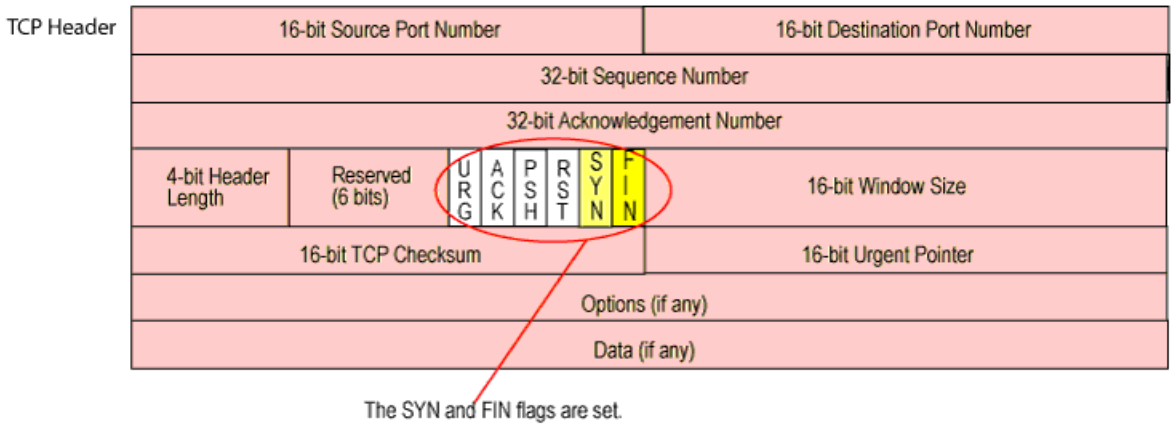


An attacker can send a segment with both flags set to see what kind of system reply is returned and thereby determine what kind of OS is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks.

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 32 on page 193 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 198.)

**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 32: TCP Header with FIN Flag Set**

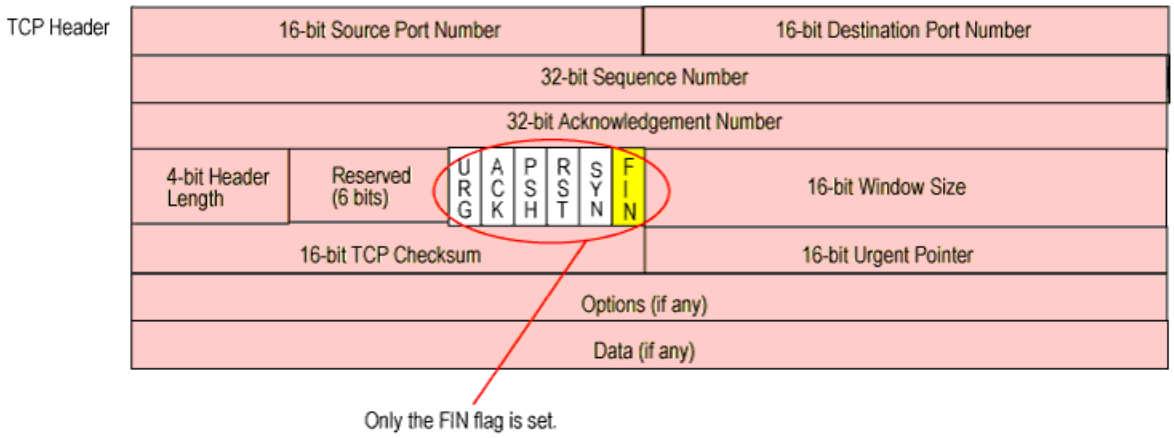


When you enable this SCREEN option, JUNOS software checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.

## *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 33 on page 193.

**Figure 33: TCP Header with No Flags Set**

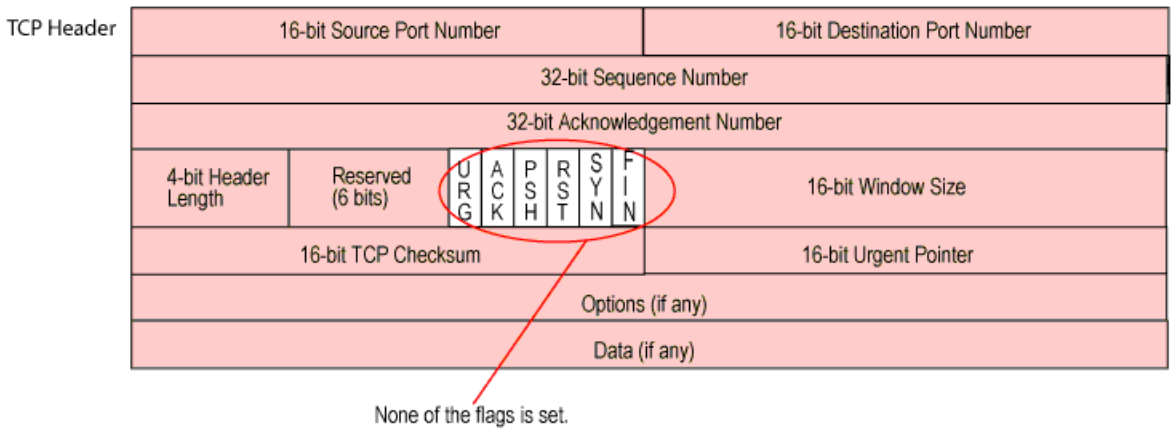


When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.

## *Related Topics*

- Blocking Packets with SYN and FIN Flags Set on page 194

- Blocking Packets with FIN Flag/No ACK Flag Set on page 195
- Blocking Packets with SYN and FIN Flags Set on page 194


## 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 191.

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 194
- CLI Configuration on page 195
- Related Topics on page 195


### *J-Web Configuration*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 195
- Blocking Packets with No Flags Set on page 196

## 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 191. |

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 195
- CLI Configuration on page 196
- Related Topics on page 196

### *J-Web Configuration*

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 194
- Blocking Packets with SYN and FIN Flags Set on page 194

## 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 191.

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 196
- CLI Configuration on page 197
- Related Topics on page 197

### J-Web Configuration

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 194
- Blocking Packets with FIN Flag/No ACK Flag Set on page 195

## 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 179
- Reconnaissance Deterrence Overview on page 181
- Understanding Operating System Probes on page 191

This topic covers:



## 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 34 on page 198 illustrates packet flow sequences when SYN flag checking is enabled and when it is disabled.

**Figure 34: SYN Flag Checking**

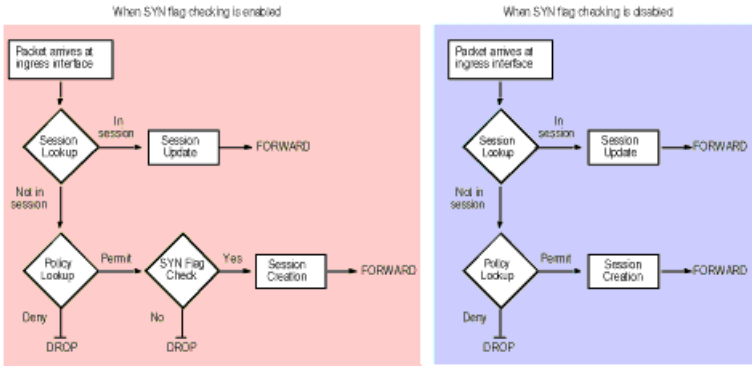


When JUNOS software with SYN flag checking enabled receives a non-SYN TCP segment that does not belong to an existing session, it drops the packet and sends the source host to a TCP RST—unless the code bit of the initial non-SYN TCP packet is also RST. In that case, JUNOS software simply drops the packet.

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 port Eth3, but JUNOS software has a route to 10.1.1.0/24 through port Eth1, checking 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 Eth1, not Eth3. 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 35 on page 201.

**Figure 35: IP Source Routing**

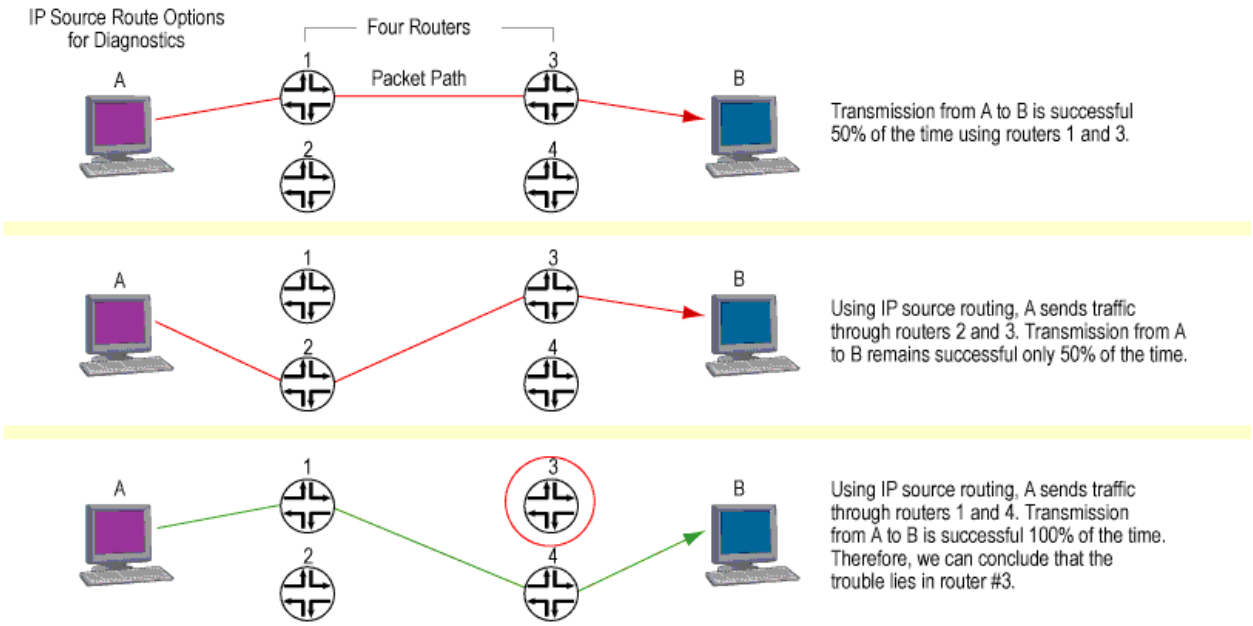


Although the uses of IP source route options were originally benign, attackers have learned to put them to more devious uses. They can use IP source route options to hide their true address and access restricted areas of a network by specifying a different path. For an example showing how an attacker can put both deceptions to use, consider the following scenario as illustrated in Figure 36 on page 201.

**Figure 36: Loose IP Source Route Option for Deception**

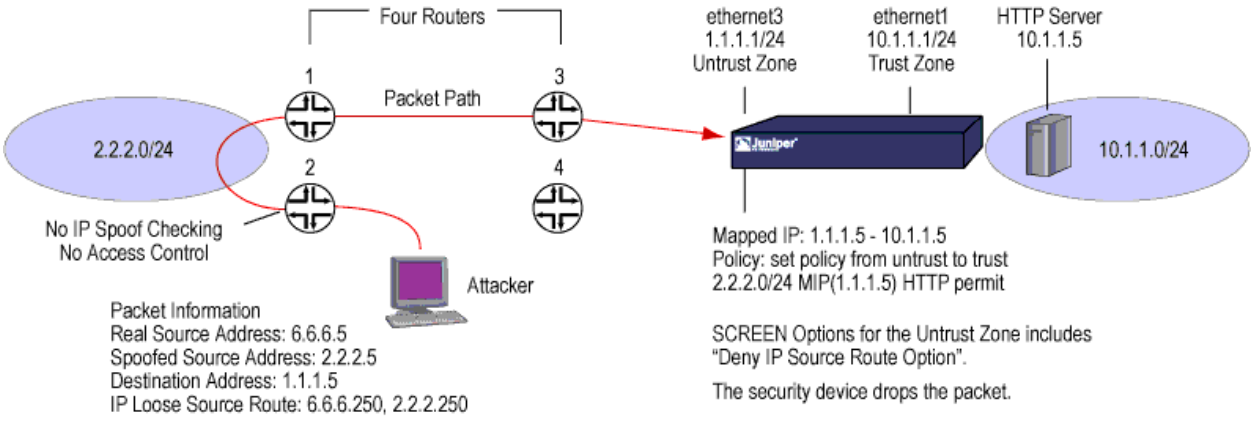


JUNOS software only allows traffic 2.2.2.0/24 if it comes through ethernet1, an interface bound to zone_external. Devices 3 and 4 enforce access controls but devices 1 and 2 do not. Furthermore, device 2 does not check for IP spoofing. The attacker spoofs the source address, and by using the loose source 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 202
- Setting TCP SYN Checking on page 203
- Blocking IP Spoofing on page 204
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 206
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 207


## 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 197. |

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 203
- Blocking IP Spoofing on page 204
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 206
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 207

## 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 197. |

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 204
- CLI Configuration on page 204
- Related Topics on page 204

### *J-Web Configuration*

To disable SYN checking:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 202
- Blocking IP Spoofing on page 204
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 206
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 207

## 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 197. |

You can use either J-Web or the CLI configuration editor to block IP spoofing.

This topic covers:

- J-Web Configuration on page 205
- CLI Configuration on page 205
- Related Topics on page 205

### *J-Web Configuration*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 202
- Setting TCP SYN Checking on page 203
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 206
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 207

## 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 197.

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*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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-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 202
- Setting TCP SYN Checking on page 203
- Blocking IP Spoofing on page 204
- Detecting Packets with Either a Loose or Strict Source Route Option Set on page 207

## 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 197. |

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 207
- CLI Configuration on page 208
- Related Topics on page 208

### *J-Web Configuration*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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-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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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-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 202
- Setting TCP SYN Checking on page 203
- Blocking IP Spoofing on page 204
- Blocking Packets with Either a Loose or Strict Source Route Option Set on page 206

## 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 179. |

All of the SCREEN options listed in "Related Topics" on page 209 block suspicious packets that might contain hidden threats.

### *Related Topics*

- Understanding ICMP Fragment Protection on page 209
- Understanding Large ICMP Packet Protection on page 211
- Understanding Bad IP Option Protection on page 214
- Understanding Unknown Protocol Protection on page 216
- Understanding IP Packet Fragment Protection on page 219
- Understanding SYN Fragment Protection on page 221

## 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 208. |

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 37 on page 210.

**Figure 37: Blocking ICMP Fragments**

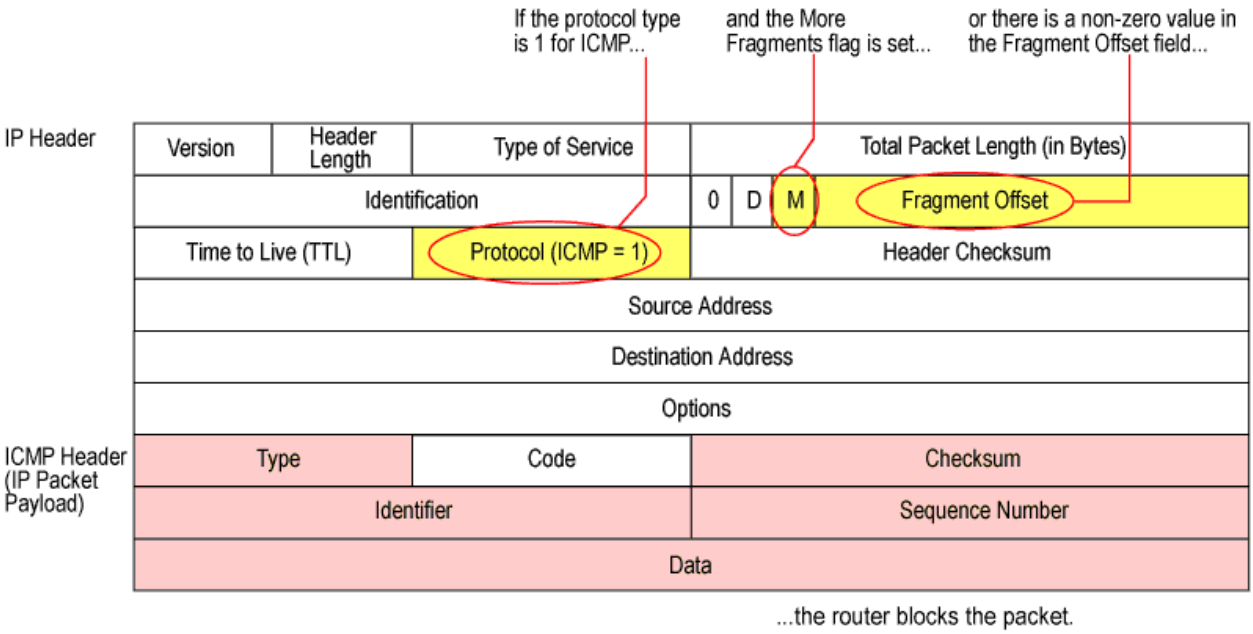


## *Related Topics*

- Blocking Fragmented ICMP Packets on page 210
- Understanding Large ICMP Packet Protection on page 211

## 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 209. |

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 211
- CLI Configuration on page 211
- Related Topics on page 211

### *J-Web Configuration*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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-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 213
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 215
- Dropping Packets Using an Unknown Protocol on page 217
- Dropping Fragmented IP Packets on page 220
- Dropping IP Packets Containing SYN Fragments on page 222

## 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 208.

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 38 on page 212.

**Figure 38: Blocking Large ICMP Packets**

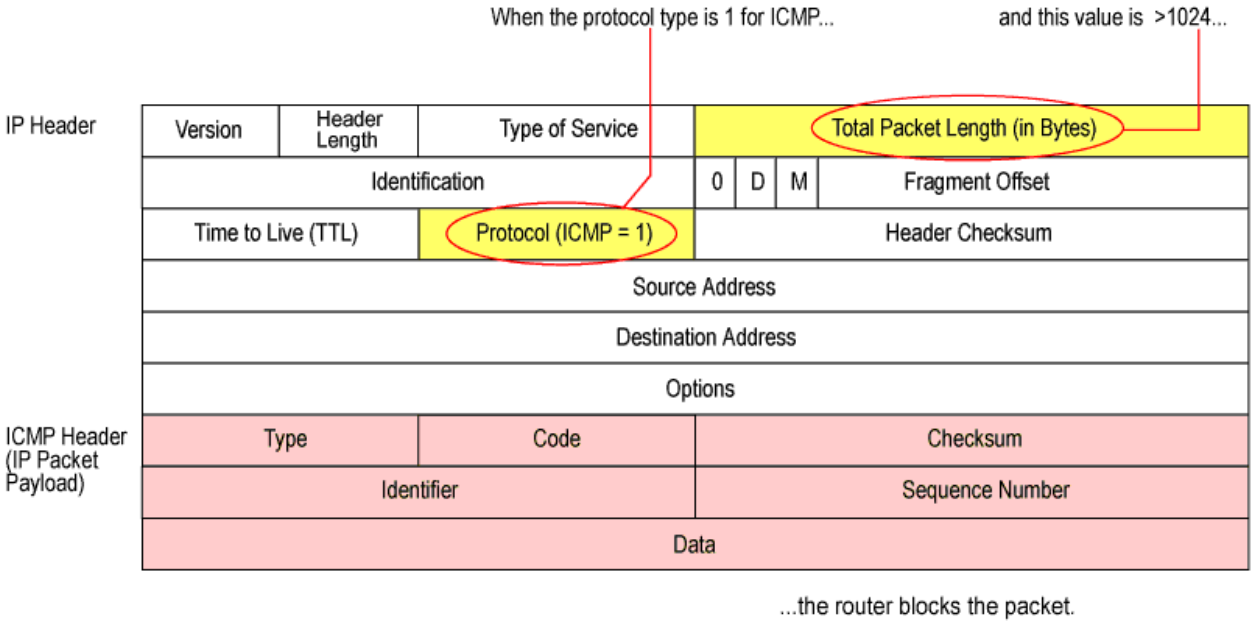


When you enable the Large Size ICMP Packet Protection SCREEN option, JUNOS software drops ICMP packets with a length greater than 1024 bytes.

### *Related Topics*

- Blocking Large ICMP Packets on page 213

## 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 211. |

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*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 210
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 215
- Dropping Packets Using an Unknown Protocol on page 217
- Dropping Fragmented IP Packets on page 220
- Dropping IP Packets Containing SYN Fragments on page 222

## 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 187.)

| **Before You Begin** |
| --- |
| For background information, read "Suspicious Packet Attributes Overview" on page 208. |

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 39 on page 215.

**Figure 39: Incorrectly Formatted IP Options**

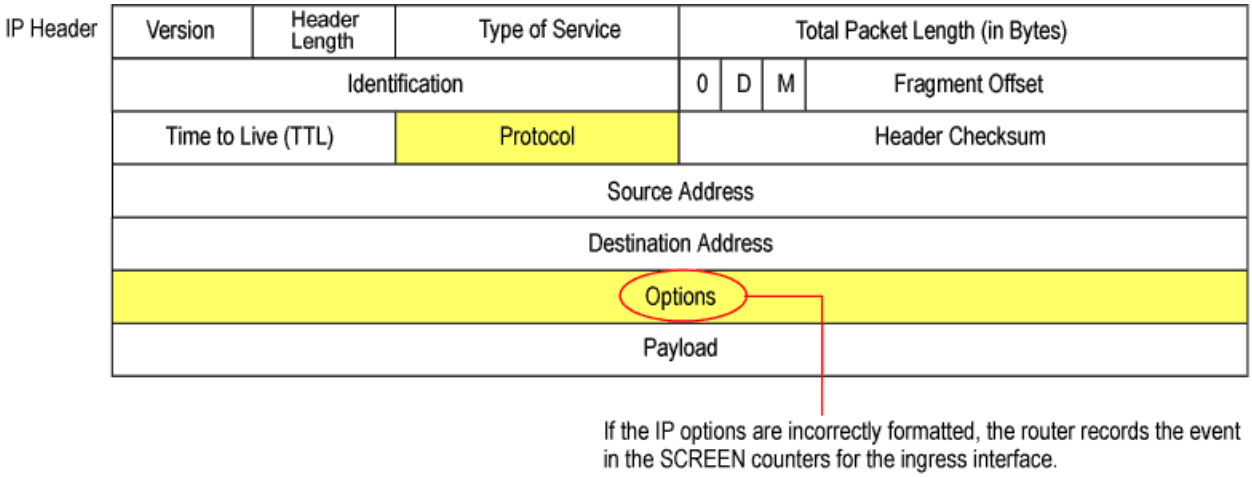


When you enable the bad IP option protection SCREEN option, JUNOS software blocks packets when any IP option in the IP packet header is incorrectly formatted. Additionally, JUNOS software records the event in the event log.

### *Related Topics*

- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 215

## 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 214. |

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 216
- CLI Configuration on page 216
- Related Topics on page 216

### J-Web Configuration

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 210
- Blocking Large ICMP Packets on page 213
- Dropping Packets Using an Unknown Protocol on page 217
- Dropping Fragmented IP Packets on page 220
- Dropping IP Packets Containing SYN Fragments on page 222

## 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 208. |

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 40 on page 217.

**Figure 40: Unknown Protocols**

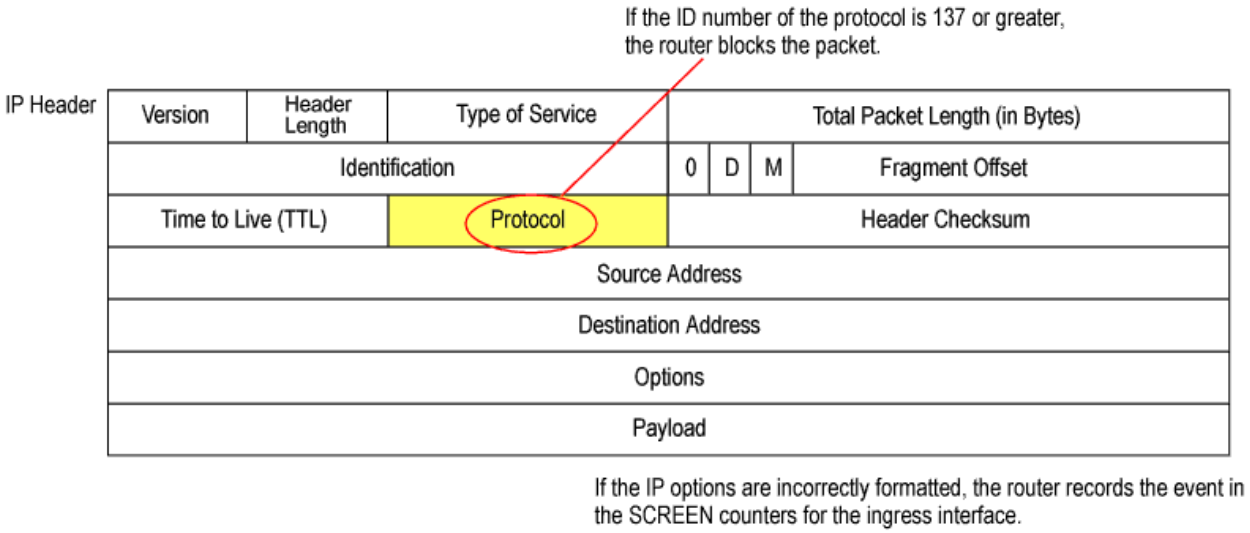


When you enable the unknown protocol protection SCREEN option, JUNOS software drops packets when the protocol field contains a protocol ID number of 137 or greater by default.

### *Related Topics*

- Dropping Packets Using an Unknown Protocol on page 217

## 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 216. |

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

To configure zones:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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

## 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 41 on page 219.

**Before You Begin**

For background information, read "Suspicious Packet Attributes Overview" on page 208.

**Figure 41: IP Packet Fragments**

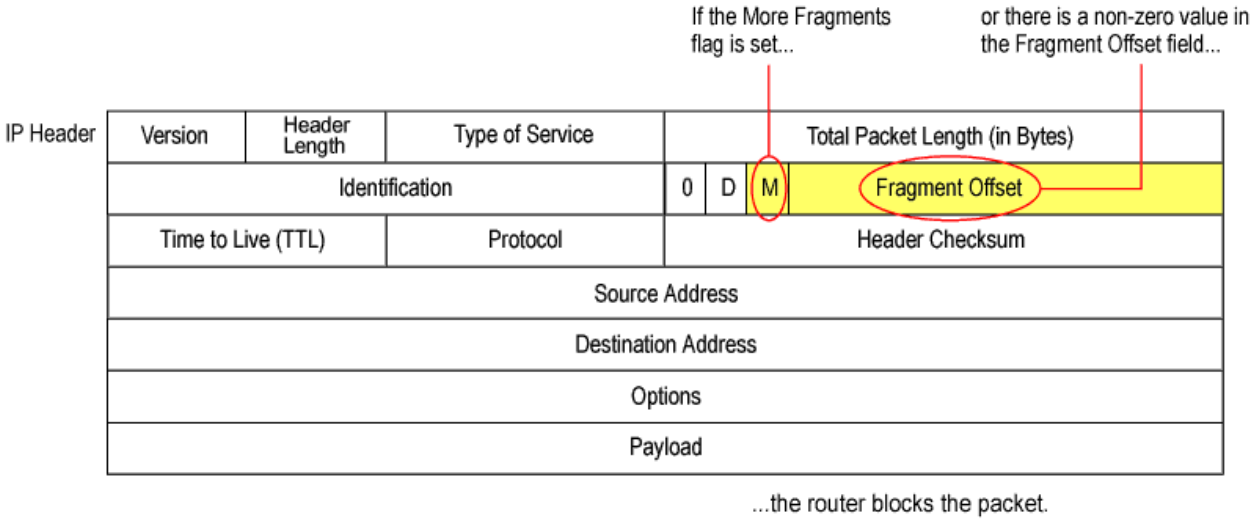


When you enable JUNOS software to deny IP fragments on a security zone, it blocks all IP packet fragments that it receives at interfaces bound to that zone.

### *Related Topics*

- "Dropping Fragmented IP Packets" on page 220

## 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 219. |

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*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 210
- Blocking Large ICMP Packets on page 213
- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 215
- Dropping Packets Using an Unknown Protocol on page 217
- Dropping IP Packets Containing SYN Fragments on page 222

## 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 208. |

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 42 on page 222.

**Figure 42: SYN Fragments**

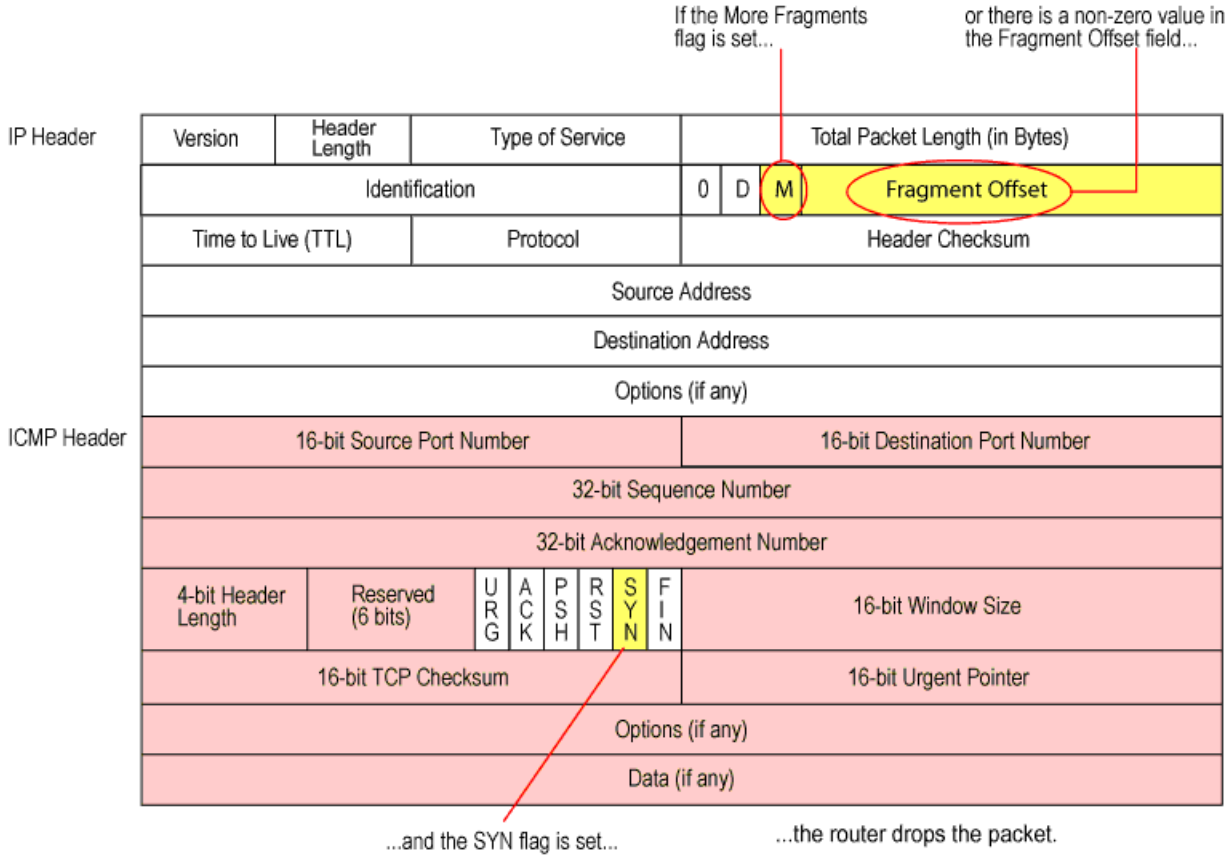


When you enable the SYN fragment detection SCREEN option, JUNOS software detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is set in the TCP header. JUNOS software records the event in the SCREEN counters list for the ingress interface.

### *Related Topics*

- Dropping IP Packets Containing SYN Fragments on page 222

## 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 221. |

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

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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

- Detecting and Blocking IP Packets with Incorrectly Formatted Options on page 215
- Dropping Packets Using an Unknown Protocol on page 217
- Dropping Fragmented IP Packets on page 220

## 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 179. |

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 224
- Network DoS Attacks Overview on page 233
- OS-Specific DoS Attacks Overview on page 257

## 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 224. |

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 225
- Understanding SYN-ACK-ACK Proxy Flood Attacks on page 230
- Network DoS Attacks Overview on page 233
- OS-Specific DoS Attacks Overview on page 257

## 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 224. |

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 225
- Destination-Based Session Limits on page 226
- Related Topics on page 227

### *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 43 on page 226.

**Figure 43: Limiting Sessions Based on Source IP Address**

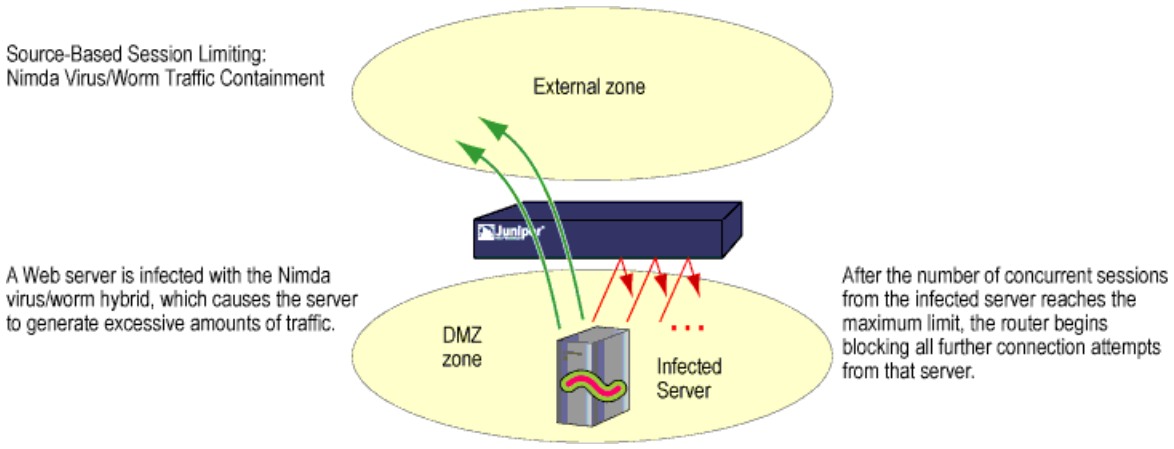


Another benefit of source-based session limiting is that it can mitigate attempts to fill up the firewall's session table, if all the connection attempts originate from the same source IP address.

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 44 on page 227.

**Figure 44: Distributed DOS Attack**

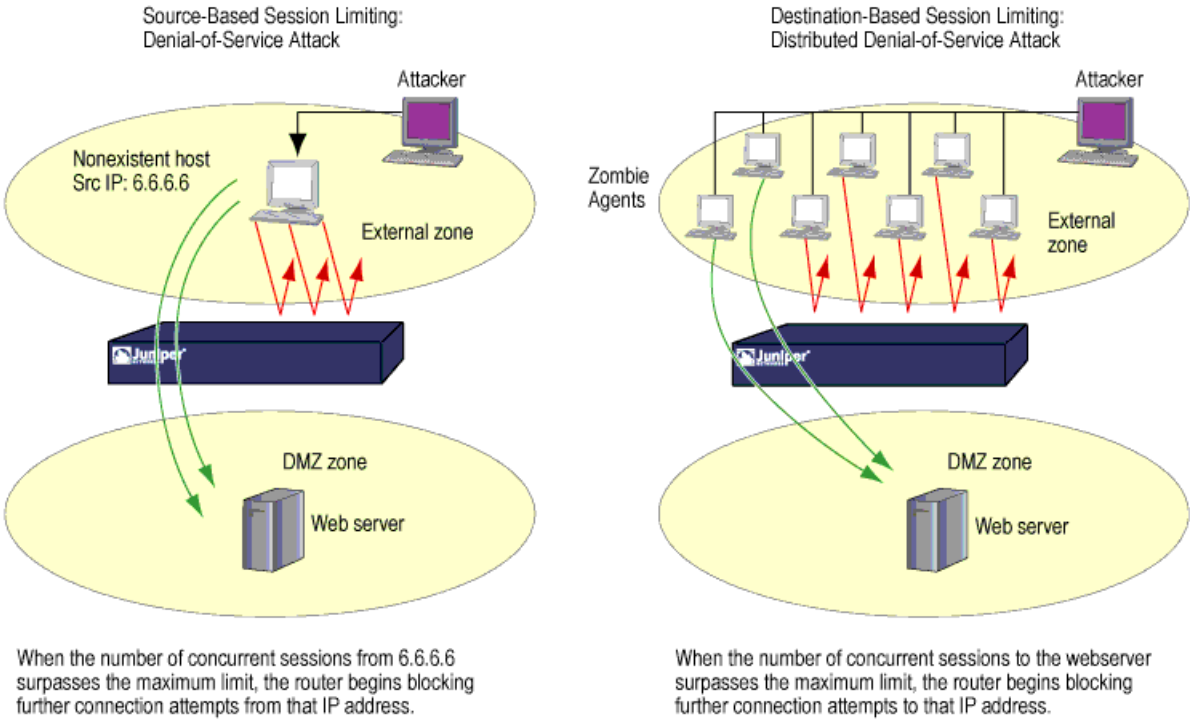


The default maximum for destination-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

### *Related Topics*

- Setting Source-Based Session Limits on page 227
- Setting Destination-Based Session Limits on page 229

## 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 225. |

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

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 229
- Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack on page 231

## 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 225. |

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 229
- CLI Configuration on page 230
- Related Topics on page 230

### J-Web Configuration

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 227
- Enabling Protection Against a SYN-ACK-ACK Proxy Flood Attack on page 231

## 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 224. |

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 231

## 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 230. |

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 232
- CLI Configuration on page 232
- Related Topics on page 232

### J-Web Configuration

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 227
- Setting Destination-Based Session Limits on page 229

## 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 224.

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 233
- Understanding SYN Cookie Protection on page 245
- Understanding ICMP Flood Attacks on page 249
- Understanding UDP Flood Attacks on page 252
- Understanding Land Attacks on page 254

## 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 233.

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 45 on page 234.

**Figure 45: SYN Flood Attack**

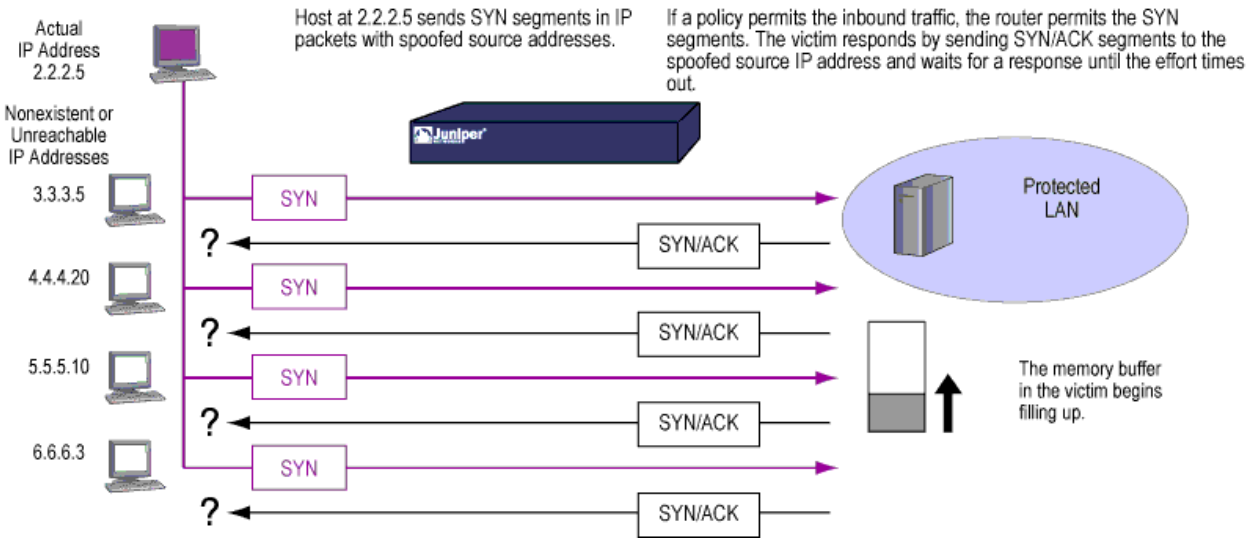


By flooding a host with incomplete TCP connections, the attacker eventually fills the memory buffer of the victim. Once this buffer is full, the host can no longer process new TCP connection requests. The flood might even damage the victim's operating system. Either way, the attack disables the victim and its normal operations.

This topic covers:



## *SYN Flood Protection*

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 46 on page 235, the SYN attack threshold has passed, and JUNOS software has started proxying SYN segments.

**Figure 46: Proxying SYN Segments**

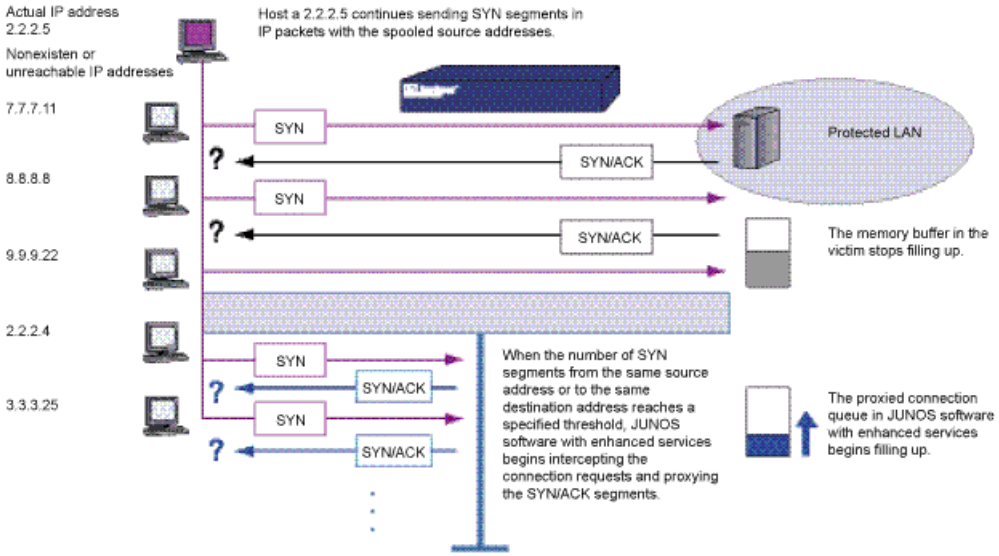


In Figure 47 on page 236, the proxied connection queue has completely filled up, and JUNOS software is rejecting new incoming SYN segments. This action shields hosts on the protected network from the bombardment of incomplete three-way handshakes.

**Figure 47: Rejecting New SYN Segments**

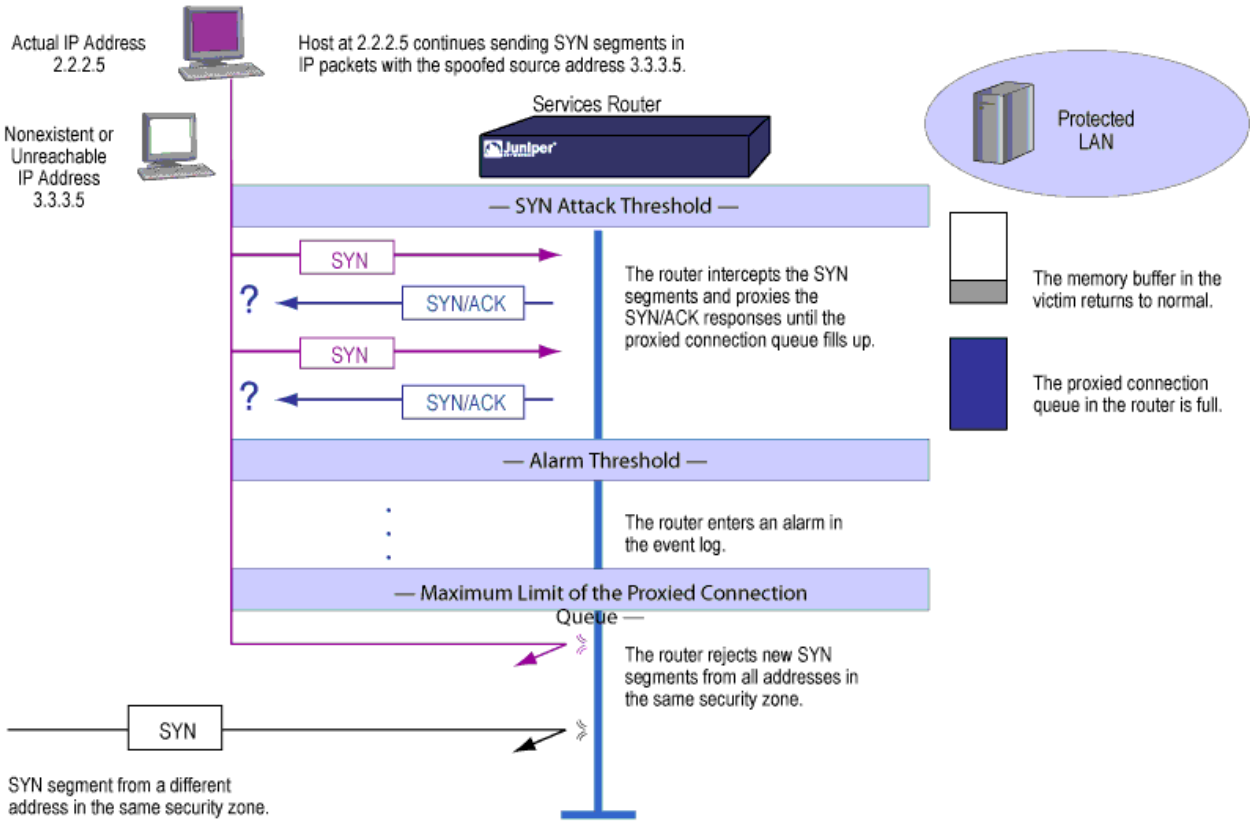


The device starts receiving new SYN packets when the proxy queue drops below the maximum limit.

**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 245
- Enabling SYN Cookie Protection on page 247
- Enabling SYN Flood Protection on page 245

## 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 233. |

**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 48: Device-Level SYN Flood Protection**

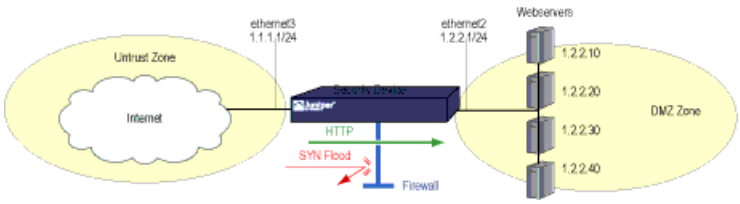


To configure the SYN flood protection parameters with appropriate values for your network, you must first establish a baseline of typical traffic flows. For one week, you run a sniffer on ethernet3—the interface bound to zone_external—to monitor the number of new TCP connection requests arriving every 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 51 on page 240.

**Table 51: 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.)<br><br>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.)<br><br>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. |
| 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.) |

**Table 51: SYN Flood Protection Parameters** *(continued)*

| Parameter | Value | Reason for Each Value |
|---|---|---|
| * 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 241
- CLI Configuration on page 244
- Related Topics on page 245


## *J-Web Configuration*

To set interfaces:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Zones, click Configure.
3. Next to Security zone, click Add new entry.

4. In the Name box, type **zone_external**.
5. In the Screen box, type **zone_external-syn-flood** and click OK.
6. 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 245

## 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 233.

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 247
- Enabling ICMP Flood Protection on page 250
- Enabling UDP Flood Protection on page 253
- Enabling Protection Against a Land Attack on page 255

## 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 233
- Understanding SYN Flood Attacks on page 233

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 49 on page 247 shows how a connection is established between an initiating host and a server when SYN Cookie is active on JUNOS software.

**Figure 49: Establishing a Connection with SYN Cookie Active**

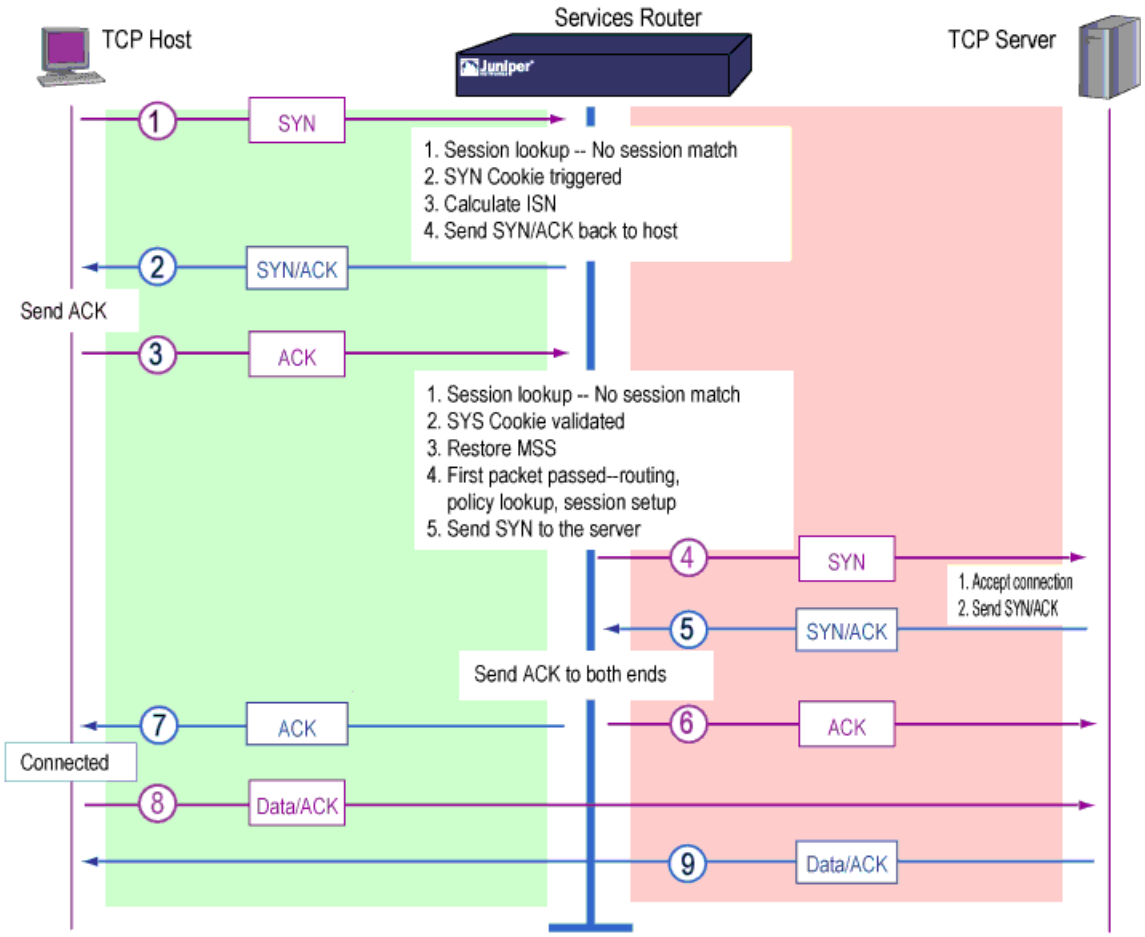


### *Related Topics*

- Enabling SYN Cookie Protection on page 247

## 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 245. |

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

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 245
- Enabling ICMP Flood Protection on page 250
- Enabling UDP Flood Protection on page 253
- Enabling Protection Against a Land Attack on page 255

## 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 233. |

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 50 on page 250.

**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 50: ICMP Flooding**

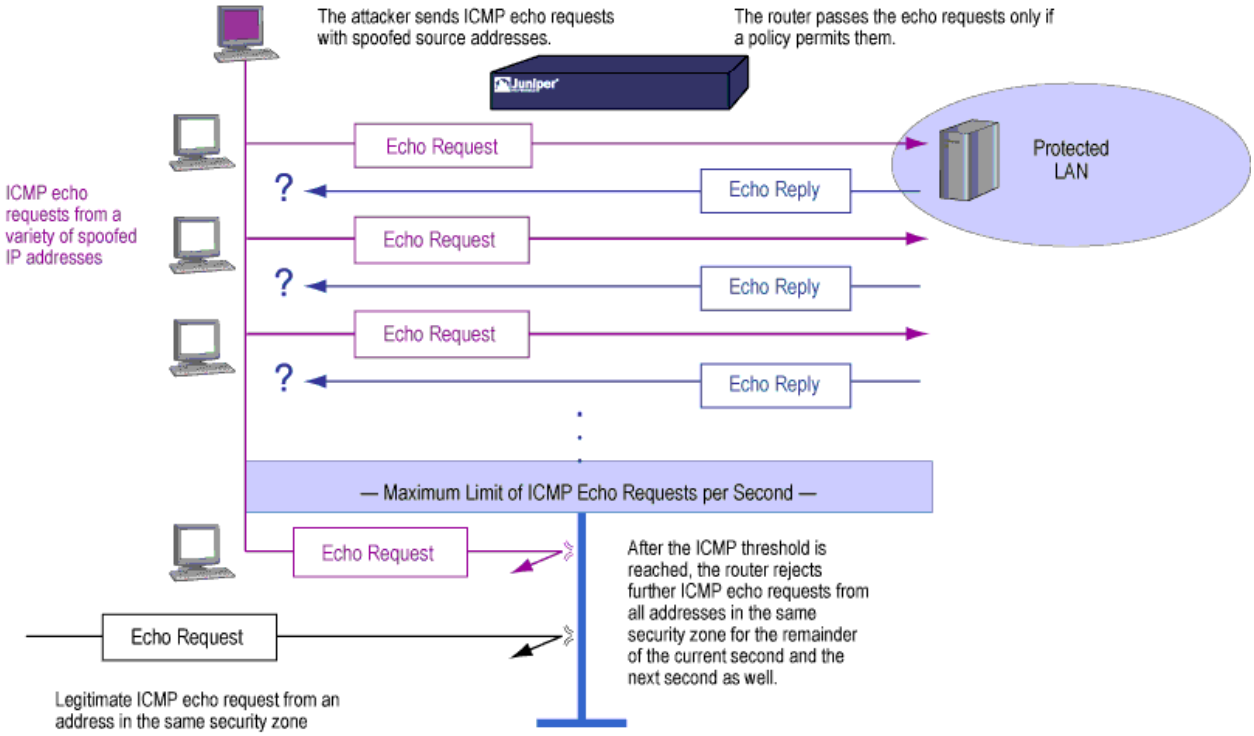


### *Related Topics*

- Enabling ICMP Flood Protection on page 250

## 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 249. |

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

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 UDP Flood Protection on page 253
- Enabling Protection Against a Land Attack on page 255

## 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 233
- Understanding ICMP Flood Attacks on page 249

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 51 on page 252.

**Figure 51: UDP Flooding**

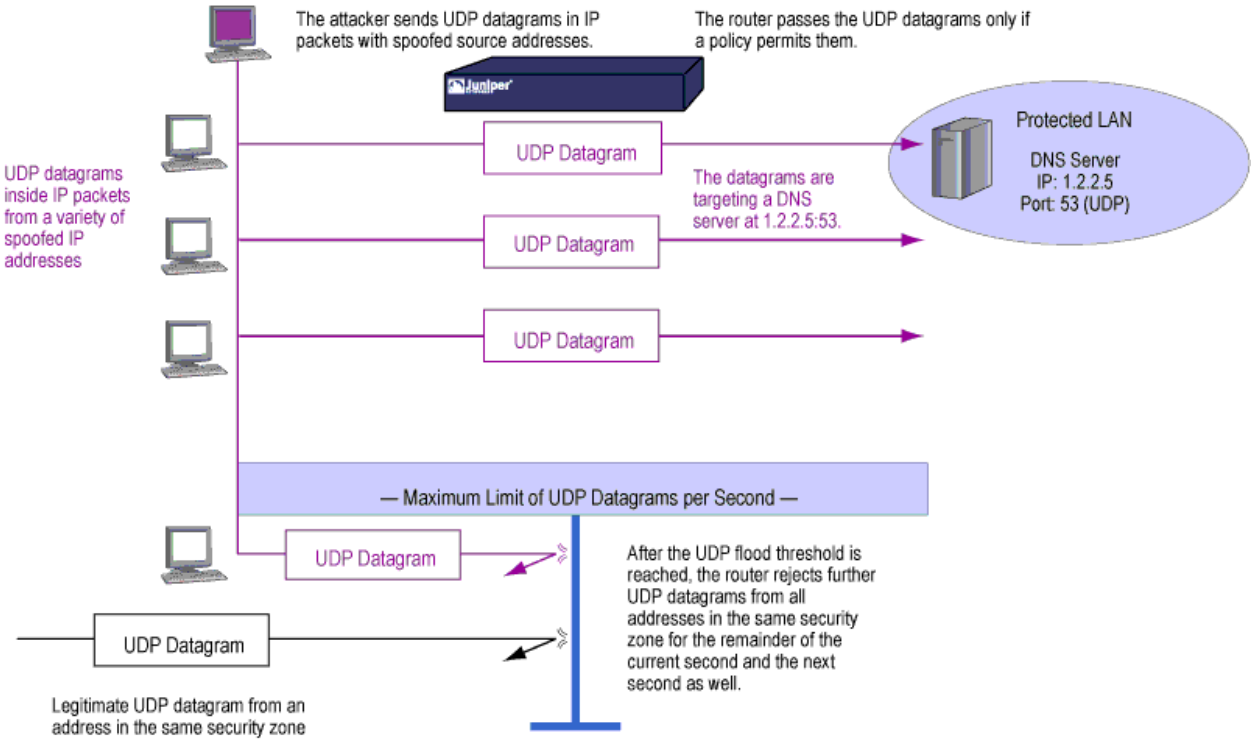

### *Related Topics*

- Enabling Protection Against a Land Attack on page 255

## 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 252. |

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 253
- CLI Configuration on page 254
- Related Topics on page 254

### *J-Web Configuration*

To configure zones:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 screen:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 245
- Enabling SYN Cookie Protection on page 247
- Enabling ICMP Flood Protection on page 250
- Enabling Protection Against a Land Attack on page 255

## 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 233
- IP Spoofing on page 200

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 52 on page 255.

**Figure 52: Land Attack**

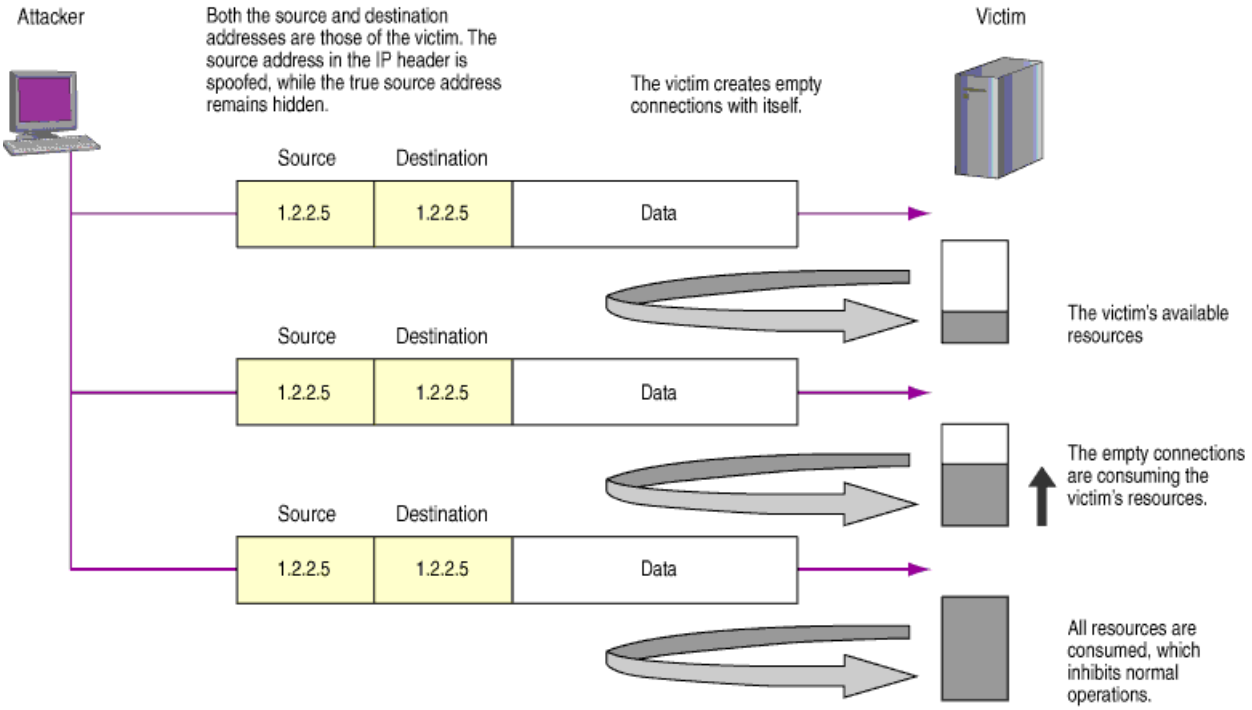


When you enable the SCREEN option to block land attacks, JUNOS software combines elements of the SYN flood defense and IP spoofing protection to detect and block any attempts of this nature.

### *Related Topics*

- Enabling Protection Against a Land Attack on page 255

## 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 254. |

You can use either J-Web or the CLI configuration editor to enable protection against a land attack.

This topic covers:



## J-Web Configuration

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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

## 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 224.

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 257
- Understanding Teardrop Attacks on page 259
- Understanding WinNuke Attacks on page 262

## 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 257.

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 53 on page 258.

**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 53: Ping of Death**

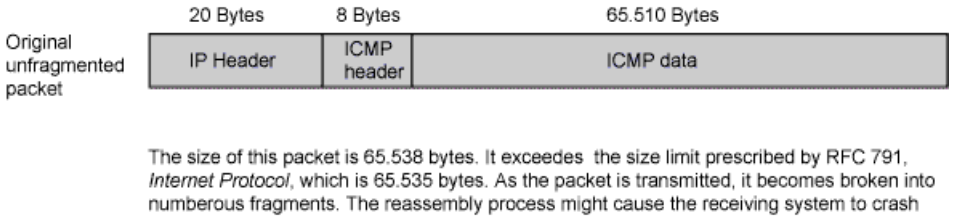


### *Related Topics*

- Enabling Protection Against a Ping of Death Attack on page 258

## 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 257. |

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 258
- CLI Configuration on page 259
- Related Topics on page 259

### *J-Web Configuration*

To configure screen:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 261
- Enabling Protection Against a WinNuke Attack on page 263

## 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 257. |

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 54 on page 260.

**Figure 54: Teardrop Attacks**

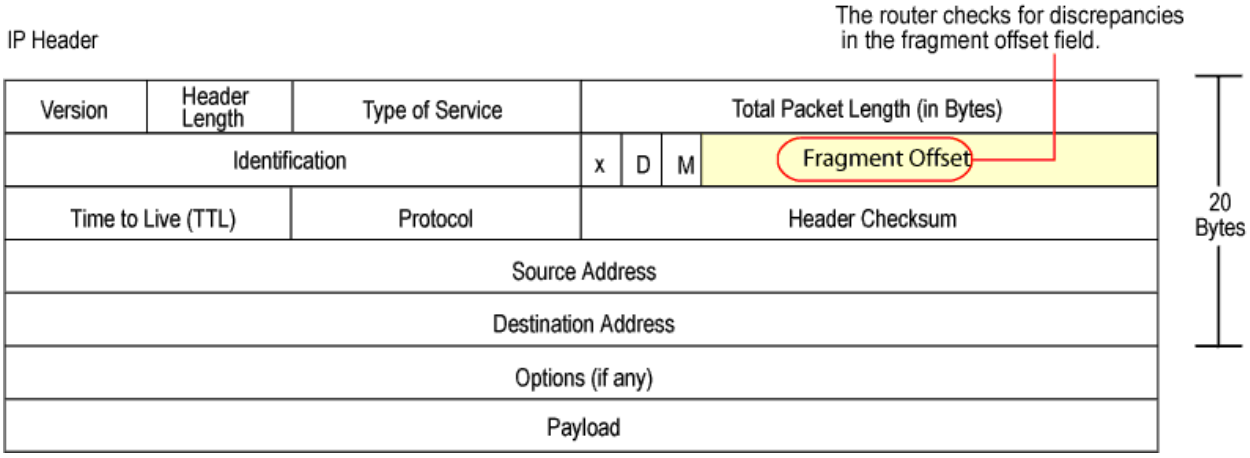


When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash, especially if it is running an older operating system that has this vulnerability. See Figure 55 on page 260.

**Figure 55: Fragment Discrepancy**

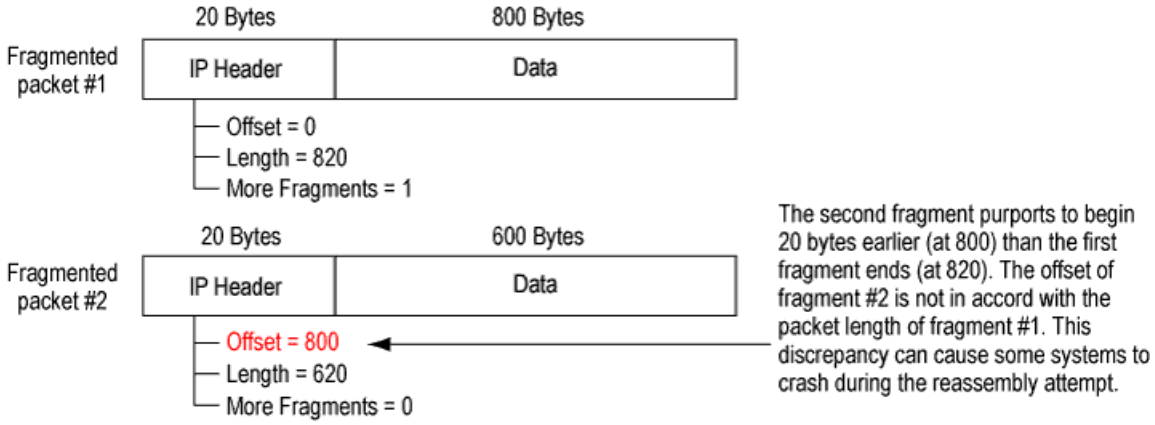


After you enable the teardrop attack SCREEN option, whenever JUNOS software detects this discrepancy in a fragmented packet, it drops it.

### *Related Topics*

- Enabling Protection Against a Teardrop Attack on page 261

## 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 259. |

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 261
- CLI Configuration on page 262
- Related Topics on page 262

### *J-Web Configuration*

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 258
- Enabling Protection Against a WinNuke Attack on page 263

## 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 257. |

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 56 on page 263). 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 56: WinNuke Attack Indicators**

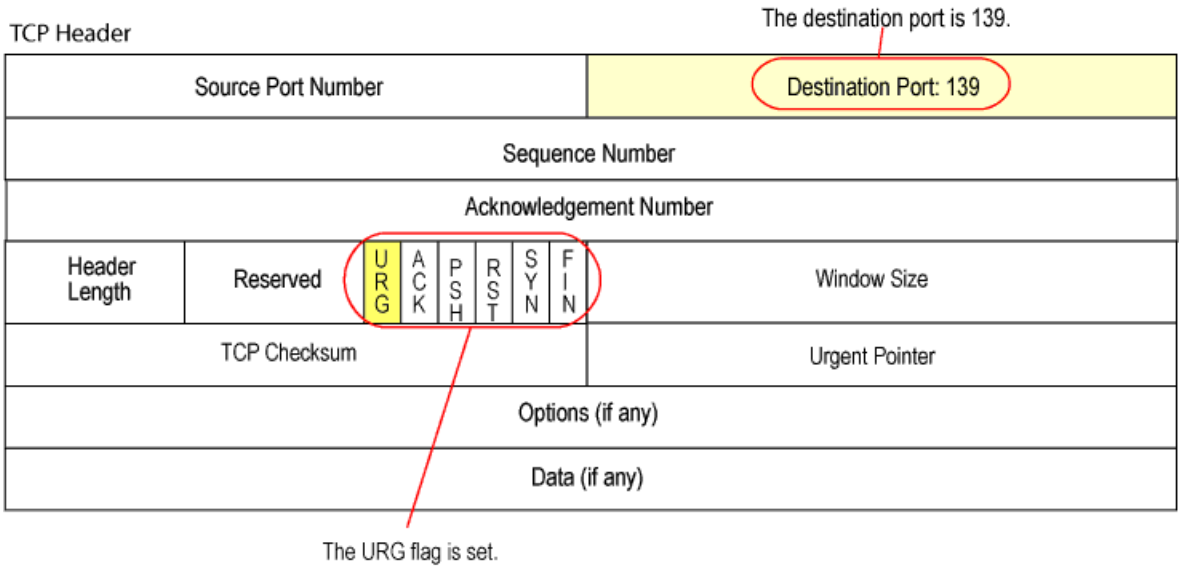


If you enable the WinNuke attack defense SCREEN option, JUNOS software scans any incoming Microsoft NetBIOS session service (port 139) packets. If JUNOS software observes that the URG flag is set in one of those packets, it unsets the URG flag, clears the URG pointer, forwards the modified packet, and 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 263

## 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 262. |

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 264
- CLI Configuration on page 264
- Related Topics on page 264

### J-Web Configuration

To configure screens:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 258
- Enabling Protection Against a Teardrop Attack on page 261

## 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 55
- "Reconnaissance Deterrence Overview" on page 181
- "Understanding Operating System Probes" on page 191
- "Understanding Attacker Evasion Techniques" on page 197

Figure 57 on page 265 shows the Quick Configuration for Firewall/NAT Screen page.

**Figure 57: Quick Configuration Page for Firewall/NAT Screen**

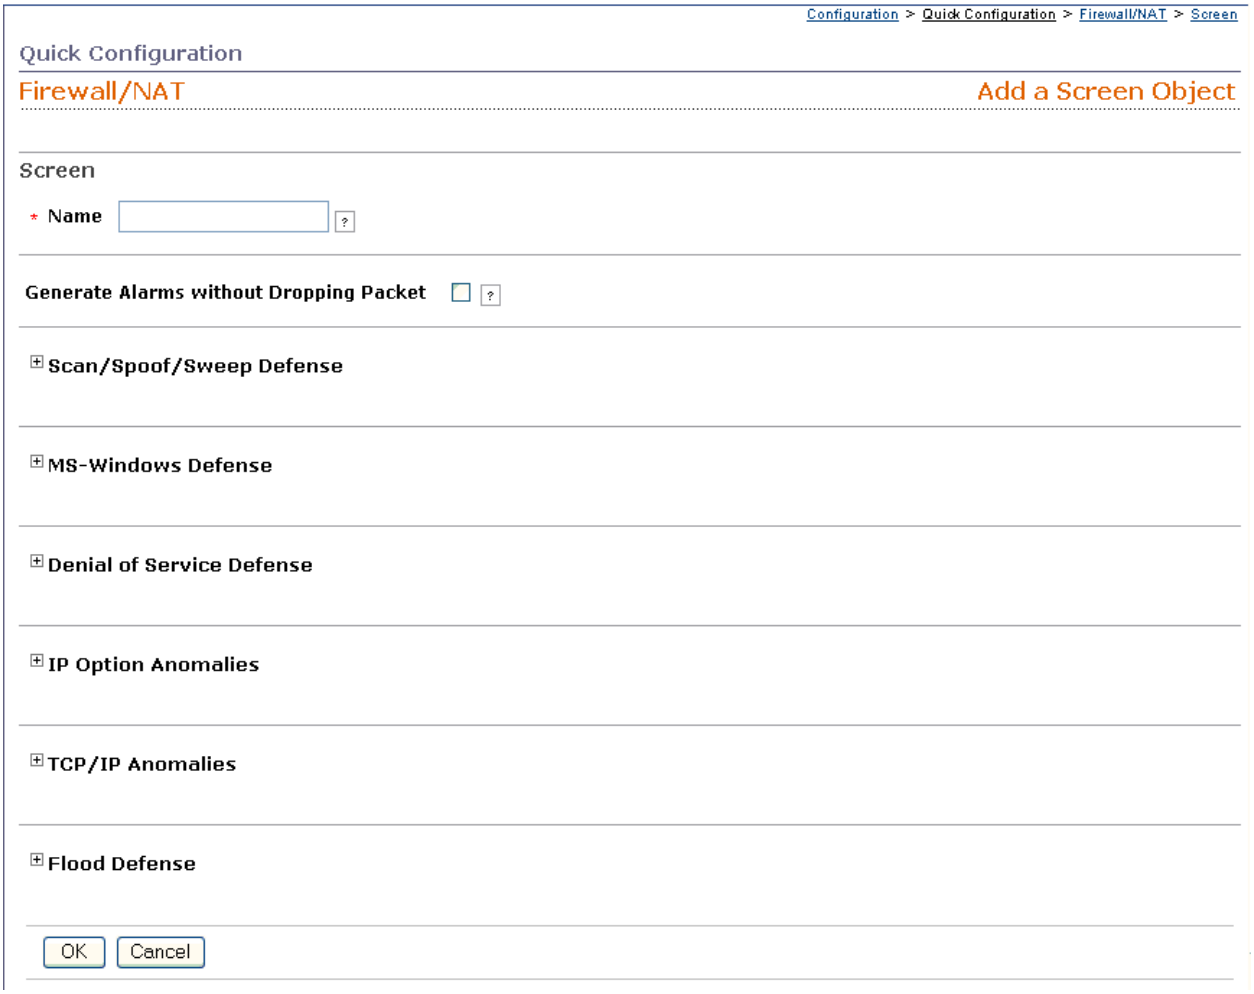

To configure screen options with Quick Configuration:

1. Select Configuration > Quick Configuration > Firewall/NAT > Screen.
2. Click Add to define screen objects; the screen objects page appears as shown in Figure 57 on page 265.
3. Fill in the Screen options as shown in Table 52 on page 266.
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 52: 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.<br><br>Configure a time threshold (in microseconds) per 10 ICMP packets. Valid values are between 1000 and 1000000 packets per micro second. The default value is 5000 ppms. |
| 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.<br><br>Configure a time threshold (in microseconds) per 10 attack packets. Valid values are between 1000 and 1000000 packets per micro second. The default value is 5000 ppms. |
| **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. |
| **Denial of Service Defense** | | |

**Table 52: Firewall/NAT Screen Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| 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.<br><br>Configure the threshold between 1 and 50000 sessions. The default value is 128 sessions.<br><br>**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.<br><br>Configure the threshold between 1 and 50000 sessions. The default value is 128 sessions.<br><br>**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 Enhanced Services rejects further connection requests from that IP address. | Select this check box to enable the SYN-ACK-ACK proxy protection SCREEN option.<br><br>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 52: 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 52: 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.<br><br>Configure threshold value for ICMP flood between 1 and 100000 ICMP packets per second (pps).<br><br>The default value is 1000 pps.<br><br>**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.<br><br>Configure threshold value for UDP flood between 1 and 100000 UDP packets with same destination address per second (pps).<br><br>The default value is 1000 pps.<br><br>**NOTE:** For SRX-series devices, the applicable range is 1 through 4000000 UDP packets per second. |

**Table 52: 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.<br><br>Alarm Threshold—Define the number of half-complete proxy connections per second at which the device makes entries in the event alarm log.<br><br>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.<br><br>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.<br><br>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.<br><br>**NOTE:** For SRX-series devices, the applicable range is 1 through 1000000 proxied requests per second.<br><br>Alarm Threshold—Configure a value between 1 and 100000 segments received per second for SYN flood alarm. The default value is 512.<br><br>**NOTE:** For SRX-series devices, the applicable range is 1 through 1000000 segments per second.<br><br>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.<br><br>**NOTE:** For SRX-series devices, the applicable range is 4 through 1000000 segments per second.<br><br>Destination Threshold—Configure a value for SYN flood to the same destination between 4 and 100000. The default value is 4000.<br><br>**NOTE:** For SRX-series devices, the applicable range is 4 through 1000000 segments per second.<br><br>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

1. Select Configuration > View and Edit > Edit Configuration > Security > Traceoptions Configure. The Configuration page appears.
2. Select the No remote trace checkbox to disable remote tracing.
3. 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.
4. Select the Use local files checkbox to write trace messages to a local file. The trace file is saved in the /var/log/directory.
5. 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.”
6. 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.
7. 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.
8. 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.
9. Select either the Yes or the No Word readable checkbox. If you select Yes, you are allowing any user to read the trace file.
10. 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

11. Click OK to return to the configuration page.
12. 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.
13. 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
  ```

Chapter 12

# 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.

This section includes:

## Understanding NAT

NAT can include one of the following attributes with or without port address translation:

- Destination IP address translation
- Source IP address translation
- Both

When a policy configuration includes Network Address Translation (NAT) in its match criteria, the J-series Services Router, acting like a Layer 3 switch (or router), translates two components in the header of an outgoing IP packet destined for the external zone: its source IP address and source port number. The router replaces the source IP address of the originating host with the IP address of the external zone interface. Also, it replaces the source port number with another random port number generated by the router. See Figure 58 on page 277.

**Figure 58: NAT Topology**

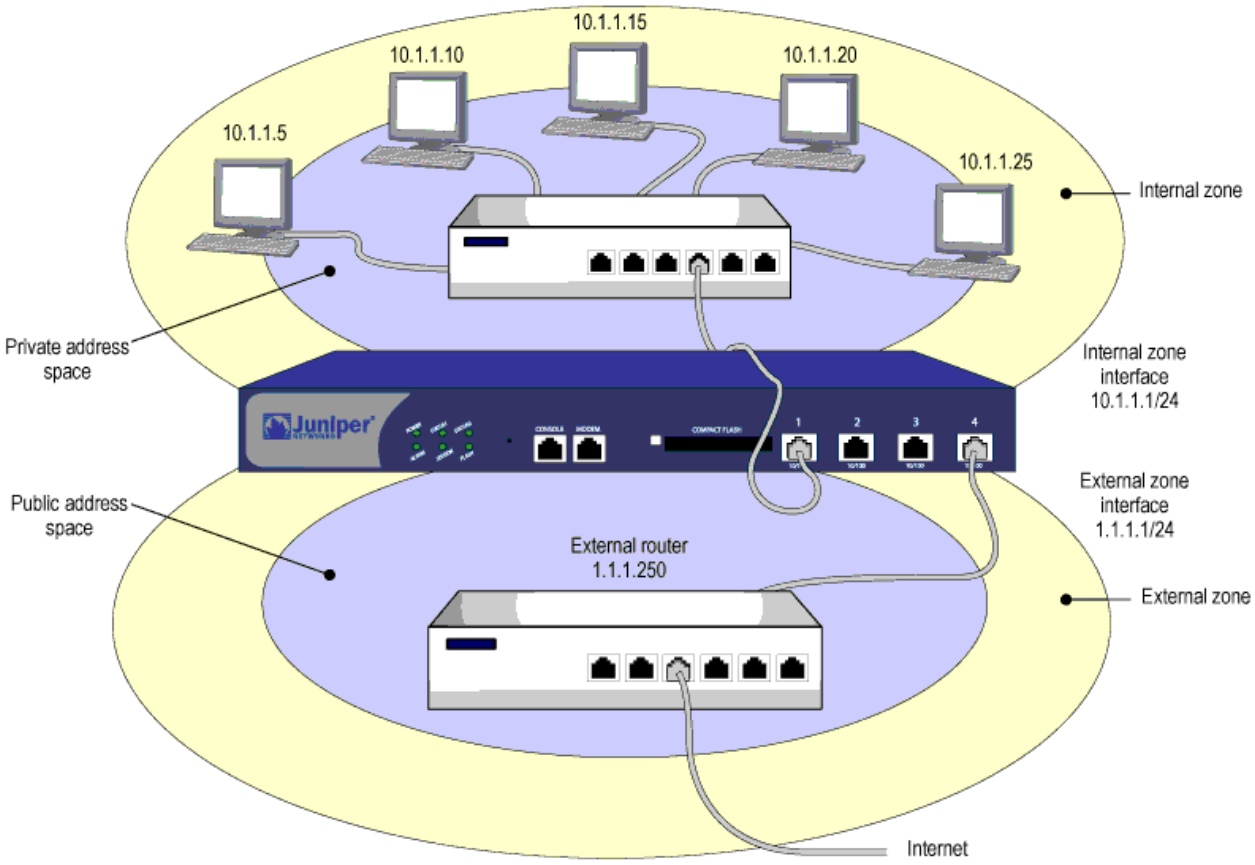


When the reply packet arrives at the device, the device translates two components in the IP header of the incoming packet (the destination address and port number) which are translated back to the original numbers. The device then forwards the packet to its destination.

The addresses of hosts sending traffic through the interface in an internal zone (where NAT is configured) are never exposed to hosts in the external zone, unless the two zones are in the same virtual routing domain and the device is advertising routes to peers through a Dynamic Routing Protocol (DRP). Even then, the internal zone addresses are only reachable if you have a policy permitting inbound traffic to them. (If you want to keep the internal zone addresses hidden while using a DRP, then put the external zone in the external-vr and the internal zone in the internal-vr, and do not export routes for internal addresses in the internal-vr to the external-vr.)

Also, NAT preserves the use of public IP addresses. In many environments, resources are not available to provide public IP addresses for all devices on the network. NAT services allow many private IP addresses to have access to Internet resources through one or a few public IP addresses. The following IP address ranges are reserved for private IP networks and must not get routed on the Internet:

- 10.0.0.0 — 10.255.255.255
- 172.16.0.0 — 172.31.255.255

- 192.168.0.0 — 192.168.255.255

NAT can involve either destination IP address translation or source IP address translation, or both, with or without port address translation.

This topic covers:



### *Inbound and Outbound NAT Traffic*

A host in a zone sending traffic through an interface in NAT mode can initiate traffic to the external zone—assuming that a policy permits it. However, if issues of privacy and private IP addresses are not a concern, traffic from the external zone can reach hosts behind an interface in NAT mode directly, without the use of a VPN. See Figure 59 on page 278.

**Figure 59: NAT Traffic Flow**

1. Interface-based NAT on traffic from an internal zone to an external zone.
2. Interface-based NAT on traffic from the user-defined zone to the external zone.

   (Note: This is possible only if the user-defined and external zones are in different virtual routing domains.)
3. No interface-based NAT on traffic between the internal and user-defined zones.

4 and 5. You can use VPNs for traffic from the external zone to reach the internal zone or the user-defined zone, but they are not required.

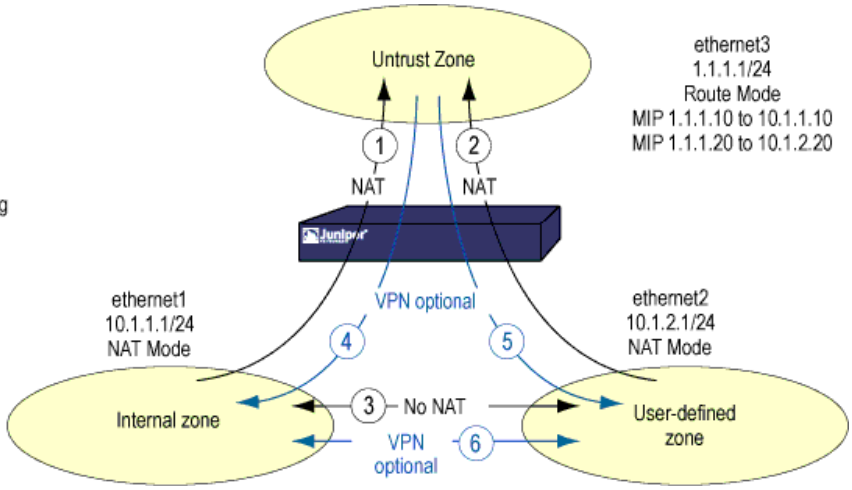


### *Related Topics*



## NAT Configuration on Different Devices

Although NAT functionality is conceptually similar on J-series Service Routers and SRX-series services gateways, there are some configuration differences to note. On J-series devices, NAT is tied closely to policies. On SRX-series devices, NAT is not

tied to policies and NAT configuration includes NAT-specific rules. Also, on SRX-series devices only, you must explicitly configure Proxy (Address Resolution Protocol) ARP.

The CLI configuration editors use slightly different commands for NAT on each device type. The following sections provide instructions for configuring NAT on both J-series Service Routers and SRX-series services gateways.

## Destination IP Address Translation Overview

When performing destination IP address translation, JUNOS software translates the original destination IP address, port number, or both to a different one.

| Before You Begin |
| --- |
| For background information, read "Understanding NAT" on page 276. |

Destination NAT can be performed in one of the following ways:

- Static NAT defined on an interface
- Policy-based NAT-dst defined within a policy on J-series Service Routers
- Rule-based destination NAT defined on SRX-series services gateways
- Allow-incoming table dynamically generated by VoIP ALGs

### *Related Topics*

- Understanding Static NAT on J-series Services Routers on page 279
- Understanding NAT-Dst Policy-Based NAT on J-series Services Routers on page 283
- Understanding NAT-Dst Allow-Incoming Table on page 287
- Understanding Rule-Based Destination NAT on SRX-series Services Gateways on page 285

## Understanding Static NAT on J-series Services Routers

Static NAT is a direct one-to-one mapping of one IP address to another without port address translation. The mapping is done by a single IP address or by a subnet.

| Before You Begin |
| --- |
| For background information, read "Destination IP Address Translation Overview" on page 279. |

When static NAT is set for an interface in any zone, an entry for the static NAT is generated in the junos-global zone address book in the following format: static_nat_w.x.y.z_n, where w.x.y.z_n is the static NAT. The junos-global zone address

book stores all the static NAT addresses, regardless of the zone to which their interface belongs. You can use these static NAT addresses as destination addresses in policies between any two zones to regulate the traffic and to perform destination IP address translation.

### *Related Topics*

- Configuring Static NAT on page 280

## Configuring Static NAT

Static NAT is a direct one-to-one mapping of one IP address to another without port address translation.

**Before You Begin**

For background information, read:

- Understanding Static NAT on J-series Services Routers on page 279
- Example: Configuring Security Policies—Detailed Configuration on page 84

In this example, you perform the following tasks:

1. Define a static NAT on an interface (ge-0/0.0.0) that maps a host IP address (10.1.1.1) to an address (1.1.1.1).
2. Define a static NAT on an interface (ge-0/0/0.0) that maps a host IP subnet (10.1.3.0/24) to a subnet (1.1.3.0/24).
3. Reference a static NAT (1.1.3.0/24) in a policy (static-nat-policy).

### *CLI Configuration*

To configure static NAT, use the CLI configuration editor:

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 junos-global policy
  incoming match source-address phone2
user@host# set security policies from-zone public to-zone junos-global policy
  incoming match destination-address static_nat_1.1.1.2_32
user@host# set security policies from-zone public to-zone junos-global policy
  incoming match application junos-sip
user@host#set security policies from-zone public to-zone junos-global policy
  incoming then permit
```

### *Related Topics*

- Example: Configuring NAT-Dst Allow-Incoming Table on page 287

## Understanding Static NAT on SRX-series Services Gateways

JUNOS software allows you to create static NAT rules that configure a one-to-one static mapping between original destination address or subnet and translated destination address or subnet. For a 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 addresses, along with other addresses in source and destination NAT pools should never be overlapped within one routing instance.

| Before You Begin |
| --- |
| For background information, read “Source IP Address Translation Overview” on page 291. |

The mapping from original destination address or subnet to translated destination address or subnet in static NAT rule's specify a destination NAT rule. The reverse mapping from translated destination address or subnet to original destination address or subnet in static NAT rule's specify a source NAT rule.

When static NAT is configured on the SRX-series device, the static NAT rules takes precedence over the destination NAT rules during mapping. Similarly it takes precedence over source NAT rules during reverse mapping.

### *Related Topics*

- Example: Configuring Static NAT on SRX-series Services Gateways on page 282

## Example: Configuring Static NAT on SRX-series Services Gateways

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:<br>■ Understanding Static NAT on SRX-series Services Gateways on page 281 |

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 NAT-Dst Policy-Based NAT on J-series Services Routers

JUNOS software allows you to define policies to translate the destination address from one IP address to another.

**Before You Begin**

For background information, read “Destination IP Address Translation Overview” on page 279.

One of the following three methods can be used to configure policy-based NAT-dst. Each configuration method provides a solution for different configuration requirements:

- Translate destination IP address(es) to a single IP address.
- Translate destination IP address(es) and port number(s) to a single IP address and a specific port number.
- Translate a range of destination IP address(es) to another range of IP address(es). The mapping is one-to-one and static.

### *Related Topics*

- Example: Configuring NAT-Dst Allow-Incoming Table on page 287

## Example: Configuring Destination NAT on J-series Services Routers

JUNOS software allows you to define policies to translate the destination address from one IP address to another.

| Before You Begin |
| --- |
| For background information, read "Understanding NAT" on page 276. |

In this example, you perform the following tasks:

- Define a destination NAT (one-addr-dst-nat) with an IP address (20.1.1.1).
- Define a destination NAT (addr-range-dst-nat) with an IP address range (30.1.1.0 to 30.1.1.255).
- Define a destination NAT (addr-and-port-dst-nat) with an IP address (40.1.1.1) and a port number (8080).
- Define a reference (one-addr-dst-nat) in a policy (dst-nat-policy-1) to map a destination IP subnet (2.1.1.0/24) to address (20.1.1.1).
- Define a reference (addr-range-dst-nat) in a policy (dst-nat-policy-2) to map a destination IP subnet (3.1.1.0/24) to a subnet (30.1.1.0/24).
- Define a reference (addr-and-port-dst-nat) in a policy (dst-nat-policy-3) to map a destination IP subnet (4.1.1.0/24) and a port number (80) to an address (40.1.1.1) and a port number (8080).

### *CLI Configuration*

user@host# **set security nat destination-nat one-addr-dst-nat address 20.1.1.1**
user@host# **set security nat destination-nat addr-range-dst-nat address-range low 30.1.1.10 high 30.1.1.255**
user@host# **set security nat destination-nat addr-and-port-dst-nat address 40.1.1.1 port 8080**
user@host# **set security policies from zone Red to-zone Green policy dst-nat-policy-1 match source-address any**
user@host# **set security policies from zone Red to-zone Green policy dst-nat-policy-1 match destination-address 2.1.1.0/24**
user@host# **set security policies from zone Red to-zone Green policy dst-nat-policy-1 match application any**
user@host# **set security policies from zone Red to-zone Green policy dst-nat-policy-1 then permit destination-nat one-address-nat**
user@host# **set security policies from zone Red to-zone Green policy dst-nat-policy-2 match source-address any**
user@host# **set security policies from zone Red to-zone Green policy dst-nat-policy-2 match destination-address 3.1.1.0/24**

```
user@host# set security policies from zone Red to-zone Green policy dst-nat-policy-2
    match application any
user@host# set security policies from zone Red to-zone Green policy dst-nat-policy-2
    then permit destination-nat addr-range-dst-nat
user@host# set security policies from zone Red to-zone Green policy dst-nat-policy-3
    match source-address any
user@host# set security policies from zone Red to-zone Green policy dst-nat-policy-3
    match destination-address 4.1.1.0/24
user@host# set security policies from zone Red to-zone Green policy dst-nat-policy-3
    match application junos-http
user@host# set security policies from zone Red to-zone Green policy dst-nat-policy-3
    then permit destination-nat addr-and-port-dst-nat
```

### *Related Topics*

- Configuring Static NAT on page 280
- Example: Configuring NAT-Dst Allow-Incoming Table on page 287

## Understanding Rule-Based Destination NAT on SRX-series Services Gateways

JUNOS software allows you to create NAT rules to translate the destination address from one IP address to another. For SRX-series devices (unlike J-series), NAT is de-coupled from policies. SRX-series NAT has its own NAT rules to regulate traffic and to perform address translation.

**NOTE:** When performing destination NAT on SRX-series devices, the address in question is first translated according to configured NAT destination rules and then security policies are applied.

**Before You Begin**

For background information, read "Destination IP Address Translation Overview" on page 279.

The following types of destination NAT are supported on SRX-series services gateway:

- Translation of the original destination IP address to an IP address from a defined pool. This type of translation is one-to-one and does not include Port Address Translation (PAT).
- Translation of the original destination IP address (and optional port number) to one specific IP address (and port number) from a defined pool.

The main configuration tasks for destination NAT on SRX-series services gateways are as follows:

- Configure a destination NAT 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.

**NOTE:** When you configure a security policy on SRX-series services gateways, each policy optionally indicates whether it allows NAT translation, does not allow NAT translation, or does not care.

## Example: Configuring Destination NAT on SRX-series Services Gateways

JUNOS software allows you to define NAT rules to translate the destination address from one IP address to another.

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
```

## Understanding NAT-Dst Allow-Incoming Table

JUNOS software supports voice over IP (VoIP) incoming calls—that is, calls that are initiated from the public network to a private network.

**Before You Begin**

For background information, read "Destination IP Address Translation Overview" on page 279.

Source pool (see "Source IP Address Translation Overview" on page 291) is used for source IP translation, port address translation, or both. You can enable the router to allow incoming calls when you define a source pool. When a policy is properly configured, JUNOS software performs the following tasks:

- Monitors outgoing VoIP registration messages
- Performs NAT on addresses within the registration messages using source pool
- Stores the translation information in the allow-incoming table

When an incoming VoIP call is received from a public network, JUNOS software uses the allow-incoming table to identify the internal host to which the incoming call should be routed.

You can take advantage of this feature by using the IP address of the egress interface as a source pool or by defining a source pool on the egress interface. We recommend that you set up a source pool for larger networks or an enterprise environment; using an interface IP address is adequate for handling incoming calls in a small office.

### *Related Topics*

- Example: Configuring NAT-Dst Allow-Incoming Table on page 287
- Understanding NAT Allow-Incoming Source Pools on page 294

## Example: Configuring NAT-Dst Allow-Incoming Table

In this example, a VoIP incoming call involves two policies, one to permit " registered" traffic, and another that permits incoming calls.

**Before You Begin**

For background information, read:

- Understanding NAT-Dst Allow-Incoming Table on page 287
- Example: Configuring Security Policies—Detailed Configuration on page 84

This topic covers:



## J-Web Configuration

To configure NAT:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Nat, click Configure.
4. Next to Interface, click Add new entry.
5. In the Name box, type **ge-0/0/1.0**.
6. Next to Source nat, click Configure.
7. Next to Pool, click Add new entry.
8. In the name box, type **incoming-nat**.
9. Next to Address, click Add new entry.
10. In the Prefix box, type **20.1.1.100** and click OK.
11. Next to Allow incoming, select the check box and click OK.
12. To save and commit the configuration, click Commit.

To configure a Security policy:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **reg-policy**.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match, 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. In 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 box, type **junos-sip** and click OK.
21. Next to Then, click Configure.
22. From the Action list, select Permit.
23. Next to Permit, click Configure.
24. Next to Source nat, select the check box and click Configure.
25. From the Source nat list, select Pool.
26. In the Pool box, type **incoming-nat** and click OK.
27. To save and commit the configuration, click Commit.

To configure another policy from zone red to zone junos-global, follow the sequence of steps listed below:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **Red**.
6. In the To zone name box, type **junos-global**.
7. Next to Policy, click Add new entry.
8. In the Policy name box, type **incoming-policy**.
9. Select the Match check box.
10. Select the Then check box.
11. Next to Match, 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. In the Value keyword list, select Enter Specific Value.

18. In the Value keyword list, type **incoming_nat_incoming-nat** 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. From the Action list, select Permit and click OK.
24. To save and commit the configuration, click Commit.

### *CLI Configuration*

```
user@host# set security nat interface ge-0/0/1.0 source-nat pool incoming-nat
  address 20.1.1.100
user@host# set security nat interface ge-0/0/1.0 source-nat pool incoming-nat
  allow-incoming
user@host# set security policies from-zone Green to-zone Red policy reg-policy match
  source-address any
user@host# set security policies from-zone Green to-zone Red policy reg-policy match
  destination-address any
user@host# set security policies from-zone Green to-zone Red policy reg-policy match
  application junos-sip
user@host# set security policies from-zone Green to-zone Red policy reg-policy then
  permit source-nat pool incoming-nat
user@host# set security policies from-zone Red to-zone junos-global policy
  incoming-policy match source-address any
user@host# set security policies from-zone Red to-zone junos-global policy
  incoming-policy match destination-address incoming_nat_incoming-nat
user@host# set security policies from-zone Red to-zone junos-global policy
  incoming-policy match application junos-sip
user@host# set security policies from-zone Red to-zone junos-global policy
  incoming-policy then permit
```

### *Related Topics*

- Configuring Static NAT on page 280
- Example: Configuring NAT-Dst Allow-Incoming Table on page 287

### *Verifying NAT Incoming-table*

**Purpose** Display NAT table information.

**Action** Use the show security nat incoming-table CLI command to verify the output for the NAT table. You get the following output:

```
user@host> show security nat incoming-table
  In use: 1, Maximum: 8, Entry allocation failed: 0
  Destination           Host                 References Timeout Source-pool
  10.1.1.26:1028        1.1.1.10:5060                 1    3600 p1
```

**What it Means** The output displays information about the host IP address and port number that the destination IP address is mapped to, the name of the source pool where translation is allocated, and the timeout, in seconds, of the entry in the NAT table.

## Source IP Address Translation Overview

During source IP address translation, the original source IP address, port number, or both are translated to a different number.

**Before You Begin**

For background information, read:

- Understanding NAT on page 276
- Example: Configuring Security Policies—Detailed Configuration on page 84

JUNOS software performs NAT in the following ways:

- Reverse mapping in static NAT—There is no explicit NAT action in a policy to trigger this type of NAT-src. As long as the traffic matches a " permit" policy and reverse mapping in static NAT is found on the egress interface, the source IP address will get translated according to reverse mapping.
- Source pool defined on egress interface—When performing source network translation (NAT-src), source pool provides JUNOS software with a supply of addresses from which to draw. When a policy requires NAT-src and references a specific source pool, JUNOS software draws addresses from that pool when translation is performed.

### *Related Topics*

- Understanding NAT Interface Source Pools on page 292
- Understanding NAT Source Pools with PAT on page 292
- Understanding NAT Source Pools Without PAT on page 293
- Understanding NAT Static Source Pools on page 294
- Understanding NAT Allow-Incoming Source Pools on page 294
- Understanding NAT Source Pool Sets on page 295

## Understanding NAT Interface Source Pools

When traffic requires source IP address translation and references a specific source pool, JUNOS software draws addresses from that pool when translation is performed. The interface source pool is predefined.

**Before You Begin**

For background information, read “Source IP Address Translation Overview” on page 291.

You can reference an interface source pool in a policy or rule by setting a NAT action. JUNOS software translates the source IP address to the address of the egress interface for the traffic matching a policy or rule that references the interface source pool. JUNOS software always applies port address translation (PAT) for the interface source pool.

### *Related Topics*

- Understanding NAT Source Pools with PAT on page 292
- Example: Configuring Source NAT on J-series Services Routers on page 295
- Example: Configuring Source NAT on SRX-series Services Gateways on page 297

## 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. Using PAT, multiple hosts can share the same IP address.

**Before You Begin**

For background information, read “Source IP Address Translation Overview” on page 291.

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 293
- Address Persistent on page 293
- Related Topics on page 293

### *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 J-series Services Routers on page 295
- Example: Configuring Source NAT on SRX-series Services Gateways on page 297
- Example: Configuring a Persistent Address and Pool Sets on page 299

## 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 “Source IP Address Translation Overview” on page 291.

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*

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 J-series Services Routers on page 295
- Example: Configuring Source NAT on SRX-series Services Gateways on page 297

## Understanding NAT Static Source Pools

You can define a one-to-one mapping from an original source IP address to a translated source IP address for a range of IP addresses using the static source pool. Such a mapping ensures that the J-series Services Router always translates a particular source IP address from within that range to the same translated address within a source pool.

| **Before You Begin** |
| --- |
| For background information, read "Source IP Address Translation Overview" on page 291. |

It is possible to use a static source pool enabled in a policy that applies to source addresses beyond the range specified in the pool. In such scenarios, the router allows traffic from all source addresses permitted in the policy, applying source IP address translation (NAT-src) to those addresses that fall within the static source pool range but leaving those addresses unchanged that fall outside the static source pool range. If you want the router to apply NAT-src to all source addresses, make sure that the range of source addresses is smaller than the range specified for the static source pool.

### *Related Topics*

- Example: Configuring Source NAT on J-series Services Routers on page 295
- Example: Configuring Source NAT on SRX-series Services Gateways on page 297

## Understanding NAT Allow-Incoming Source Pools

Source pool is designed to support source IP address translation, port address translation, or both. You can define a source pool with PAT to support VoIP incoming calls.

| **Before You Begin** |
| --- |
| For background information, read "Source IP Address Translation Overview" on page 291. |

The allow-incoming calls feature is supported on an interface source pool. When the allow-incoming feature is enabled on a source pool (p1), the following format *incoming_nat_p1* (or as *incoming_nat_ge-0/0/0.0* for an interface source pool on an interface *ge-0/0/0.0*) is referenced in a policy. Similar to static NAT, such a reference can be used as a destination address in the policy configuration to regulate the VoIP incoming traffic and to perform destination IP and port number translations.

### *Related Topics*

- Example: Configuring Source NAT on J-series Services Routers on page 295
- Example: Configuring Source NAT on SRX-series Services Gateways on page 297
- Example: Configuring a Persistent Address and Pool Sets on page 299
- Understanding NAT-Dst Allow-Incoming Table on page 287

## Understanding NAT Source Pool Sets

JUNOS software NAT is interface based, whereas policy and rule are zone based. Each zone can contain multiple interfaces. You can define a source pool on each egress interface in the zone, group these source pools into a pool set, and then reference the pool set in one policy or rule.

| **Before You Begin** |
| --- |
| For background information, read “Source IP Address Translation Overview” on page 291. |

The device can perform source IP or port address translation using proper source pools within the pool set for the traffic to match the policy or rule and to exit on certain egress interfaces in the zone.

### *Related Topics*

- Example: Configuring a Persistent Address and Pool Sets on page 299

## Example: Configuring Source NAT on J-series Services Routers

When performing source Network Address Translation (NAT-src), source pool provides JUNOS software with a supply of addresses from which to draw. When a policy

requires NAT-src and references a specific source pool, JUNOS software draws addresses from that pool when translation is performed.

**Before You Begin**

For background information, read:

- Source IP Address Translation Overview on page 291
- Understanding NAT Interface Source Pools on page 292
- Understanding NAT Source Pools with PAT on page 292
- Understanding NAT Source Pools Without PAT on page 293
- Understanding NAT Static Source Pools on page 294
- Understanding NAT Allow-Incoming Source Pools on page 294
- Example: Configuring Security Policies—Detailed Configuration on page 84

In this example, you perform the following tasks:

- Define a source pool (src-nat-with-pat) on an interface (ge-0/0/0.0) that contains an address range (10.1.1.20 10.1.1.30) and an address (10.1.1.2) with port address translation.
- Define a source pool (src-nat-wo-pat) on an interface (ge-0/0/0.0) that contains an address range (10.1.1.3 10.1.1.5) without port address translation and uses (src-nat-with-pat) as the overflow pool.
- Define a source pool (src-nat-static) on an interface (ge-0/0/1.0) that maps an IP address range (1.1.1.10 1.1.120) to an IP address range (20.1.1.10 20.1.1.20).
- Define a source pool (src-nat-incoming) on an interface (ge-0/0/1.0) that contains an IP address range (20.1.1.25 20.1.1.30) with allow-incoming enabled.
- Reference source pool (src-nat-with-pat) in a policy (src-nat-policy-1).
- Reference an interface source pool in a policy (src-nat-policy-2).

## *CLI Configuration*

user@host# **set security nat interface ge-0/0/0.0 source-nat pool src-nat-with-pat address 10.1.1.2**
user@host# **set security nat interface ge-0/0/0.0 source-nat pool src-nat-with-pat address-range low 10.1.1.20 high 10.1.1.30**
user@host# **set security nat interface ge-0/0/0.0 source-nat pool src-nat-wo-pat address-range low 10.1.1.3 high 10.1.1.5**
user@host# **set security nat interface ge-0/0/0.0 source-nat pool src-nat-wo-pat no-port-translation**
user@host# **set security nat interface ge-0/0/1.0 source-nat pool src-nat-static address-range low 20.1.1.10 high 20.1.1.20**
user@host# **set security nat interface ge-0/0/1.0 source-nat pool src-nat-static host-address-low 1.1.1.10**
user@host# **set security nat interface ge-0/0/1.0 source-nat pool src-nat-incoming address low 20.1.1.25 high 20.1.1.50**
user@host#**set security nat interface ge-0/0/1.0 source-nat pool src-nat-incoming allow-incoming**

```
user@host# set security policies from-zone Red to-zone Green policy src-nat-policy-1
   match source-address any
user@host#set security policies from-zone Red to-zone Green policy src-nat-policy-1
   match destination-address any
user@host#set security policies from-zone Red to-zone Green policy src-nat-policy-1
   match application any
user@host#set security policies from-zone Red to-zone Green policy src-nat-policy-1
   then permit source-nat pool src-nat-with-pat
user@host# set security policies from-zone Red to-zone Green policy src-nat-policy-2
   match source-address any
user@host# set security policies from-zone Red to-zone Green policy src-nat-policy-2
   match destination-address any
user@host# set security policies from-zone Red to-zone Green policy src-nat-policy-2
   match application any
user@host# set security policies from-zone Red to-zone Green policy src-nat-policy-2
   then permit source-nat interface
```

### *Related Topics*

- Example: Configuring a Persistent Address and Pool Sets on page 299

## Example: Configuring Source NAT on SRX-series Services Gateways

When performing source Network Address Translation, 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. For SRX-series devices (unlike J-series), NAT is de-coupled from policies. SRX-series NAT has its own NAT rules to regulate traffic and to perform address translation.

**NOTE:** When performing source NAT on SRX-series devices, 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 IP Address Translation Overview on page 291
- Understanding NAT Interface Source Pools on page 292
- Understanding NAT Source Pools with PAT on page 292
- Understanding NAT Source Pools Without PAT on page 293
- Understanding NAT Static Source Pools on page 294
- Understanding NAT Allow-Incoming Source Pools on page 294
- Example: Configuring Security Policies—Detailed Configuration on page 84

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:** 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 to do source NAT 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 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*

```
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

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
```

## Verifying Static NAT Summary

**Purpose** Display static NAT summary information.

**Action** Use the show security nat static-nat summary CLI command to verify static NAT l information. You get the following output:

```
user@host> show security nat static-nat summary
  Total static NAT mappings: 3, Maximum: 300
  Ingress Interface    Destination    Host         Virtual Router
  ge-0/0/0.0           1.1.1.1/32     10.1.1.1/32  trust-vr
  ge-0/0/0.0           1.1.3.0/24     10.1.3.0/24  trust-vr
  ge-0/0/0.1           2.2.2.1/32     20.1.1.1/32  trust-vr
```

**What it Means** The output displays information about the total static NAT mappings, destination IP address and subnet mask, name of the interface on which static NAT is defined, and the name of the virtual router that performs route lookup for the host IP address and the subnet mask.

## Example: Configuring a Persistent Address and Pool Sets

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 293
- Understanding NAT Source Pool Sets on page 295

In this example, you perform the following tasks:

- Enable a persistent address globally.
- Configure pool utilization thresholds.
- Define a pool-set (pool-set-example) that contains two members (src-nat-with-pat) and src-nat_incoming).
- Reference a pool-set (pool-set-example) in a policy (pool-set-policy).

To configure a persistent address and pool sets, use the CLI configuration editor.

This topic covers:

- CLI Configuration on page 300
- Related Topics on page 300

## CLI Configuration

```
user@host# set security nat source-nat address-persistent
user@host# set security nat source-nat pool-utilization-alarm raise-threshold 90
user@host# set security nat source-nat pool-utilization-alarm clear-threshold 80
user@host# set security nat source-nat pool-set pool-set-example pool src-nat-with-pat
user@host# set security nat source-nat pool-set pool-set-example pool
   src-nat-incoming
user@host# set security policies from-zone Red to-zone Green policy pol-set-policy
   match source-address any
user@host# set security policies from-zone Red to-zone Green policy pol-set-policy
   match destination-address any
user@host# set security policies from-zone Red to-zone Green policy pol-set-policy
   match application any
user@host# set security policies from-zone Red to-zone Green policy pol-set-policy
   then permit source-nat pool-set pool-set-example
```

## Related Topics

- Example: Configuring Source NAT on J-series Services Routers on page 295
- Example: Configuring Source NAT on SRX-series Services Gateways on page 297

## Configuring Proxy ARP on SRX-series Services Gateways

On SRX-series devices, 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 interface on which to configure proxy ARP. It should be the logical interface. 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 on SRX–series Services Gateways

The NAT trace options hierarchy configures trace file and flags for verification purposes. 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.

**NOTE:** The trace data is written to/var/log/security-trace by default.

### *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**

## Configuring Source NAT—Quick Configuration

You can use J-Web Quick Configuration to quickly configure source NAT.

**Before You Begin**

For background information, read

- "Configuring Host Inbound Traffic" on page 55
- "Understanding NAT Interface Source Pools" on page 292
- "Example: Configuring Source NAT on J-series Services Routers" on page 295

Figure 60 on page 302 shows the Quick Configuration for Source NAT page.

**Figure 60: Quick Configuration Page for Source NAT**

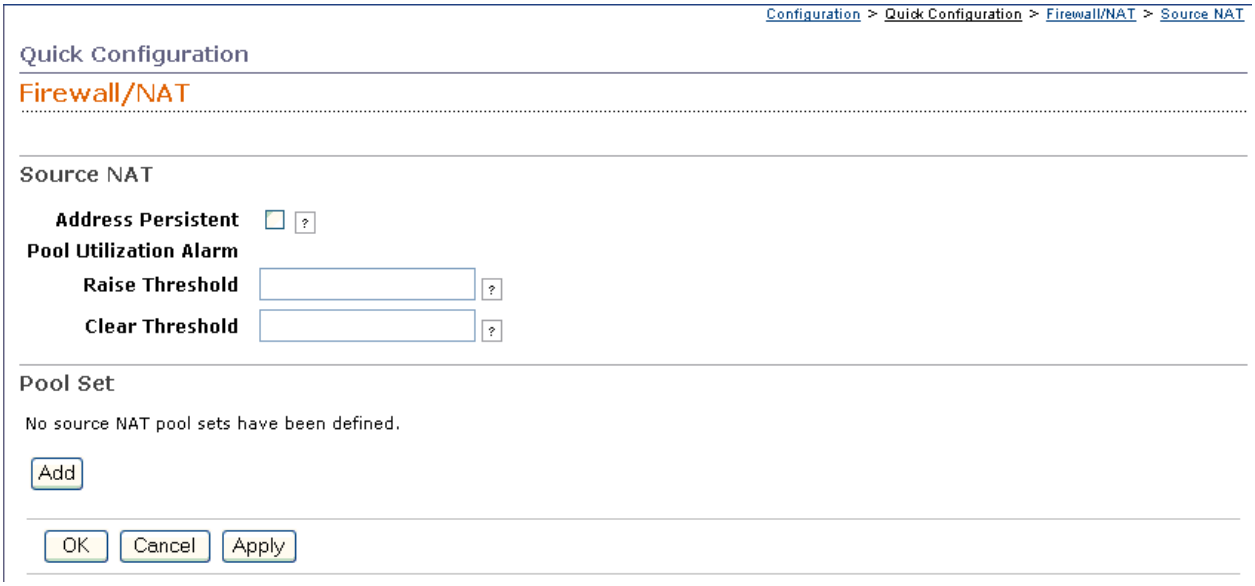


To configure source NAT with Quick Configuration:

1. Select Configuration > Quick Configuration > Firewall/NAT > Source NAT. See Figure 60 on page 302.
2. Fill in the options as shown in Table 53 on page 303.
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 53: Source NAT Configuration Options**

| Field | Function | Action |
|---|---|---|
| **Source NAT** | | |
| 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.<br><br>Clear Threshold—Enter a number between 40-100 to clear the threshold for the pool utilization alarm. |
| **Pool Set** | | |
| Add Pool Sets | Defines one or more source pool sets (group of source pools). | Click **Add.** |
| Pool Set Name | Descriptive name for the pool set. | Type a name for the pool set. |
| Pool | Name of the pools to be included in the address pool set. | Select the pools to be added and click **Add**. |

## Configuring Destination NAT—Quick Configuration

You can use J-Web Quick Configuration to quickly configure destination NAT.

**Before You Begin**

For background information, read

- Understanding NAT on page 276
- Destination IP Address Translation Overview on page 279
- Example: Configuring NAT-Dst Allow-Incoming Table on page 287
- Example: Configuring Destination NAT on SRX-series Services Gateways on page 286

Figure 61 on page 304 shows the Quick Configuration for Destination NAT page.

Figure 62 on page 304 shows the Destination NAT page.

**Figure 61: Quick Configuration Page for Destination NAT**

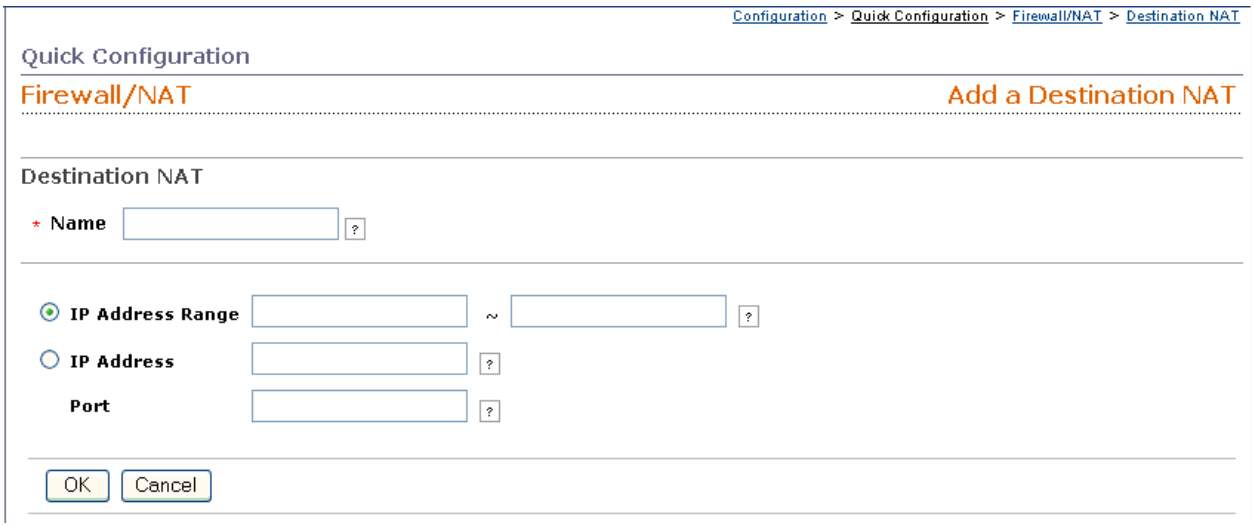


**Figure 62: Destination NAT**

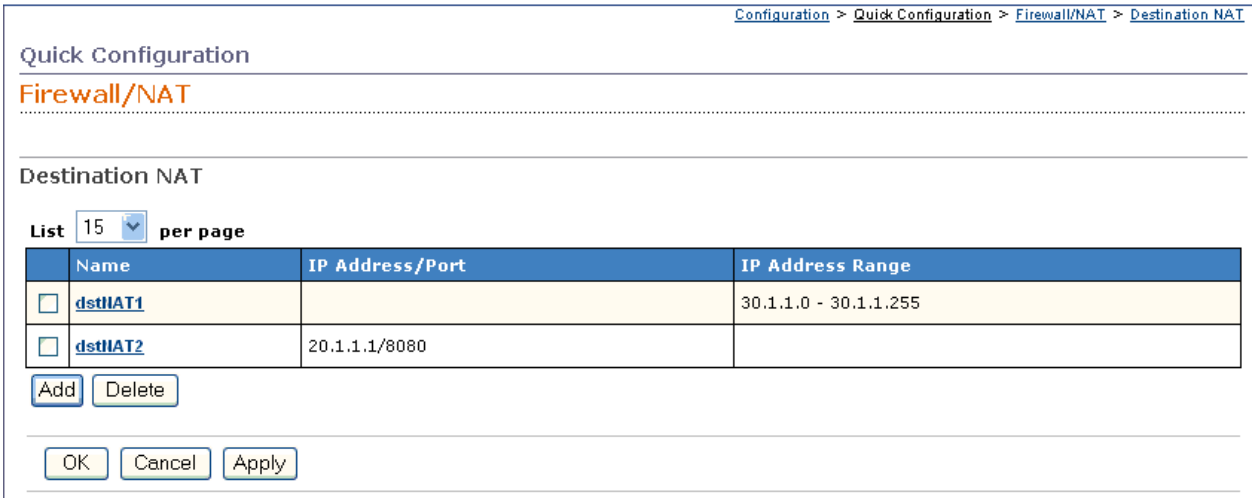


To configure destination NAT with Quick Configuration:

1. Select Configuration > Quick Configuration > Firewall/NAT > Destination NAT. See Figure 61 on page 304.
2. Click Add. The Quick Configuration screen for destination NAT appears as shown in Figure 62 on page 304.
3. Fill in either an IP address range or a single IP address with its port number. See Table 54 on page 305 for more information on these options.
4. Click one of the following buttons:
   - To apply the configuration, click OK. The Destination NAT screen appears as shown in Figure 62 on page 304.

- To cancel your entries and return to the main page, click Cancel.

5. Click one of the following buttons:
   - To add another destination NAT configuration, click Add. You will return to the configuration screen as shown in Figure 61 on page 304.
   - To delete an already configured destination NAT, select the desired dst_NAT from the list and click Delete.
   - 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 54: Destination NAT Configuration Options**

| Field | Function | Action |
|---|---|---|
| **Destination NAT** | | |
| Name | Name for the destination NAT. | Specify a name for the destination NAT. |
| IP Address Range | A range of IP addresses for destination address translation. Multiple address ranges can be specified in a pool. | Specify an address range, a low IP address to specify the start address in the range and a high IP address to specify the end address in the range.<br><br>For example, 30.1.1.0 to 30.1.1.255. |
| IP Address | IP address to which destination IP addresses are translated. | Specify an IP address. This can be a fixed IP address for a subnet. |
| Port | (Optional) Port number to which destination IP addresses are translated. | Specify a port number. |

## Configuring Interface NAT—Quick Configuration

You can use J-Web Quick Configuration to quickly configure interface NAT.

**Before You Begin**

For background information, read

- "Understanding NAT" on page 276
- "Understanding NAT Interface Source Pools" on page 292
- "Understanding Security Zone Interfaces" on page 63

Figure 63 on page 306 shows the Quick Configuration for Interface NAT page.

**Figure 63: Quick Configuration Page for Interface NAT**

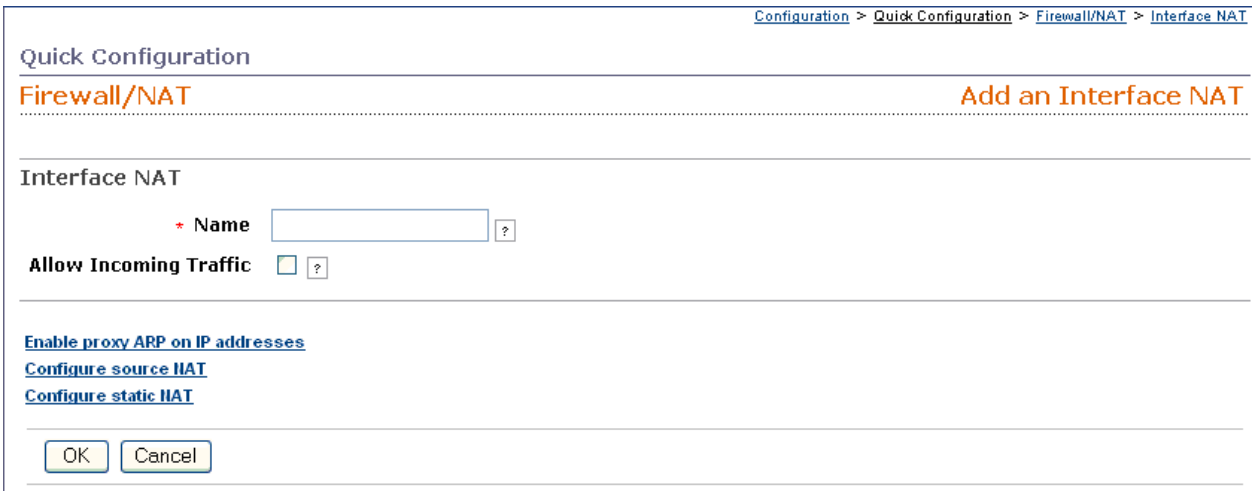


To configure interface NAT with Quick Configuration:

1. Select Configuration > Quick Configuration > Firewall/NAT > Interface NAT.
2. Click Add. The Quick Configuration screen for interface NAT appears as shown in Figure 63 on page 306.
3. Fill in either an IP address range or a single IP address with its port number. See Table 55 on page 307 for more information on these options.
4. 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.

**Figure 64: Source NAT Pool Configuration**

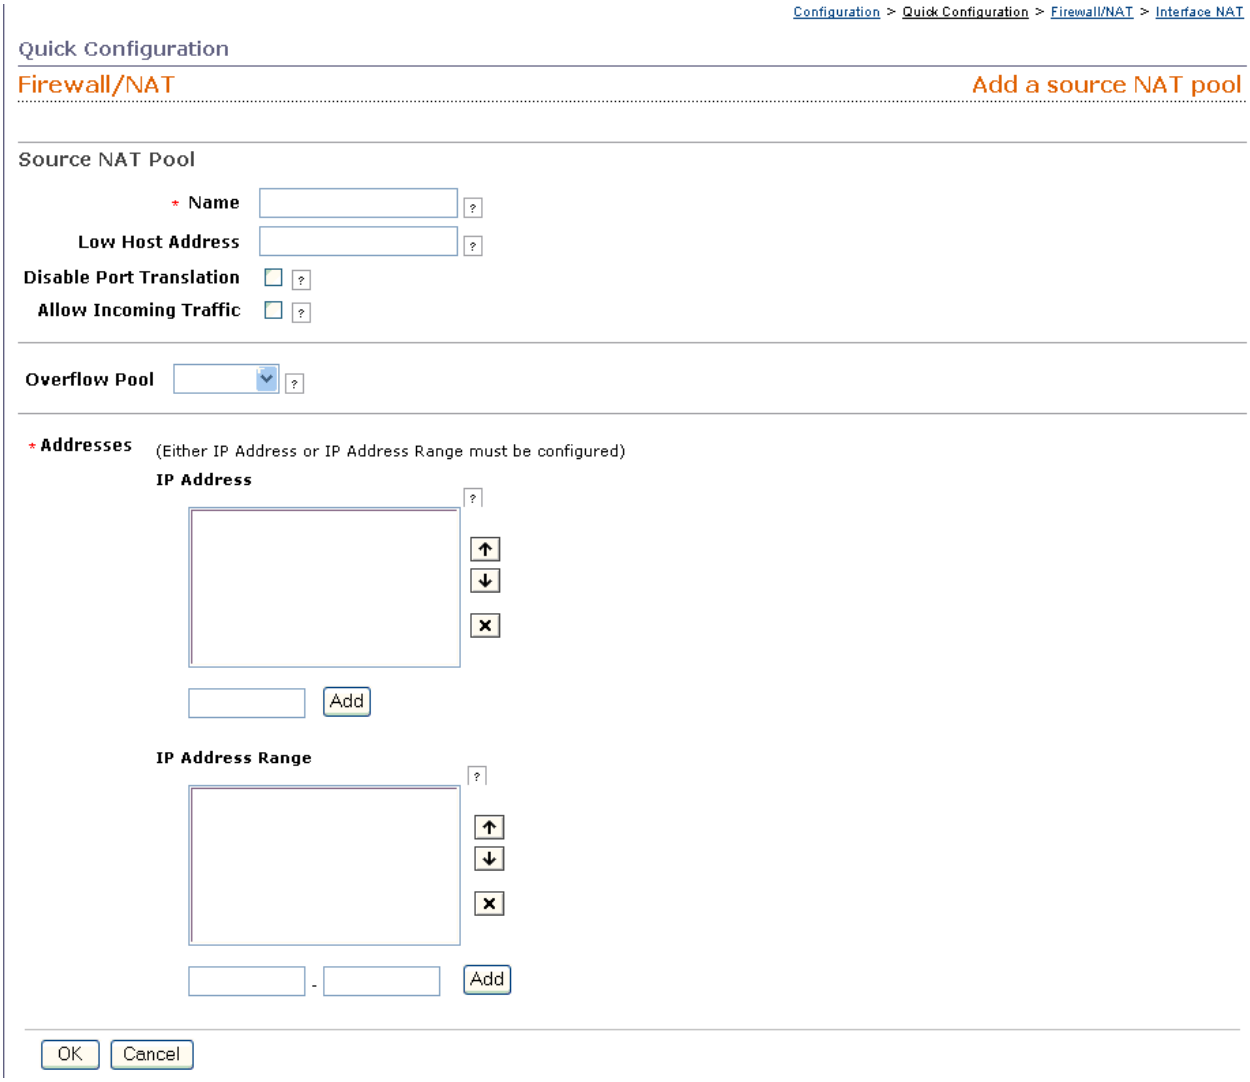


**Table 55: Interface NAT Configuration Options**

| Field | Function | Action |
|---|---|---|
| **Interface NAT** | | |
| Name | Indicates name of the interface. | Provide an already defined logical interface name. For example, ge-0/0/0.0. |
| Allow Incoming Traffic | Supports VoIP incoming calls—that is, the calls that are initiated from the public network to the private network. | Select this option to enable incoming traffic. |
| **Enable Proxy ARP** | | |

**Table 55: Interface NAT Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| IP Address | Configures proxy ARP for a single IP address. | Specify the IP address for which you want to configure proxy ARP. |
| IP Address Range | Configures proxy ARP for a range of IP addresses. | Specify the IP address range for which you want to configure proxy ARP. For example, a low IP address (10.1.1.150) to specify the start address in the range and a high IP address (10.1.1.200) to specify the end address in the range.<br><br>Click OK. |
| Configure Source NAT | Indicates name of the source address pool. | Specify a name for the source address pool. |
| Interface Source NAT Pools | Indicates predefined source pool. | Click Add. The Source NAT Pool configuration screen displays as shown in Figure 64 on page 307. |
| **Name** | | |
| Low Host Address | Indicates lower limit of the original source address range that is used to define the static source pool. | Specify the start address of the IP address range. |
| Disable Port Translation | Makes the pool perform address translation only. | Select the checkbox to disable port translation. |
| Allow Incoming Traffic | Allows VoIP incoming calls. | Select the checkbox to allow this pool to support incoming traffic. |
| Overflow Pool | Indicates a source pool to use when the current address pool is exhausted. | From the dropdown list, select the overflow pool. |
| Addresses | Indicates IP address(es) for the source NAT pool. | Configure either a single IP address or an IP address range for the source NAT pool. For more information, see "Configuring Source NAT—Quick Configuration" on page 302.<br><br>Click OK. |
| **Configure Static NAT** | | |
| Static NAT | Indicates a direct one-to-one mapping of one IP address to another<br><br>without port address translation. | Click Add to define static NAT. |
| Mapped Address | Indicates IP address to map. | Specify the IP address to be mapped. |
| Host Address | Indicates IP address of host | Specify the IP address of the host to which another address is mapped. |
| Virtual Router | Indicates virtual router name. | Specify the name of the virtual router that is to perform route lookup for the host address(es).<br><br>Click OK. |

## Configuring Firewall/NAT Flow—Quick Configuration

You can use J-Web Quick Configuration to quickly configure stateful firewall or NAT flow

**Before You Begin**

For background information, read

- "Configuring Host Inbound Traffic" on page 55
- "Destination IP Address Translation Overview" on page 279
- "Understanding Static NAT on J-series Services Routers" on page 279
- "Understanding NAT Interface Source Pools" on page 292

Figure 65 on page 310 shows the Quick Configuration for Stateful Firewall or NAT Flow page.

**Figure 65: Quick Configuration Page for Stateful Firewall or NAT Flow**

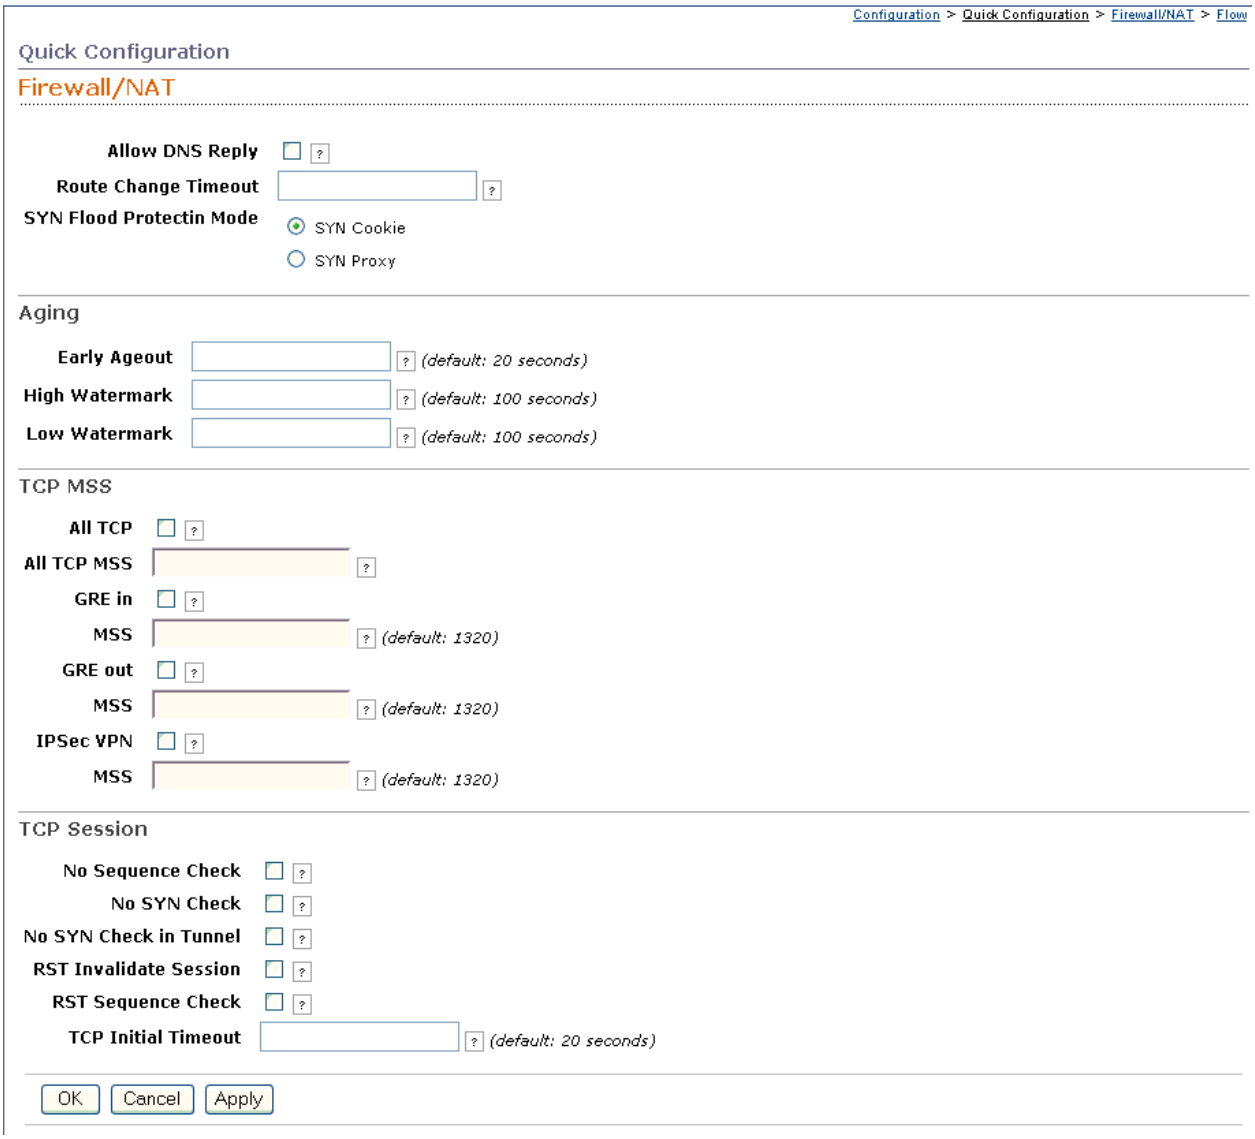


To configure Firewall/NAT Flow with Quick Configuration:

1. Select Configuration > Quick Configuration > Firewall/NAT > Flow. See Figure 65 on page 310.
2. Fill in the options as shown in Table 56 on page 311.
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 56: Firewall or NAT Flow Configuration Options**

| Field | Function | Action |
|---|---|---|
| **Firewall NAT** | | |
| Allow DNS Reply | Allows an incoming DNS reply packet without a matched request. By default, if the query request does not match, the router drops the packet, does not create a session, and increments the illegal packet flow counter for the interface. Using the allow-dns-reply statement directs the router to skip the check. | Select this checkbox to enable DNS replies. |
| Route Change Timeout | Applies the session timeout value on a route change to a nonexistent route. By default, this feature is disabled. If the timeout is not defined, sessions discovered to have no route are aged out using their current session timeout values. | Specify a value between 6 and 1800 seconds. |
| SYN Flood Protection Mode | Enables SYN-cookie defenses or SYN-proxy defenses against SYN attacks. Sets the flow from traditional SYN Proxy mode to SYN Cookie mode. SYN Cookie is enabled globally on the security router and is activated when the configured syn-flood attack-threshold is exceeded. | Select SYN Cookie or SYN Proxy. |
| **Aging** | | |
| Early Ageout | Defines the ageout value before the router aggressively ages out a session from its session table. | Specify a value between 1 and 65535 seconds. The default value is 20 seconds. |
| High Watermark | Sets percentage of session table capacity at which the aggressive aging-out process begins. | Specify a value between 0 and 100 percent. The default value is 100 percent. |
| Low Watermark | Sets percentage of session-table capacity at which aggressive aging-out ends. | Specify a value between 0 and 100 percent. The default value is 100 percent. |
| **TCP MSS** | | |
| All TCP | Sets all TCP packets for network traffic. | Select the All TCP check box to enable all TCP packets. |
| All TCP MSS | Sets the TCP-maximum segment size (TCP-MSS) value for all TCP packets for network traffic. | Specify a value between 64 and 65535. |
| GRE in | Enables MSS override for all Generic Routing Encapsulation (GRE) packets exiting an IPsec tunnel. | Select the GRE in check box to enable TCP-MSS for GRE. |
| MSS | Enables and specifies the TCP-MSS for GRE packets that are about to go into an IPsec VPN tunnel. By default, a TCP-MSS for GRE packets is not set. | Specify a value between 64 and 65,535 bytes. The default value is 1320 bytes. |
| GRE out | Enables MSS override for all GRE packets entering an IPsec tunnel. | Select the GRE out check box to enable. |

**Table 56: Firewall or NAT Flow Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| MSS | Enables and specifies the TCP-MSS for GRE packets that are leaving an IPsec VPN tunnel. By default, a TCP-MSS for GRE packets is not set. | Specify a value between 64 and 65,535 bytes. The default value is 1320 bytes. |
| IPsec VPN | Enables MSS override for all packets entering an IPsec tunnel. | Select the IPSec VPN check box to enable MSS override for all packets that enter an IPsec tunnel. |
| MSS | Enables and specifies the TCP-MSS for all packets that are entering an IPsec VPN tunnel. | Specify a value between 64 and 65,535 bytes. The default value is 1320 bytes. |
| **TCP Session** | | |
| No Sequence Check | Disables the checking of sequence numbers in TCP segments during stateful inspection. By default, the router monitors the sequence numbers in TCP segments. | Select the checkbox to disable sequence number checking. |
| No SYN Check | Disables the checking of the TCP SYN bit before creating a session. By default, the router checks that the SYN bit is set in the first packet of a session. If it is not set, the router drops it. | Select the checkbox to disable creation time SYN-flag check. |
| No SYN Check in Tunnel | Disables the checking TCP SYN bit before creating a session for tunneled packets. By default, the router checks that the SYN bit is set in the first packet of a VPN session. If it is not set, the router drops it. | Select the checkbox to disable creation time SYN-flag check for tunnel packets. |
| RST Invalidate Session | Marks a session for immediate termination when it receives a TCP reset (RST) segment. By default, this statement is unset. When unset, the router applies the normal session timeout interval—for TCP, session timeout is 30 minutes; for HTTP, it is 5 minutes; and for UDP, it is 1 minute. | Select this checkbox to immediately end session on receipt of reset (RST) segment. |
| RST Sequence Check | Checks that the TCP sequence number in a TCP segment with the RST bit enabled matches the previous sequence number for a packet in that session or is the next higher number incrementally. By default, this check is disabled. | Select this checkbox to enable checking of sequence numbers in a RST statement. |
| TCP Initial Timeout | Defines the length of time (in seconds) that the router keeps an initial TCP session in the session table before dropping it, or until the router receives a FIN or RST packet. | Specify a value between 20 and 300 seconds. The default value is 20 seconds. |

## Configuring Stateful Firewall or NAT Screen—Quick Configuration

You can use J-Web Quick Configuration to quickly configure stateful firewall or NAT screen options.

**Before You Begin**

For background information, read

- Understanding NAT on page 276
- Understanding IP Address Sweeps on page 182
- Understanding Network Reconnaissance Using IP Options on page 187
- Understanding Operating System Probes on page 191
- Understanding Attacker Evasion Techniques on page 197

For more information on configuring firewall screen options using Quick Configuration, see "Configuring Firewall Screen Options—Quick Configuration" on page 265.

Chapter 13

# Chassis Cluster

This section includes:

- Disabling Chassis Cluster on page 365
- Chassis Cluster Configuration Scenarios on page 366

## Understanding Chassis Cluster

Chassis clustering provides network node redundancy by grouping a pair of the same kind of supported J-series devices or SRX-series devices into a cluster. The devices must be running JUNOS software. The two nodes back up each other, with one node acting as the primary and the other as the secondary, ensuring stateful failover of processes and services in the event of system or hardware failure. If the primary node fails, the secondary takes over processing of traffic. Nodes in a cluster are interconnected over Gigabit Ethernet links and synchronize configuration, kernel, and session state across the cluster to facilitate high availability of interfaces and services.

**NOTE:** For an SRX-series chassis cluster, nodes can also be interconnected over 10-Gigabit Ethernet links.

**NOTE:** For a chassis cluster with SRX 5600 or SRX 5800 devices, nodes must be interconnected with a fiber optic cable.

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.

### *Related Topics*

- Understanding Chassis Cluster Formation on page 316

## Understanding Chassis Cluster Formation

To form a chassis cluster, a pair of the same kind of supported J-series devices or SRX-series devices combines to act as a single system that enforces the same overall security. For J-series chassis clusters, although the devices must be the same kind, they can contain different Physical Interface Modules (PIMs). For SRX-series chassis clusters, components including interfaces and processing units must match on the two boxes. 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 316 |

### *Related Topics*

- Understanding Redundancy Groups on page 317
- Understanding Redundant Ethernet Interfaces on page 322
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326
- Setting the Node ID and Cluster ID on page 339
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342
- What Happens When You Enable Chassis Cluster on page 331
- Understanding Failover on page 328

## 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 316 |

This topic includes:

## 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.

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 255 redundancy groups.

**NOTE:** In this release, an SRX-series chassis cluster supports the creation of only one redundancy group beyond redundancy group 0.

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 it 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 chapter 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. Redundancy group 0 is primary, and the Routing Engine is active on the node with the higher priority. By default, both nodes have the same priority for redundancy group 0, but you can change the default setting to specify which node is primary for redundancy group 0. Here is how redundancy group 0 primacy is determined:

- If both nodes of a cluster are initialized at the same time and you have not changed the default setting for redundancy group 0 node priority, node 0 takes precedence.

- If you have not changed the default setting for redundancy group 0 node priority and one node of the cluster is initialized before the other, the first node to be initialized takes precedence. In this case, the Routing Engine on the first initialized node is the active one and the node is considered primary. (The primary node is not necessarily node 0. If you boot node 1 before node 0, node 1's Routing Engine takes precedence.)

  The other node is considered secondary. The secondary node's Routing Engine is synchronized with state information from the primary node so that it is ready to take over if the primary node fails.

- If you set the redundancy group 0 node priority, the Routing Engine on the node with the higher priority takes precedence.

**NOTE:** In addition to redundancy group 0, on J-series chassis clusters you configure other redundancy groups that manage the interfaces of the cluster nodes (only one such redundancy group supported on SRX-series chassis clusters). If all these redundancy groups fail over from one node to the other, the node whose redundancy group 0 group is currently primary remains the primary node. The node whose Routing Engine is active is always 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.

**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 primary 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 255*

**NOTE:** In this release, an SRX-series chassis cluster supports the creation of only one redundancy group beyond redundancy group 0.

You can configure one or more redundancy groups numbered 1 through 255, referred to in this chapter as redundancy group *x*. 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 pseudo interface that contains a pair of physical Gigabit Ethernet interfaces or a pair of Fast Ethernet interfaces. Redundancy groups can contain one or more redundant Ethernet interfaces. A redundant Ethernet interface has two child links, one from each node. 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 J-series chassis clusters, you can configure multiple redundancy groups to load-share traffic across the cluster (only one additional redundancy group supported on SRX-series chassis clusters). 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.

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 57 on page 320 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 57: Redundancy Groups Example for a J-series Chassis Cluster**

| Group | Primary | Priority | Objects | Interface | Interface |
|---|---|---|---|---|---|
| Redundancy group 0 | Node 1 | Node 0: 50<br>Node 1: 100 | Routing Engine on node 0<br>Routing Engine on node 1 | —<br>— | —<br>— |
| Redundancy group 1 | Node 1 | Node 0: 50<br>Node 1: 100 | Redundant Ethernet interface 0<br>Redundant Ethernet interface 1 | fe-1/0/0<br>fe-1/0/1 | fe-8/0/0<br>fe-8/0/1 |
| Redundancy group 2 | Node 1 | Node 0: 50<br>Node 1: 100 | Redundant Ethernet interface 2<br>Redundant Ethernet interface 3 | ge-2/0/0<br>ge-2/0/1 | ge-9/0/0<br>ge-9/0/1 |
| Redundancy group 3 | Node 0 | Node 0: 100<br>Node 1: 50 | Redundant Ethernet interface 4<br>Redundant Ethernet interface 5 | ge-3/0/0<br>ge-3/0/1 | ge-10/0/0<br>ge-10/0/1 |

As the example for a J-series chassis cluster in Table 57 on page 320 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.

Table 58 on page 321 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*.

**Table 58: Redundancy Groups Example for an SRX-series Chassis Cluster**

| 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 | — | — |
| Redundancy group 1 | Node 0 | Node 0: 254 | Redundant Ethernet interface 0 | ge-1/0/0 | ge-23/0/0 |
| | | Node 1: 2 | | ge-1/3/0 | ge-23/3/0 |
| | | | Redundant Ethernet interface 1 | | |

As the example for an SRX-series chassis cluster in Table 58 on page 321 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 Interface Monitoring*

For a redundancy group *x* to automatically fail over to another node, its interfaces must be monitored. When you configure a redundancy group *x*, you can specify a set of interfaces the redundancy group *x* 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 *x* to monitor, you give it a weight.

Every redundancy group *x* has a threshold tolerance value initially set to 255. When an interface 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 interfaces has brought its threshold value to 0. When the monitored interfaces 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 322
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342
- Configuring Redundant Ethernet Interfaces on page 355
- Initiating a Manual Redundancy Group Failover on page 357

## Understanding Redundant Ethernet Interfaces

A redundant Ethernet interface is a pseudo interface that includes a physical interface from each node of the cluster.

| Before You Begin |
| --- |
| Read "Understanding Chassis Cluster Formation" on page 316 and "Understanding Redundancy Groups" on page 317 |

A redundant Ethernet interface can contain either a pair of physical 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.

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. Consequently, a redundant Ethernet interface can remain active even if one of its child interfaces becomes unavailable. Only one child interface of a redundant Ethernet interface can accept and send traffic at a time. A redundancy group *x* can contain up to 15 redundant Ethernet interfaces.

### *Related Topics*

- Understanding Chassis Cluster Formation on page 316
- Understanding Redundancy Groups on page 317
- Configuring Redundant Ethernet Interfaces on page 355
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326
- Understanding Failover on page 328

## Understanding the Control Plane

The control plane 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 primary Routing Engine fails, the secondary one is ready to assume control.

| **Before You Begin** |
| --- |
| Read “Understanding Chassis Cluster Formation” on page 316 |

This topic includes:



The control plane:

- 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 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 running on the primary Routing Engine maintains state for the entire cluster, and only processes running on its node can update state information. The primary 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. In a J-series chassis cluster, the control link is a physical link connecting the control interfaces on the two nodes of the cluster (the pair ge-0/0/0 and ge-0/0/3).

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.

On an SRX 5600 and SRX 5800 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 SRX 5600 or SRX 5800 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.

### 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.

### *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, and you must reboot it. When you reboot it, the node synchronizes its 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.

### *Related Topics*

- Understanding the Control Plane on page 323
- Understanding Redundancy Groups on page 317
- Understanding Redundant Ethernet Interfaces on page 322
- Understanding Failover on page 328

## Understanding the Data Plane

The data plane 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 it is not the local one.)

To provide for session (or flow) redundancy, the data plane synchronizes its state by sending special payload packets called real-time 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 gives RTOs transmission priority over transit traffic.

| **Before You Begin** |
| --- |
| Read “Understanding Chassis Cluster Formation” on page 316 |

This topic includes:

- About Session RTOs on page 326
- About the Fabric Data Link on page 327
- About Data Forwarding on page 327
- About Fabric Data Link Failure and Recovery on page 328
- Related Topics on page 328

### *About Session RTOs*

The data plane creates RTOs for UDP and TCP sessions and tracks state changes. It also synchronizes traffic for IPv4 passthrough protocols such as 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 Gigabit 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.

The fabric data link does not support fragmentation. To accommodate this state, jumbo frame support is enabled by default on the link with a maximum 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

**NOTE:** In this release, active/active chassis cluster (that is, cross-box data forwarding over the fabric interface) is not supported on an SRX-series chassis cluster.

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. (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:** For Intrusion Detection and Prevention (IDP) running on a chassis cluster, when there is a fabric data link failure, IDP processing will not be continued for existing sessions. (IDP processing will resume for sessions that are created after data link recovery.)

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 real-time objects (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 323
- Understanding Chassis Cluster Formation on page 316
- Understanding Redundancy Groups on page 317

## 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 316
- Understanding Redundancy Groups on page 317
- Understanding Redundant Ethernet Interfaces on page 322
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326

This topic includes:



### *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. For more information, refer to “Redundancy Group Interface Monitoring” on page 321.

### *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.

**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 primary Routing Engine. This failover could result in loss of state, such as routing state, and degrade performance by introducing system churn.

## 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 66 on page 330 shows two J-series devices connected using the built-in interfaces for both the fabric and control links.

**Figure 66: Connecting J-series Devices in a Cluster (J6350 Devices)**

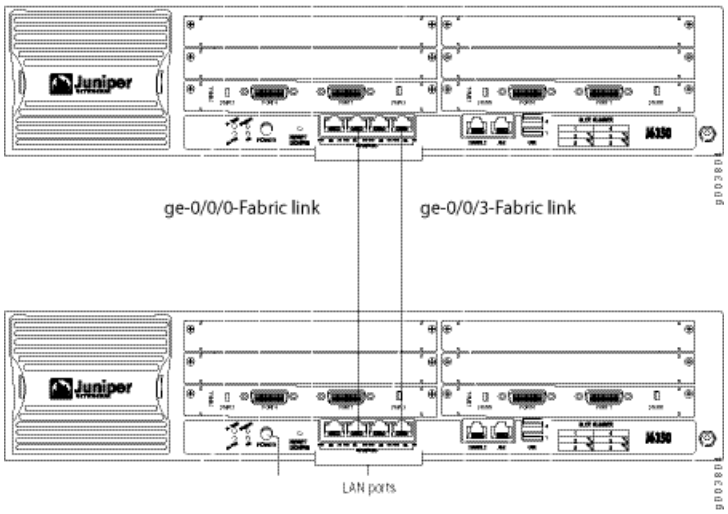


## Hardware Setup for SRX-series Chassis Clusters

To create an SRX-series chassis cluster, you must physically connect a pair of the same kind of supported SRX-series devices back-to-back over a pair of Gigabit Ethernet connections or a pair of 10-Gigabit Ethernet connections. The connection that serves as the control link must be the built-in controller port on each device. The fabric link connection can be a combination of any pair of Gigabit Ethernet interfaces on the devices. Figure 67 on page 330 shows a pair of SRX-series devices connected using the built-in interfaces for both the fabric and control links.

**Figure 67: Connecting SRX-series Devices in a Cluster (SRX 5800 Devices)**

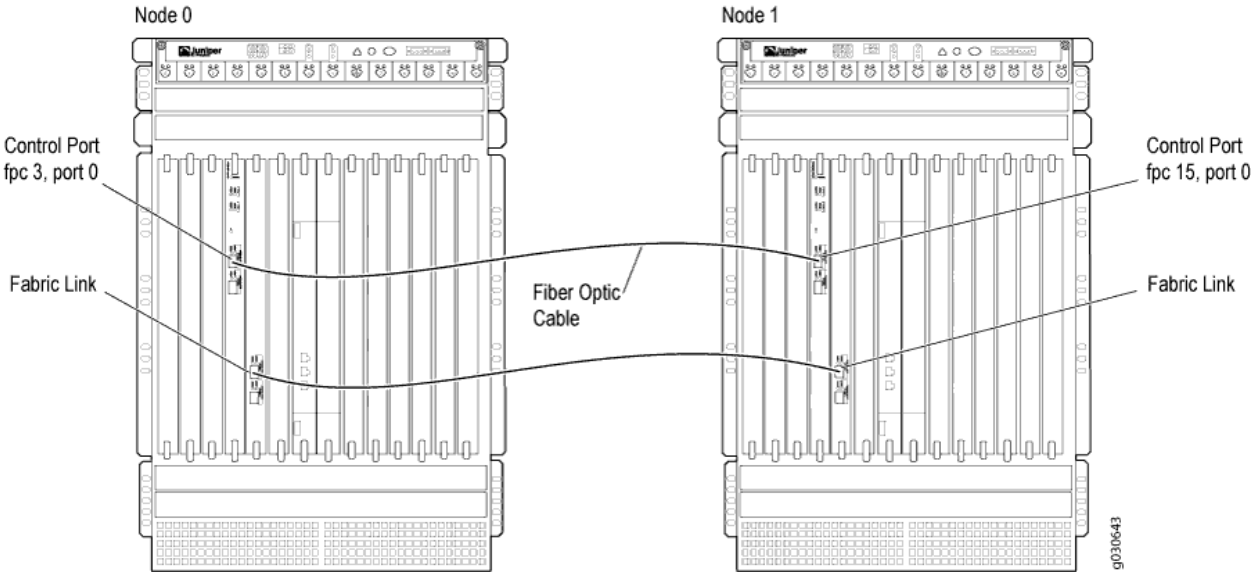

## What Happens When You Enable Chassis Cluster

After wiring the two devices together as described in "Hardware Setup for J-series Chassis Clusters" on page 329 or "Hardware Setup for SRX-series Chassis Clusters" on page 330, 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 317.)

Figure 68 on page 332 shows how the PIM slots are numbered on two nodes in a J-series chassis cluster. Figure 69 on page 334 shows how the SPC slots are numbered on two nodes in an SRX-series chassis cluster.

This topic includes:



### *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, two of the built-in interfaces are repurposed as the management and control interfaces and are automatically renamed fxp0 and fxp1, respectively. Table 59 on page 332 shows how these interfaces are renamed and how they are renumbered for a chassis cluster for the various J-series devices.

**Table 59: J-series Chassis Cluster Interface Naming Scheme**

| Model | Chassis | Interface | Interface | Management Interface | Control Interface |
|---|---|---|---|---|---|
| J2320 | Node 0 | ge-0/0/0 | ge-0/0/1 | fxp0 | fxp1 |
| | Node 1 | ge-4/0/0 | ge-4/0/1 | fxp0 | fxp1 |
| J2350 | Node 0 | ge-0/0/0 | ge-0/0/1 | fxp0 | fxp1 |
| | Node 1 | ge-6/0/0 | ge-6/0/1 | fxp0 | fxp1 |
| J4350 and J6350 | Node 0 | ge-0/0/0 | ge-0/0/1 | fxp0 | fxp1 |
| | Node 1 | ge-7/0/0 | ge-7/0/1 | fxp0 | fxp1 |

**NOTE:** See the *JUNOS Software with Enhanced Services 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. Internally, the cluster models a dual Routing Engine device. The primary Routing Engine propagates all its network and configuration settings and the current session information to the secondary. The secondary Routing Engine serves as the backup. As a single system, the cluster now has twice as many PIM slots. (See Figure 68 on page 332.)

**Figure 68: PIM Slot Numbering in a J-series Chassis Cluster (J6350 Devices)**

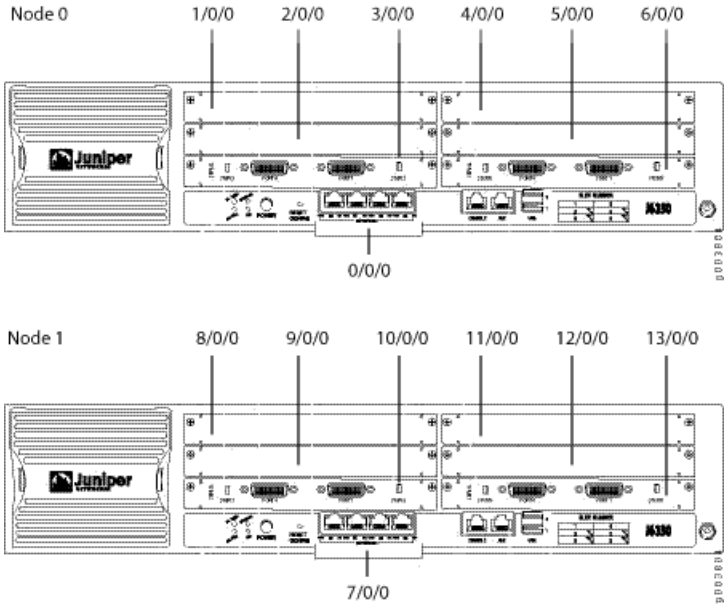

### *Node Interfaces on SRX-series Chassis Clusters*

Normally, on SRX-series devices, the built-in interfaces are numbered as follows:

| | | | |
|---|---|---|---|
| reth0 | reth1 | reth2 | reth3 |

For chassis clustering, SRX-series devices have a built-in management interface named fxp0. For SRX 5600 and 5800 devices, control interfaces are configured on SPCs. Table 60 on page 333 shows how an example of how these interfaces could be named and numbered for SRX-series chassis clusters.

**Table 60: Example of SRX-series Chassis Cluster Interface Naming Scheme**

| Model | Chassis | Interface | Interface | Management Interface | Control Interface |
|---|---|---|---|---|---|
| 5600 | Node 0 | ge-0/0/0 | ge-0/0/1 | fxp0 | fpc3, port 1 |
| | Node 1 | ge-6/0/0 | ge-6/0/1 | fxp0 | fpc9, port 1 |
| 5800 | Node 0 | ge-0/0/0 | ge-0/0/1 | fxp0 | fpc3, port 1 |
| | Node 1 | ge-12/0/0 | ge-12/0/1 | fxp0 | fpc 15, port 1 |

**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. Internally, the cluster models a dual Routing Engine device. The primary Routing Engine propagates all its network and configuration settings and the current session information to the secondary. The secondary Routing Engine serves as the backup. As a single system, the cluster now has twice as many FPC slots. (See Figure 69 on page 334.)

**Figure 69: FPC Slot Numbering in an SRX-series Chassis Cluster (SRX 5800 Devices)**

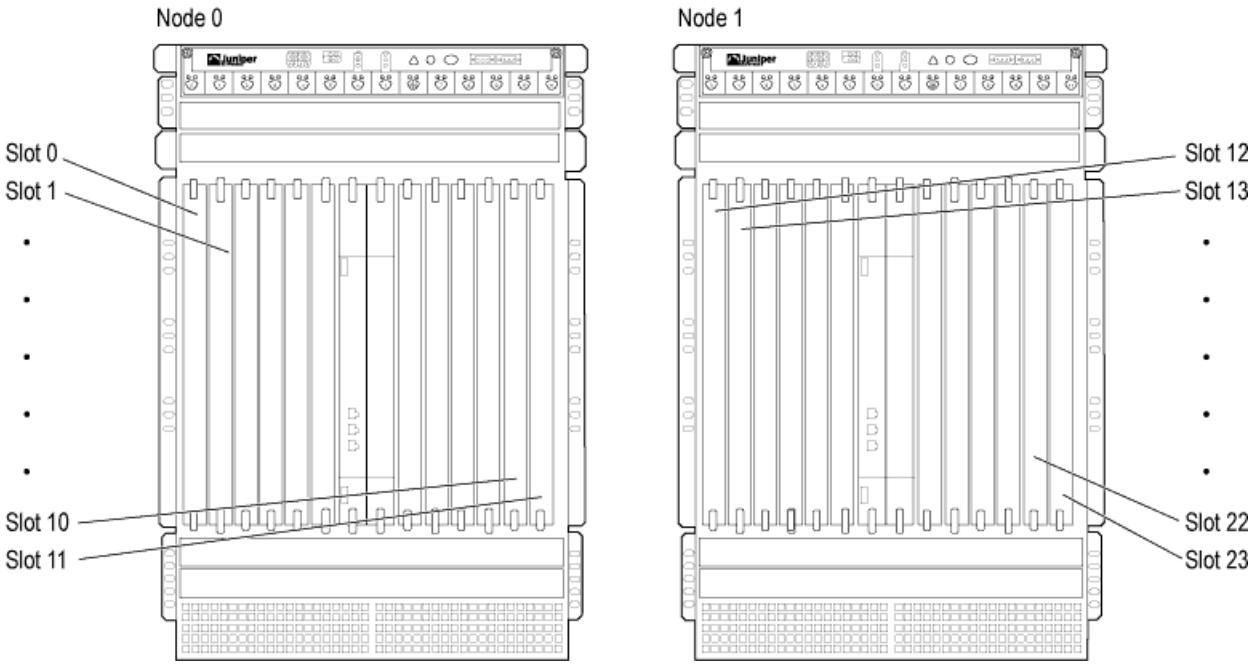


### *Management Interfaces on J-series Chassis Clusters*

The fxp0 interfaces function like standard management interfaces on 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.

### *Management Interfaces on SRX-series Chassis Clusters*

The fxp0 interfaces, when configured for active/active operations, function like standard management interfaces on SRX-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. The fabric link can be any pair of Gigabit Ethernet interfaces spanning the cluster.

### *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.

### *Related Topics*

- Understanding Chassis Cluster Formation on page 316
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326
- Disabling Chassis Cluster on page 365

## 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 337.

**Before You Begin**

For background information, read:

- Understanding Chassis Cluster Formation on page 316
- Hardware Setup for J-series Chassis Clusters on page 329
- What Happens When You Enable Chassis Cluster on page 331

**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 66 on page 330 shows nodes connected using the built-in ge-0/0/1 interface for the fabric link.

b. Connect the ge-0/0/2 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 *JUNOS Software with Enhanced Services 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.

4. On the other device, connect to the console port and use CLI operational mode commands to enable clustering:

   a. Identify the cluster the device is joining by setting the same cluster ID you set on 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 341.

6. Configure the cluster:

   To use J-Web Quick Configuration, see:

   a. Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342

   b. Configuring Redundant Ethernet Interfaces—Quick Configuration on page 345

   c. Configuring a Gigabit Interface—Quick Configuration on page 348

   To configure the cluster with the CLI, see:

   a. Configuring Chassis Cluster Information on page 352

   b. Configuring the Fabric on page 352

   c. Configuring Redundancy Groups on page 354

   d. Configuring Redundant Ethernet Interfaces—Quick Configuration on page 345

   e. Configuring Interface Monitoring on page 356

7. To initiate manual failover, see “Initiating a Manual Redundancy Group Failover” on page 357.

8. To configure conditional route advertisement over redundant Ethernet interfaces, see “Configuring Conditional Route Advertising” on page 358.

9. To verify the configuration, see “Verifying the Chassis Cluster Configuration” on page 361.

### *Related Topics*

- Understanding Chassis Cluster Formation on page 316
- What Happens When You Enable Chassis Cluster on page 331
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342
- Configuring Redundant Ethernet Interfaces—Quick Configuration on page 345
- Configuring Chassis Cluster Information on page 352
- Configuring the Fabric on page 352
- Understanding Redundancy Groups on page 317
- Understanding Redundant Ethernet Interfaces on page 322

## 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 335.

**Before You Begin**

For background information, read:

- Understanding Chassis Cluster Formation on page 316
- Hardware Setup for SRX-series Chassis Clusters on page 330
- What Happens When You Enable Chassis Cluster on page 331

**NOTE:** For SRX-series chassis clusters, the two nodes in a cluster must be identical models, but can have any combination of FPCs installed.

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 Gigabit Ethernet interfaces (or pair of 10-Gigabit Ethernet interfaces). The only requirement is that both interfaces be Gigabit Ethernet interfaces (or 10-Gigabit Ethernet interfaces).

      Figure 67 on page 330 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). For SRX 5600 and SRX 5800 devices, control ports should be

on corresponding slots in the two devices, with the following slot numbering offsets:

| Model | Offset |
|---|---|
| 5600 | 6 |
| 5800 | 12 |

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.
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.
5. On the other device, connect to the console port and use CLI operational mode commands to enable clustering:
   a. Identify the cluster the device is joining by setting the same cluster ID you set on first node.
   b. Identify the node by giving its own node ID and then reboot the system.
6. Configure the management interfaces on the cluster.

   See “Configuring the Management Interface” on page 341.

7. Configure the cluster:

   To use J-Web Quick Configuration, see:

   a. Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342
   b. Configuring Redundant Ethernet Interfaces—Quick Configuration on page 345
   c. Configuring a Gigabit Interface—Quick Configuration on page 348

   To configure the cluster with the CLI, see:

   a. Configuring Chassis Cluster Information on page 352
   b. Configuring the Fabric on page 352
   c. Configuring Redundancy Groups on page 354
   d. Configuring Redundant Ethernet Interfaces—Quick Configuration on page 345
   e. Configuring Interface Monitoring on page 356

8. To initiate manual failover, see “Initiating a Manual Redundancy Group Failover” on page 357.
9. To configure conditional route advertisement over redundant Ethernet interfaces, see “Configuring Conditional Route Advertising” on page 358.
10. To verify the configuration, see “Verifying the Chassis Cluster Configuration” on page 361.

### *Related Topics*

- Understanding Chassis Cluster Formation on page 316
- What Happens When You Enable Chassis Cluster on page 331
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342
- Configuring Redundant Ethernet Interfaces—Quick Configuration on page 345
- Configuring Chassis Cluster Information on page 352
- Configuring the Fabric on page 352
- Understanding Redundancy Groups on page 317
- Understanding Redundant Ethernet Interfaces on page 322

## 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 316
- Understanding Redundancy Groups on page 317

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 352. 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

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 node 0 cluster-id 1
  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 node 1 cluster-id 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 59 on page 332). 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 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 316
- Setting the Node ID and Cluster ID on page 339

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 and the other as the secondary.

In an SRX-series chassis cluster, the fxp0 interface is a port on the Routing Engine (RE) card.

**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 341
- Related Topics on page 342

### *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 to 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 352
- Configuring the Fabric on page 352
- Configuring Redundant Ethernet Interfaces on page 355
- Configuring Conditional Route Advertising on page 358
- Verifying the Chassis Cluster Configuration on page 361

## 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 70 on page 343.

**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 339 for details on information you must configure first
- Configuring Chassis Cluster Information on page 352

Figure 70 on page 343 shows the Chassis Cluster page.

**Figure 70: Chassis Cluster Page**

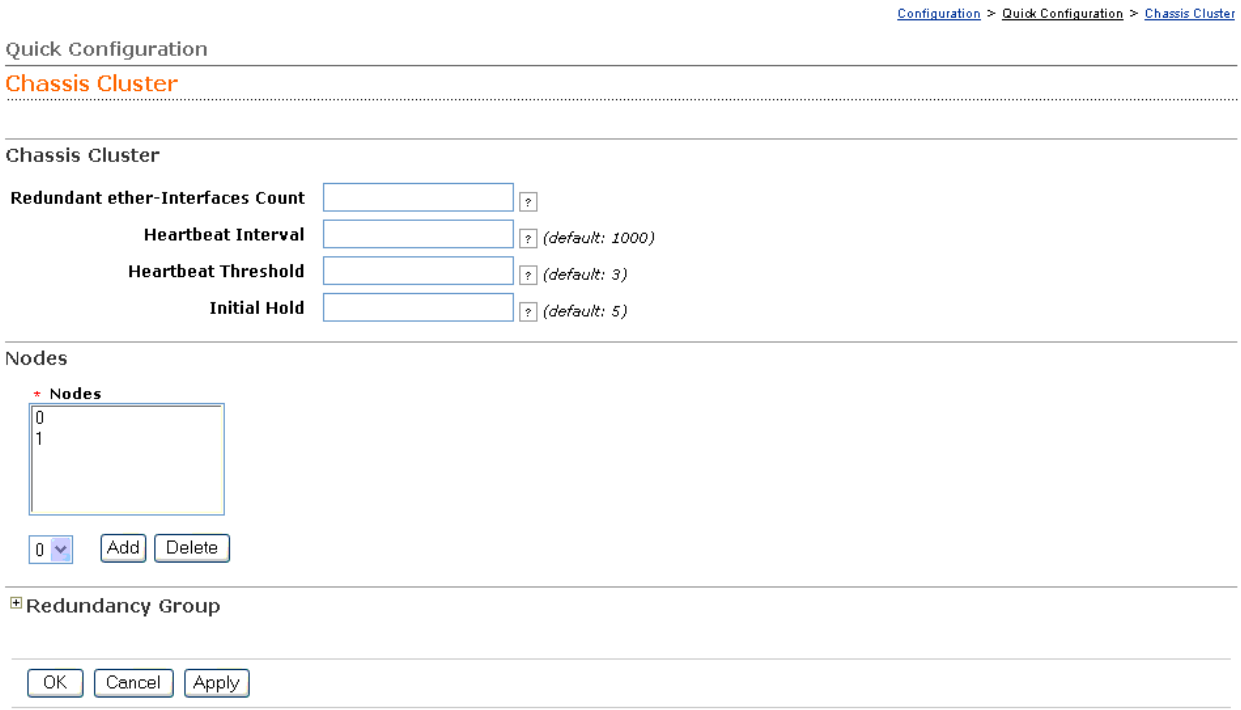


To configure chassis cluster with J-Web Quick Configuration:

1. Select Configuration > Quick Configuration > Chassis Cluster.
2. Fill in the parameter settings as described in Table 61 on page 343.
3. Click one of the following buttons:
   - To apply the configuration and stay on the Chassis Cluster 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.

Go on to "Configuring Redundant Ethernet Interfaces—Quick Configuration" on page 345.

**Table 61: 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. |

**Table 61: Chassis Cluster Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Heartbeat Interval | Specifies, in milliseconds, how often heartbeats are sent out to the other nodes in the cluster.<br><br>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.<br><br>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.<br><br>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.<br>■ To modify information for a redundancy group, select the corresponding group number. Specify a value between 1 and 255.<br>■ To create a new redundancy group:<br>  ■ Assign a group ID to the redundancy group. Specify a value between 1 and 255.<br>  ■ Configure the redundancy group with information described in the rest of this table. |
| 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.<br><br>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.<br><br>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. |

**Table 61: Chassis Cluster Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| 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 317
- Setting the Node ID and Cluster ID on page 339
- Configuring Redundant Ethernet Interfaces on page 355
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326

## 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 pseudo interface 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 322 |

Figure 71 on page 346 shows the Interface page.

**Figure 71: Interface Page**

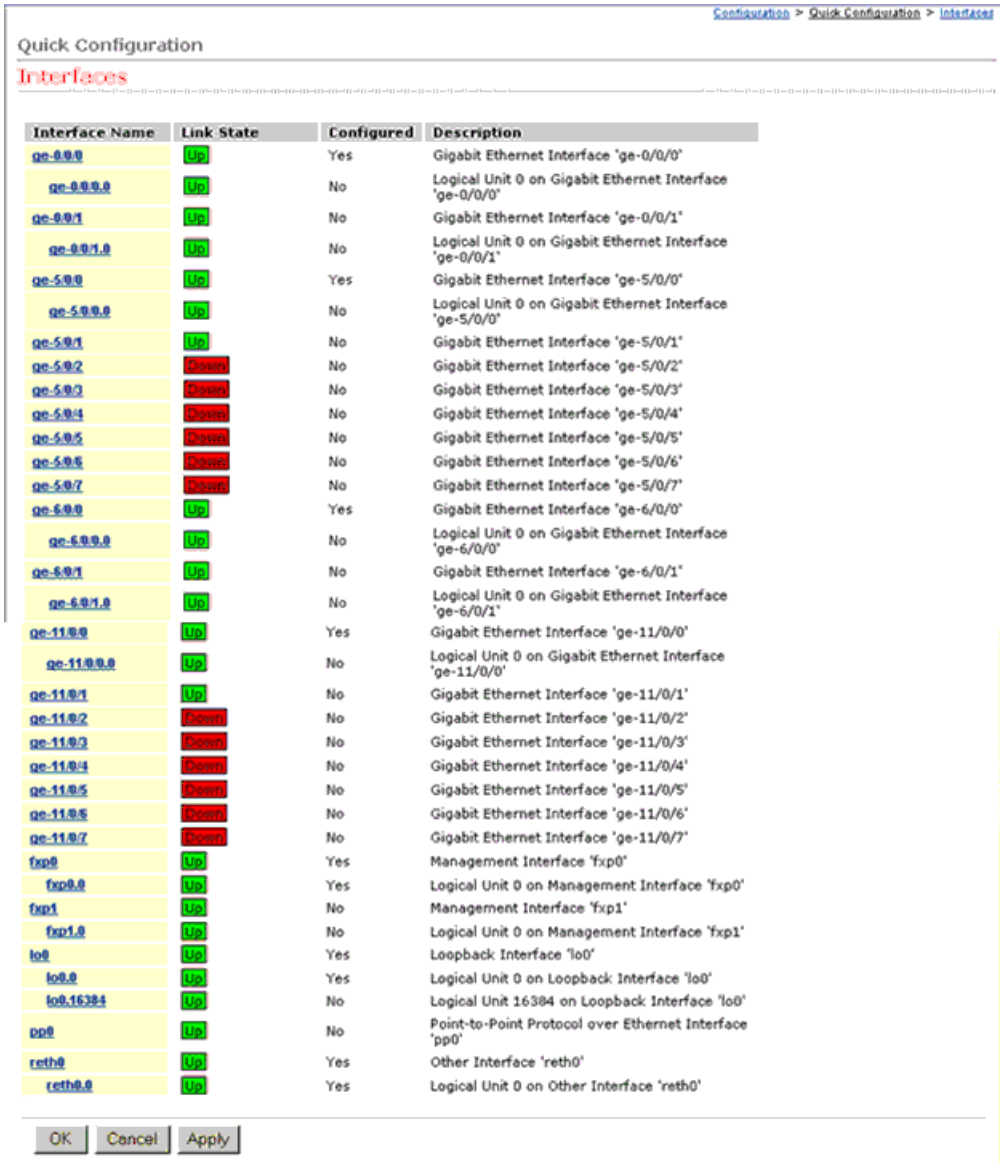


Figure 72 on page 347 shows the Redundant Ethernet Interface Configuration page.

**Figure 72: Redundant Ethernet Interface Configuration page**

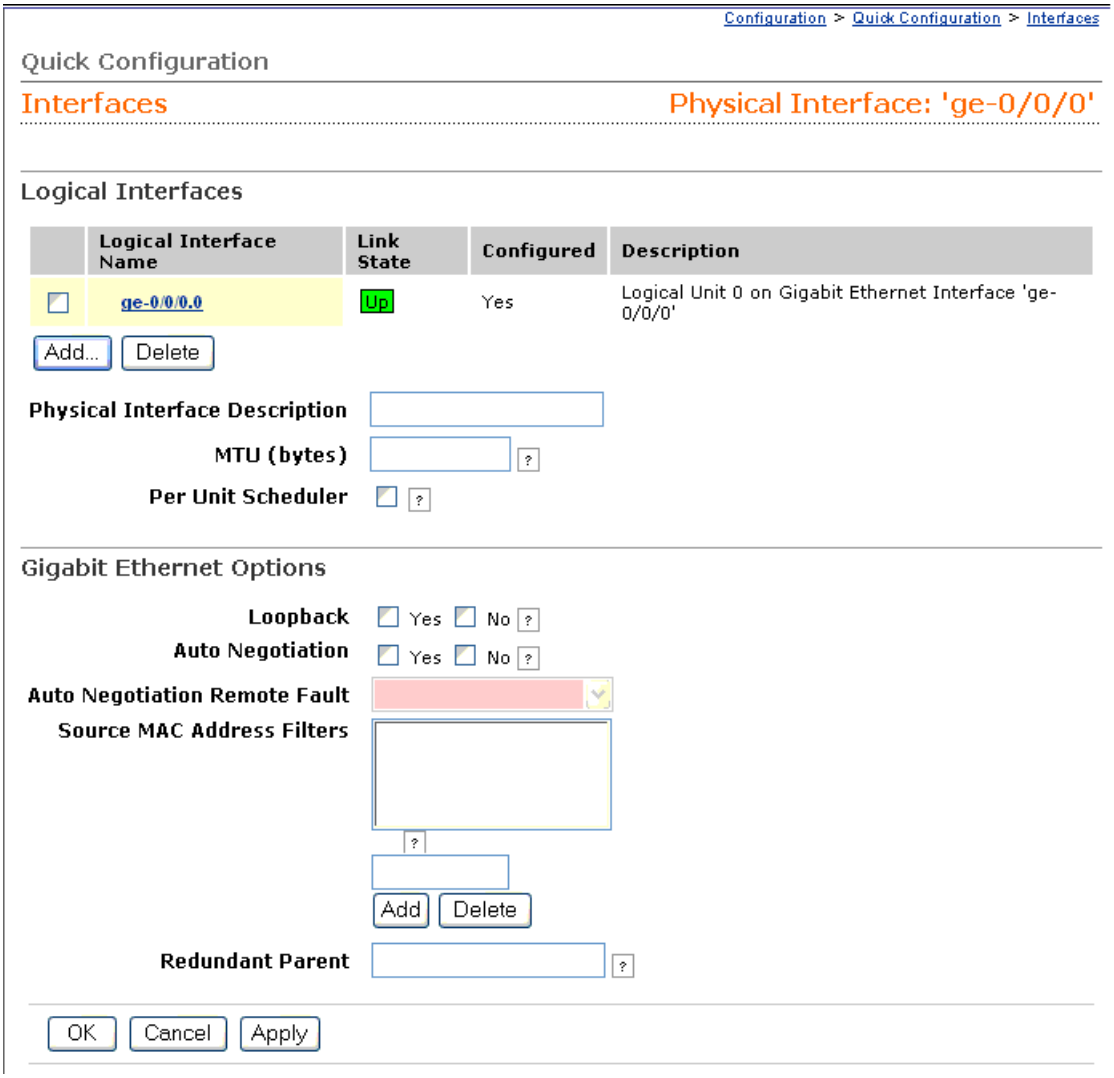


To configure redundant Ethernet interfaces with J-Web Quick Configuration:

1. Select Configuration > Quick Configuration > Interfaces.
2. Click an interface name to group physical Ethernet interfaces for redundancy. See Figure 72 on page 347.
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 62 on page 348. 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 62: 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.<br><br>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 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 a Gigabit Interface—Quick Configuration

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

Figure 73 on page 349 shows the Gigabit Ethernet Interface Quick Configuration page.

**Figure 73: Gigabit Ethernet Interface Quick Configuration**

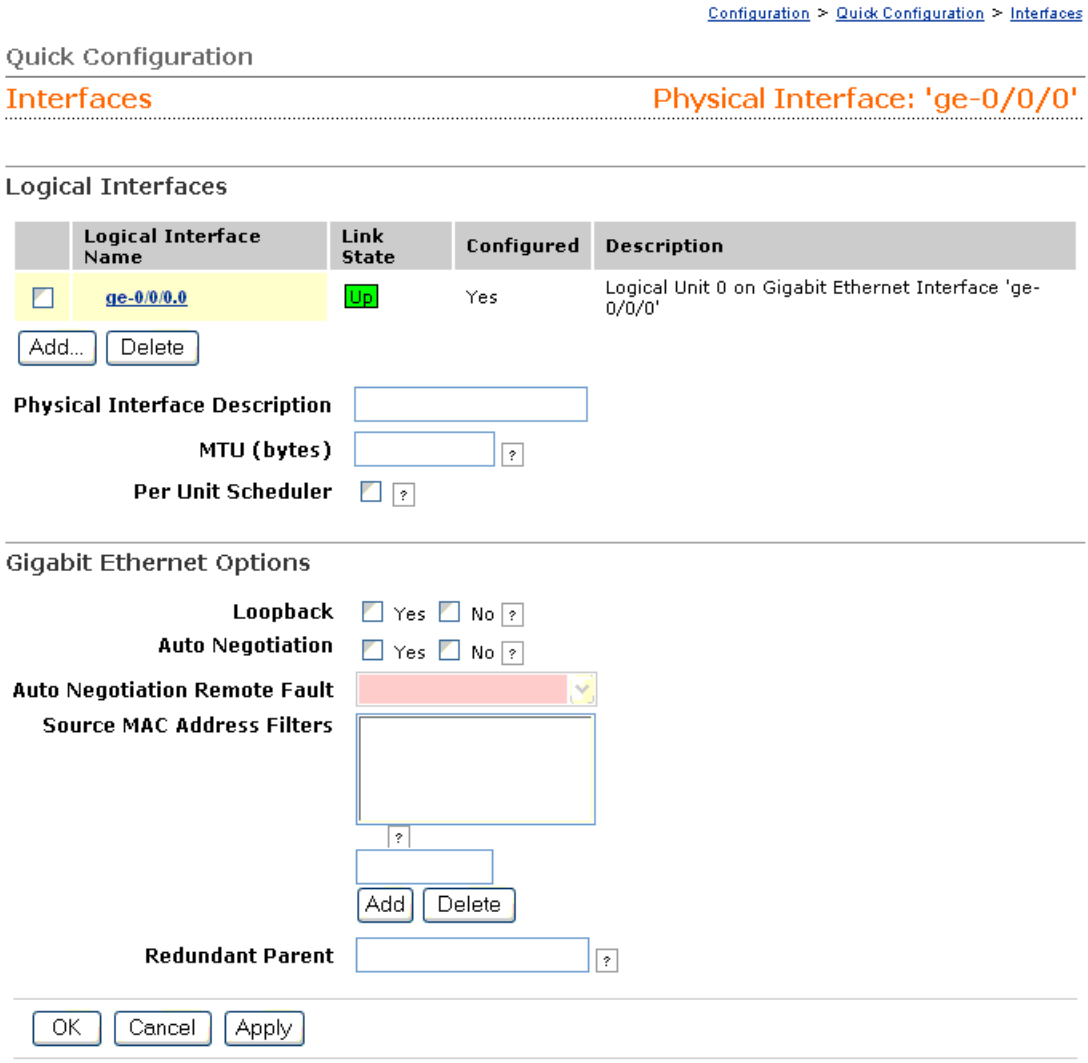


1. Select Configuration > Quick Configuration > Interfaces.
2. Click a Gigabit Ethernet interface. The properties you can configure on a Gigabit Ethernet interface appear, as shown in Figure 11 on page 66.
3. Fill in the information as described in Table 63 on page 350.
4. Click one of the following 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.

5. Verify that the Gigabit Ethernet interface is configured correctly by seeing the *JUNOS Software Interfaces and Routing Configuration Guide*

**Table 63: Gigabit Ethernet Quick Configuration Page Summary**

| Field | Function | Action |
|---|---|---|
| Logical Interfaces | | |
| 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:<br>10.10.10.10/24<br>2. Click Add.<br>3. Click OK.<br>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.<br><br>Static ARP entries associate the IP addresses and MAC addresses of nodes on the same subnet, enabling a 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:<br>10.10.10.1 |
| MAC Address | Specifies the hardware media access control (MAC) address associated with the ARP address.<br><br>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:<br>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.<br>■ 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. |

**Table 63: Gigabit Ethernet Quick Configuration Page Summary** *(continued)*

| Field | Function | Action |
|---|---|---|
| MTU (bytes) | Specifies, in bytes, the maximum transmission unit size for the Gigabit Ethernet interface. | Type a value between 256 and 9014. The default MTU for Gigabit Ethernet interfaces is 1514. |
| Per Unit Scheduler | Enables scheduling on logical interfaces.<br><br>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.<br>■ 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 autonegotiation.<br><br>By default, Gigabit Ethernet interfaces automatically negotiate the link mode and speed settings. If you disable autonegotiation and do not manually configure link mode and speed, the link is negotiated at 1000 Mbps, full duplex.<br><br>When you configure both the link mode and the speed, the link negotiates with the manually configured settings whether autonegotiation is enabled or disabled. | Select Yes to enable autonegotiation, or select No to disable it. By default, autonegotiation is enabled. |
| Auto Negotiation Remote Fault | Indicates the autonegotiation remote fault value. | Select the autonegotiation remote fault value from the list of options given. This field is enabled only if autonegotiation 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. |
| Redundant Parent | Specifies the parent redundant Ethernet interface (reth) of the physical child Ethernet interface. | 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 316
- Setting the Node ID and Cluster ID on page 339
- Configuring a Chassis Cluster and Redundancy Groups—Quick Configuration on page 342

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

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

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 316
- Setting the Node ID and Cluster ID on page 339
- Understanding the Control Plane on page 323

In a J-series chassis cluster, you can configure any pair of Gigabit Ethernet interfaces to serve as the fabric between nodes. In an SRX-series 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. 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

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 352
- Configuring Redundant Ethernet Interfaces on page 355
- Configuring Conditional Route Advertising on page 358
- Verifying the Chassis Cluster Configuration on page 361

## 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 317
- Setting the Node ID and Cluster ID on page 339
- Configuring Chassis Cluster Information on page 352
- Configuring the Fabric on page 352

Before you can create redundant Ethernet interfaces you must create their redundancy groups.

This topic includes:



### *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 just 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 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
```

## Configuring Redundant Ethernet Interfaces

A redundant Ethernet interface is a pseudo interface 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 339
- Configuring Chassis Cluster Information on page 352
- Configuring the Fabric on page 352
- Configuring Redundancy Groups on page 354
- Understanding Redundant Ethernet Interfaces on page 322

This topic includes:



### *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 355
- Configuring Conditional Route Advertising on page 358
- Verifying the Chassis Cluster Configuration on page 361

## 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 317
- Understanding Redundant Ethernet Interfaces on page 322

This topic includes:

- CLI Configuration on page 356
- Related Topics on page 357

### *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 355
- Configuring Redundancy Groups on page 354

## 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 354
- Configuring Redundant Ethernet Interfaces on page 355

After a manual failover, the new primary continues in that role until there is a failback. If there is a failback, the manual failover is lost and state election is made based on priority and preempt settings. A failback in manual failover mode can occur if the primary node fails or if the threshold of a redundancy group 0 reaches 0.

**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 primary 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 357

### *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: 3
Node name                  Priority     Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 0
    node0                   254         primary   no       no
    node1                   2           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 the primary:

```
{primary:node1}
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.

```
{primary:node1}
user@host> show chassis cluster status redundancy-group 0
Cluster ID: 1
    Node name              Priority     Status        Preempt  Manual failover

Redundancy-group: 0, Failover count: 1
    node0                  254          secondary     no       yes
    node1                  2            primary       no       yes
```

Output to this command shows that node 1 is now primary.

You can reset the failover for redundancy groups using the request command. This change is propagated across the cluster.

```
{primary:node1}
user@host> request chassis cluster failover reset redundancy-group 0 node 0
node0:
- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
Successfully reset manual failover for redundancy group 1
node1:
- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
```

With back-to-back failovers, 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 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 316

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 74 on page 359 shows a typical scenario, with a redundant Ethernet interface connecting the corporate LAN, through a chassis cluster, to an external network segment.

**Figure 74: Conditional Route Advertising**

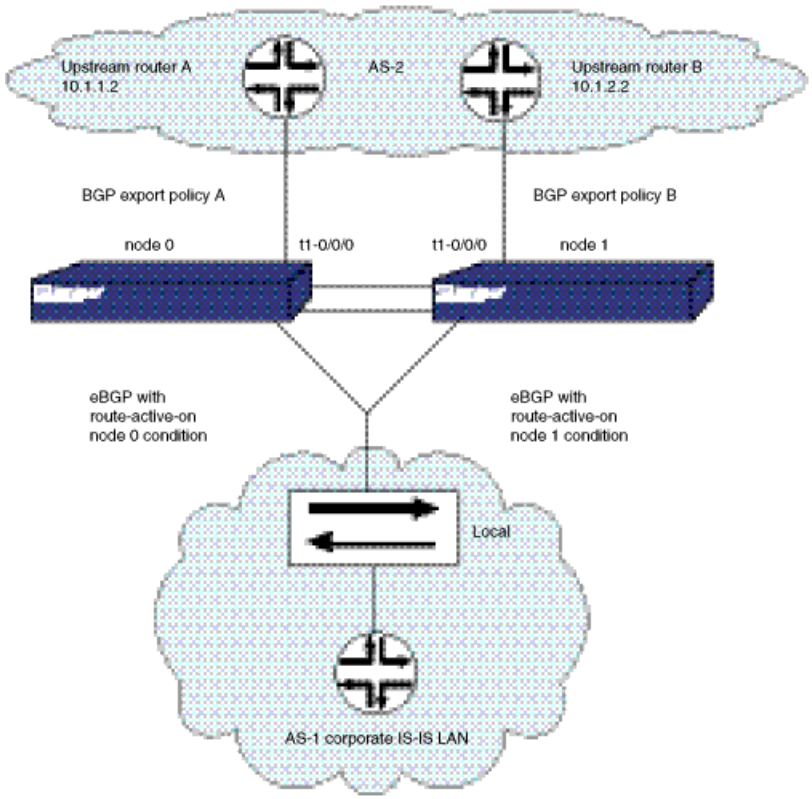


## *CLI Configuration*

As illustrated in Figure 74 on page 359, 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 using this same concept.

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 361

## Verifying the Chassis Cluster Configuration

To verify the chassis cluster configuration, perform the following tasks:



### *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 361
- Verifying Chassis Cluster Statistics on page 362
- Verifying Chassis Cluster Status on page 364
- Verifying Chassis Cluster Redundancy Group Status on page 364


### *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 name: em0.0

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 output shows the state of fxp1, the control link between the nodes, and provides information about the link state. The output also shows the state of the fabric interface between the nodes and provides information about traffic on that link.

### *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:
    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
```

**What it Means** The 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.

**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:
    Heartbeat packets sent: 124
    Heartbeat packets received: 125
Fabric link statistics:
    Probes sent: 124
    Probes received: 125
```

**What it Means** The output shows the control link statistics (heartbeats sent and received) and the fabric link statistics (probes sent and received).

**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
```

**What it Means** The output shows the failover status of the chassis cluster in addition to information about the chassis cluster redundancy groups.

**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.

## Upgrading Chassis Cluster

To upgrade a chassis cluster:

**NOTE:** During cluster upgrade, a service disruption of about three to five 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.

### *Related Topics*

- Understanding Chassis Cluster Formation on page 316
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326
- What Happens When You Enable Chassis Cluster on page 331

## 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 316
- Understanding the Control Plane on page 323
- Understanding the Data Plane on page 326
- What Happens When You Enable Chassis Cluster on page 331

## Chassis Cluster Configuration Scenarios

This section describes the following deployment scenarios for J-series chassis clusters:

- Active/Passive Chassis Cluster Scenario on page 366
- Asymmetric Routing Chassis Cluster Scenario on page 371
- Active/Active Full Mesh Chassis Cluster Scenario on page 375

**Before You Begin**

For background information, read:

- Understanding Chassis Cluster Formation on page 316
- Hardware Setup for J-series Chassis Clusters on page 329
- What Happens When You Enable Chassis Cluster on page 331
- Creating a J-series Chassis Cluster—Overview on page 335

### *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. When a failure occurs, the backup device becomes master and controls all forwarding.

**Figure 75: Active/Passive Chassis Cluster Scenario**

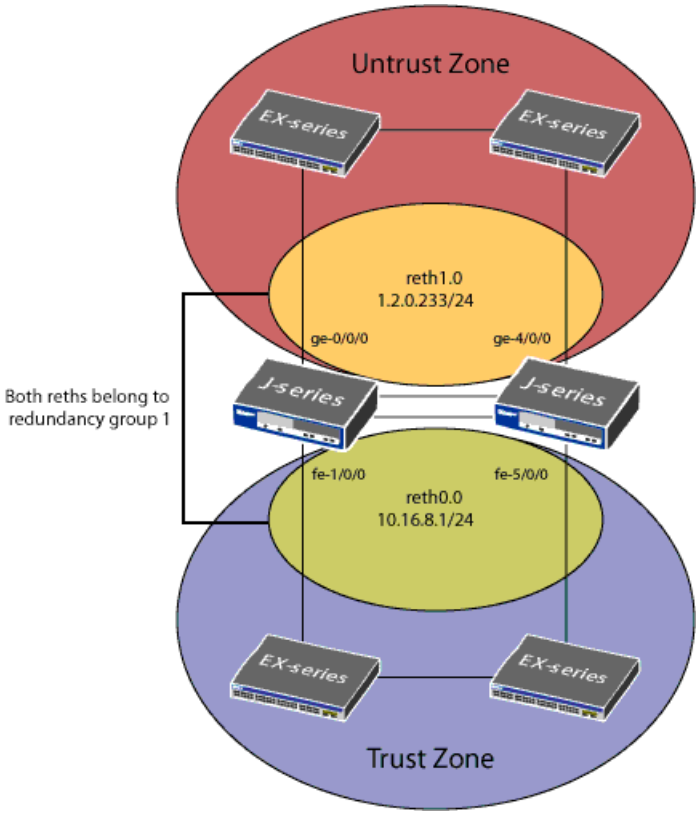


An active/passive chassis cluster can be achieved using redundant Ethernet interfaces reth0 and reth1. If any of the interfaces in a reth fails, the group is declared inactive and all its interfaces fail over to the other group. This configuration minimizes the traffic over the fabric link because only one node in the cluster will forward traffic at any given time.

## CLI

1. Basic Chassis Cluster

```
user@host> set chassis cluster node 0 cluster-id 1
warning: A reboot is required for chassis cluster to be enabled
user@host> set chassis cluster node 1 cluster-id 1 reboot
Successfully enabled chassis cluster. Going to reboot now.
```

2. Management Interface

```
{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 preempt
{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
```

(Optional) Redundancy Group 0 (to monitor physical interfaces if data processing and control plane functions are in same node)

```
{primary:node1}
user@host# set chassis cluster redundancy-group 0 interface-monitor fe-1/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 0 interface-monitor fe-5/0/0
  weight 255
```

```
{primary:node1}
user@host# set chassis cluster redundancy-group 0 interface-monitor ge-0/0/0
  weight 255
{primary:node1}
user@host# set chassis cluster redundancy-group 0 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 Zone

```
{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. Access

```
{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

1. Basic Chassis Cluster

See “CLI” on page 367.

2. Management Interface

See “CLI” on page 367.

3. Fabric Interface

See “CLI” on page 367.

4. Redundancy Groups

Configuration > Quick Configuration > 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
Preempt: Select checkbox.
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

(Optional) Redundancy Group 0 (to monitor physical interfaces if data processing and control plane functions are in same node)
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

Configuration > Quick Configuration > Interfaces

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 “CLI” on page 367

6. Security Zone

   See “CLI” on page 367.

7. Access

   See “CLI” on page 367.

### *Asymmetric Routing Chassis Cluster Scenario*

In this case, chassis cluster makes use of its asymmetric routing capability. Traffic received by a node is matched against that node’s session table. The result of this lookup determines whether that node processes the session or forwards it to the other node over the fabric link. Sessions are anchored on the node in which the first packet created the session, and the session is synced to the other node. If traffic is received on the node in which the session is not anchored, those packets are forwarded over the fabric link.

**Figure 76: Asymmetric Routing Chassis Cluster Scenario**

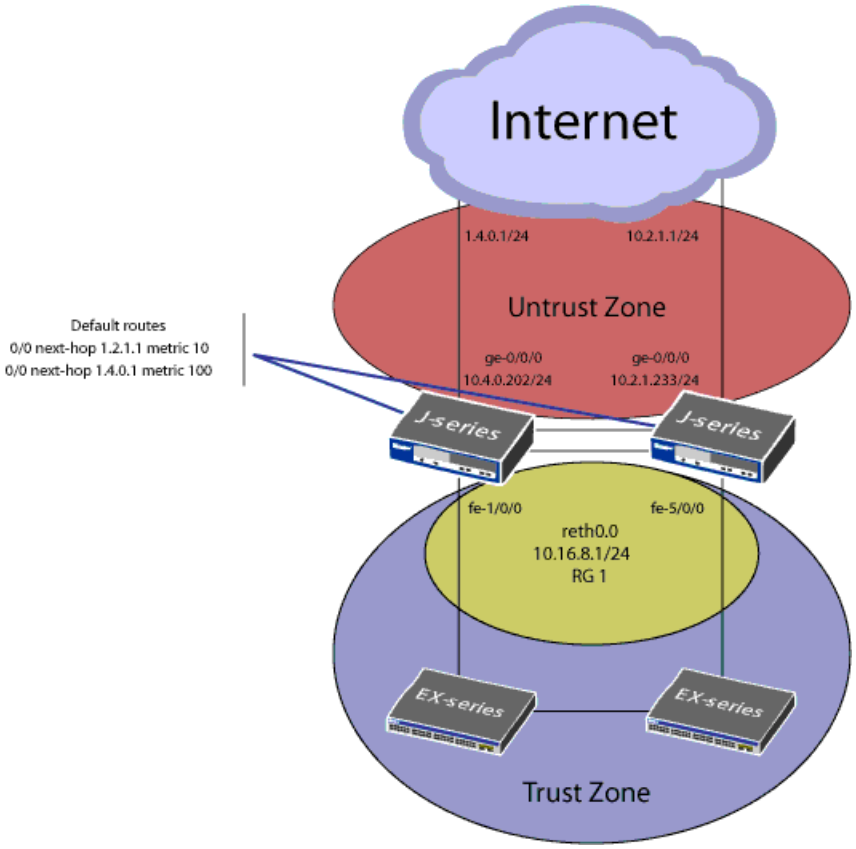


Figure 76 on page 372 illustrates how asymmetric routing is supported in a chassis cluster. In this scenario, two Internet connections are used, with one being preferred. The connection to the Trust zone is done using a reth interface to provide LAN redundancy for the devices in the Trust zone. This 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 to interface fe-1/0/0 belonging to reth0.0 on node 0. Because the primary Internet connection is in node 0, sessions are created in both node 0 and node 1, but active in only 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—traffic arrives at node 0, is processed at node 0 because the session is anchored to this node, and is sent to node 1 through the fabric interface, where node 1 forwards it through 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-4/0/0 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-4/0/0.

In this chassis cluster configuration, redundancy group 1 is used to control the reth interface connected to the Trust zone. This redundancy group (and, therefore, reth0) fails over only if interface fe-1/0/0 or fe- 5/0/0 fails, but not if any of the interfaces connected to the Internet fail.

#### CLI

**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-4/0/1
```

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 fe-1/0/1 disable
{primary:node1}
user@host# set interfaces ge-4/0/0 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 Zone

```
{primary:node1}
user@host# set security zones security-zone Untrust interfaces ge-0/0/0.0
  host-inboundtraffic system-services dhcp
{primary:node1}
user@host# set security zones security-zone Untrust interfaces ge-4/0/0.0
  host-inboundtraffic system-services dhcp
```

6. Access

```
{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

**NOTE:** First, do basic chassis cluster and management interfaces setup.

1. Fabric Interface

See “CLI” on page 373.

2. Redundancy Groups

Configuration > Quick Configuration > 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

   Configuration > Quick Configuration > 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 of Step 5, see "CLI" on page 367

4. Static Routes (one to each ISP, with preferred route through ge-0/0/0)

   Configuration > Quick Configuration > Routing and Protocols > 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 Zone

   See "CLI" on page 373.

6. Access

   See "CLI" on page 373.

### *Active/Active Full Mesh Chassis Cluster Scenario*

This scenario is found in medium to large deployments in which services routers are placed between two pairs of routers. OSPF is used to control the traffic flow through

the nodes in the cluster, and JSRP is used to synchronize the sessions between the two nodes. Since asymmetric routing is supported, traffic does not have to be forced to a particular node. If a failure occurs and return traffic for a session arrives asymmetrically, the fabric link is used to send the traffic back to the node in which sessions are active, which will be the node hosting the egress interface for that particular session.

This scenario benefits from the use of full-mesh connectivity between the devices, which improves the resiliency of the network while eliminating the need to add extra switches between the firewalls and routers and also reducing the number of network failure points.

**Figure 77: Active/Active Full Mesh Chassis Cluster Scenario**

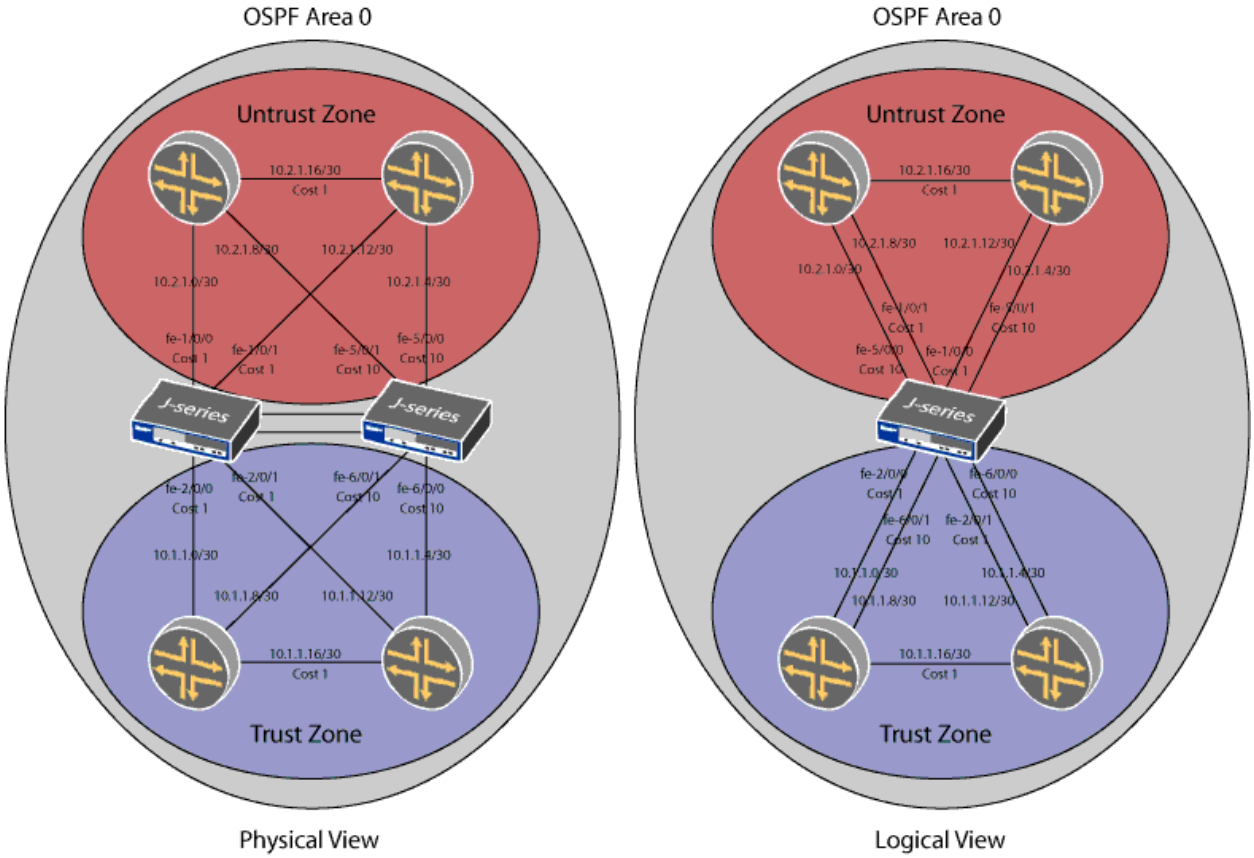

Chapter 14

# Internet Protocol Security (IPsec)

IP Security (IPsec) is a suite of related protocols for cryptographically securing communications at the IP Packet Layer. IPsec consists of two modes and two main protocols:

- Transport and tunnel modes
- The Authentication Header (AH) protocol for authentication and the Encapsulating Security Payload (ESP) protocol for encryption (and authentication)

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). See RFC 2407 and RFC 2408. See Figure 78 on page 377.

**Figure 78: IPsec Architecture**

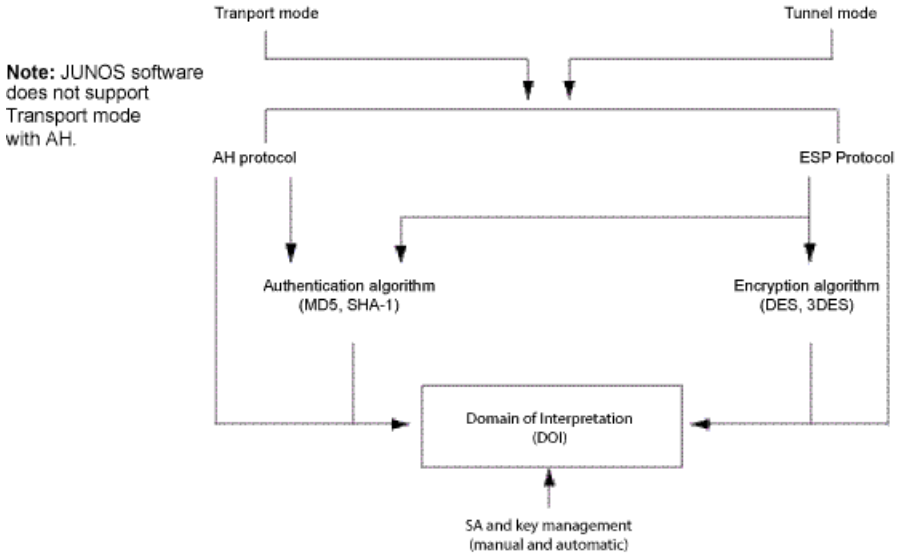


**NOTE:** The IPsec domain of interpretation (DOI) is a document containing definitions for all the security parameters required for the successful negotiation of a VPN tunnel—essentially, all the attributes required for SA and IKE negotiations.

This section includes:



## Virtual Private Networks (VPNs)

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) 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 either Transport or Tunnel mode (see "Understanding IPsec Operational Modes" on page 380). It refers to the IPsec connection.

### *Security Associations (SAs)*

An IPsec tunnel consists of a pair of unidirectional security associations (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.

Through the SA, an IPsec tunnel can provide the following security functions:

- Privacy (thourgh encryption)
- Content integrity (thourgh data authentication)
- Sender authentication and—if using certificates—nonrepudiation (thourgh 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.

### *Key Management*

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

### *Related Topics*


- Understanding IPsec Operational Modes on page 380
- Understanding IPsec Security Protocols on page 382
- Understanding IPsec Security Associations (SAs) on page 384
- Understanding IPsec Key Management on page 385
- Understanding IKE and IPsec Packets on page 386
- Understanding IPsec Tunnel Negotiation on page 391

## Understanding IPsec Operational Modes

IPsec operates in one of two modes—Transport or Tunnel.

| Before You Begin |
| --- |
| For background information, read “Internet Protocol Security (IPsec)” on page 377. |

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 router or firewall, you must use Tunnel mode. Juniper Networks devices always operate in Tunnel mode for IPsec tunnels.

### *Transport Mode*

The original IP packet is not encapsulated within another IP packet, as shown in Figure 79 on page 380. The entire packet can be authenticated with the Authentication Header protocol (AH), the payload can be encrypted with Encapsulating Security Payload protocol (ESP), and the original header remains in plaintext as it is sent across the WAN.

**Figure 79: Transport Modes**

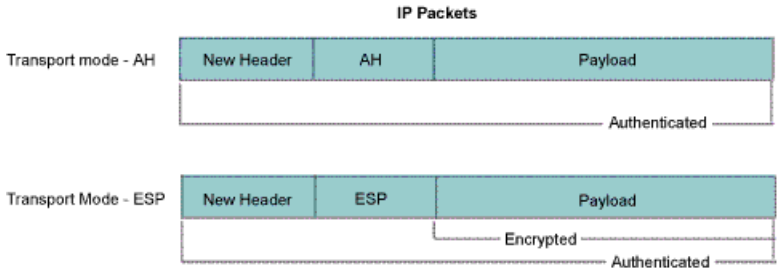


### *Tunnel Mode*

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 80 on page 381. 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 80: Tunnel Modes**

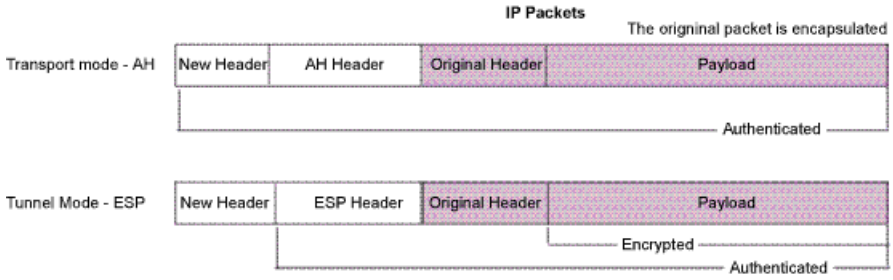


In a site-to-site VPN, the source and destination addresses used in the new header are the IP addresses of the outgoing interface (in NAT or Route mode) or the VLAN1 IP address (in Transparent mode); the source and destination addresses of the encapsulated packets are the addresses of the ultimate endpoints of the connection. See Figure 81 on page 381.

**Figure 81: Site-to-Site VPN in Tunnel Mode**

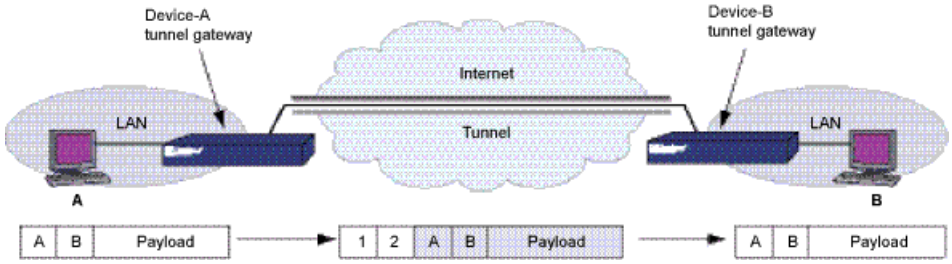


In a dial-up VPN, there is no tunnel gateway on the VPN dial-up client end of the tunnel; the tunnel extends directly to the client itself (see Figure 82 on page 382). In this case, on packets sent from the dial-up client, both the new header and the encapsulated original header have the same IP address: that of the client's computer. See Figure 82 on page 382.

**NOTE:** Some VPN clients such as the NetScreen-Remote allow you to define a virtual inner IP address. 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 dia-lup client is the source IP address in the outer header.

**Figure 82: Dial-up VPN in Tunnel Mode**

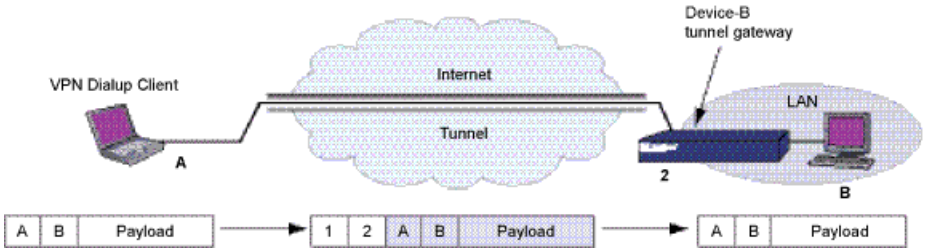


### *Related Topics*

- Understanding IPsec Security Protocols on page 382
- Understanding IPsec Security Associations (SAs) on page 384
- Understanding IPsec Key Management on page 385
- Understanding IKE and IPsec Packets on page 386
- Understanding IPsec Tunnel Negotiation on page 391
- Configuring an IKE IPsec Tunnel—Overview on page 399

## Understanding 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)

| **Before You Begin** |
| --- |
| For background information, read "Understanding IPsec Operational Modes" on page 380. |

This topic covers:

- Authentication Header (AH) Protocol on page 382
- Encapsulating Security Payload (ESP) Protocol on page 383
- Related Topics on page 383

### *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 thourgh 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 and SHA-1 hashing algorithms, see the following RFCs: for more information on MD5, see RFCs 1321, 2403; for more information on SHA-, 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.

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 it is possible to select NULL for authentication, it has been demonstrated that IPsec might be vulnerable to attack under such circumstances. Therefore select NULL for authentication.

### *Related Topics*

- Understanding IPsec Security Associations (SAs) on page 384
- Understanding IPsec Key Management on page 385

- Understanding IKE and IPsec Packets on page 386
- Understanding IPsec Tunnel Negotiation on page 391
- Configuring an IKE IPsec Tunnel—Overview on page 399

## Understanding IPsec 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. Full bidirectional communication requires at least two SAs, one for each direction.

**Before You Begin**

For background information, read

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Security Protocols" on page 382.

An SA groups together the following components for securing communications:

- Security algorithms and keys
- Protocol mode, either Transport or Tunnel (see "Understanding IPsec Operational Modes" on page 380)
- Key-management method, either manual key or AutoKey IKE (see "Understanding IPsec Key Management" on page 385)
- SA lifetime

For outbound VPN traffic, the policy invokes the SA associated with the VPN tunnel. For inbound traffic, JUNOS software looks up the SA by using the following triplet:

- Destination IP
- Security protocol, either AH or ESP (see "Understanding IPsec Security Protocols" on page 382)
- Security parameter index (SPI) value

In SRX-series services gateways, the IKE provides tunnel management for IPsec and authenticates end entities . The IKE performs a Diffie-Hellman key exchange to establish an IPsec tunnel between network devices.

### *Related Topics*

- Understanding IKE and IPsec Packets on page 386
- Understanding IPsec Tunnel Negotiation on page 391
- Configuring VPN Global Settings—Quick Configuration on page 397

## Understanding IPsec Key Management

The distribution and management of keys are critical to using VPNs successfully. IPsec supports both manual and automatic key-distribution methods..

**Before You Begin**

For background information, read

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Security Protocols" on page 382.
- "Understanding IPsec Security Associations (SAs)" on page 384.

This topic covers:



### *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.

### Distributed VPN in SRX-series Services Gateway

In an SRX-series services gateway, 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 tables (source IP address, destination IP addresses, and UDP ports). 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 Packets on page 386
- Understanding IPsec Tunnel Negotiation on page 391
- Configuring an IKE IPsec Tunnel—Overview on page 399

## Understanding IKE and IPsec Packets

An IPsec VPN tunnel consists of two major elements:

- Tunnel Setup—The peers first establish security associations (SAs), which define the parameters for securing traffic between themselves. The admins at each end can define the SAs manually, or they can configure the endpoints to define SAs dynamically through IKE Phase 1 and Phase 2 negotiations. Phase 1 can occur in either Main mode or Aggressive mode. Phase 2 always occurs in Quick mode.
- Applied Security—IPsec protects traffic sent between the two tunnel endpoints by using the security parameters defined in the SAs that the peers agree to during the tunnel setup. IPsec can be applied in either Transport mode or Tunnel mode.

Both modes support the Encapsulating Security Payload (ESP) and Authentication Header (AH) protocols.

**Before You Begin**

For background information, read

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Security Protocols" on page 382.
- "Understanding IPsec Security Associations (SAs)" on page 384.
- "Understanding IPsec Key Management" on page 385.

This topic covers:



## *IKE Packets*

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 83 on page 388.

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 83: IKE Packet for Phases 1 and 2**

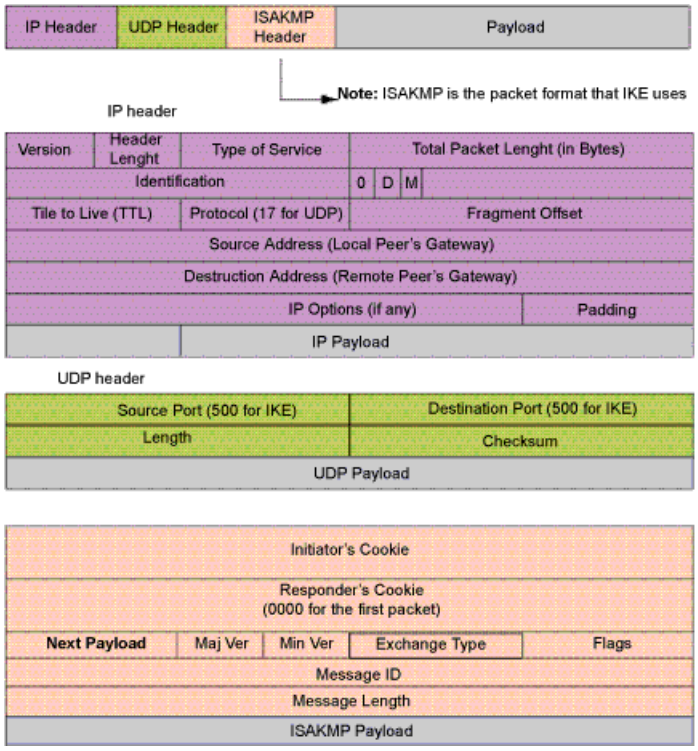


The Next Payload field contains a number indicating one of the following payload types:

- 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 84 on page 389.

**Figure 84: Generic ISAKMP Payload Header**

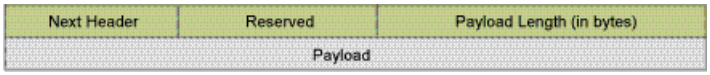


There can be multiple ISAKMP payloads chained together, with each subsequent payload type indicated by the value in the Next Header field. A value of **0000** indicates the last ISAKMP payload. See Figure 85 on page 389 for an example.

**Figure 85: ISAKMP Header with Generic ISAKMP Payloads**

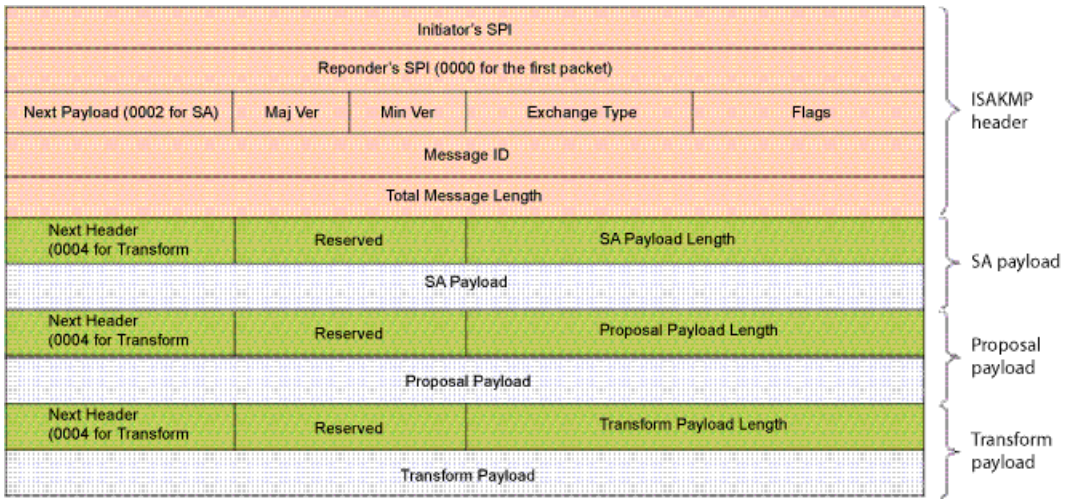


## *IPsec Packets*

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 383. For information about Tunnel mode, see “Tunnel Mode” on page 380.

As shown in Figure 86 on page 390, the packet that the initiating host constructs includes the payload, the TCP header, and the inner IP header (IP1).

**Figure 86: IPsec Packet—ESP in Tunnel Mode**

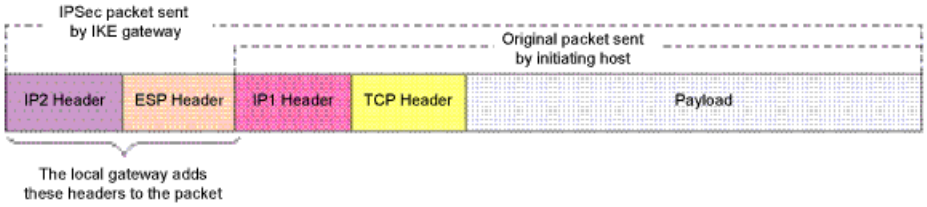


The outer IP header (IP2), which JUNOS software adds, contains the IP address of the remote gateway as the destination IP address and the IP address of the local router as the source IP address. JUNOS software also adds an ESP header between the outer and inner IP headers. The ESP header contains information that allows the remote peer to properly process the packet when it receives it. This is illustrated in Figure 87 on page 390.

**Figure 87: Outer IP Header (IP2) and ESP Header**

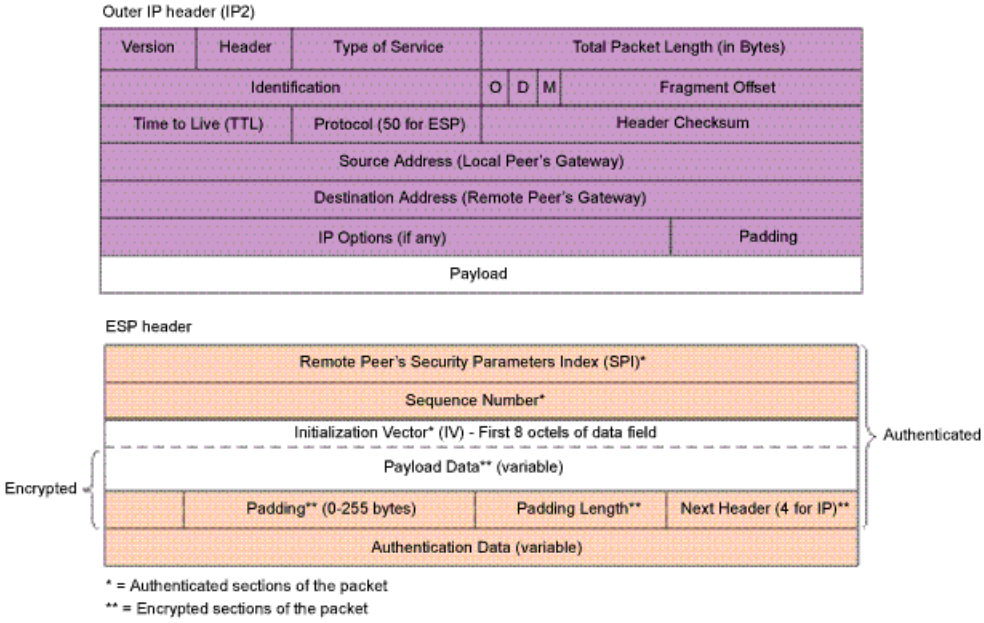


The Next Header field indicates the type of data in the payload field. In Tunnel mode, this value is 4, indicating IP-in-IP. If ESP is applied in Transport mode, this value indicates a Transport Layer protocol such as 6 for TCP or 17 for UDP. See Figure 88 on page 391.

**Figure 88: Inner IP Header (IP1) and TCP Header**

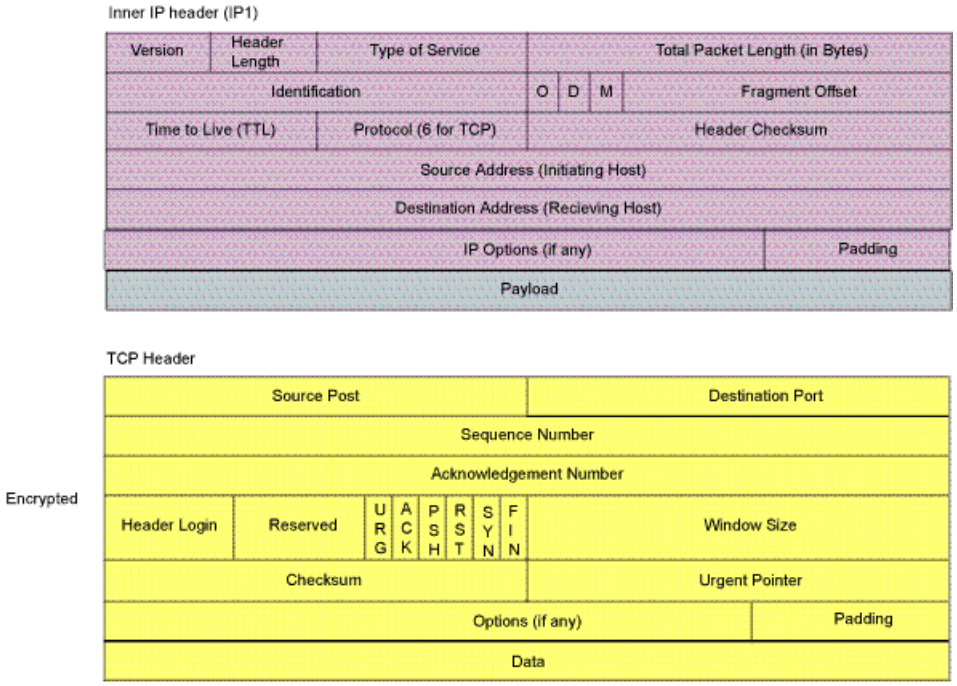


## *Related Topics*

- Understanding IPsec Tunnel Negotiation on page 391
- Configuring an IKE IPsec Tunnel—Overview on page 399

## Understanding IPsec Tunnel Negotiation

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.

In SRX series services gateways, The IKE provides tunnel management for IPsec. 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.

**Before You Begin**

For background information, read

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Security Protocols" on page 382.
- "Understanding IPsec Security Associations (SAs)" on page 384.
- "Understanding IPsec Key Management" on page 385.
- "Understanding IKE and IPsec Packets" on page 386.

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 SAs.
- In Phase 2, the participants negotiate the IPsec SAs for encrypting and authenticating the ensuing exchanges of user data.

This topic covers:



## *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 exchange can be in one of two modes: Aggressive or Main. Using either mode, the participants exchange proposals for acceptable security services such as:

- Encryption algorithms (DES and 3DES) and authentication algorithms (MD5 and SHA-1)—See "Understanding IPsec Security Protocols" on page 382.
- A Diffie-Hellman group—See "Diffie-Hellman Exchange" on page 393.
- Preshared Key or RSA/DSA certificates—see "AutoKey IKE" on page 385.

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.

### Main and Aggressive Modes

Phase 1 can take place in either Main or Aggressive mode.

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—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. Note also that a dialup VPN user can use an email address, a fully qualified domain name (FQDN), or an IP address as its IKE ID. A dynamic peer can use either an email address or FQDN, but not an IP address.

### 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, and 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.

### *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.

J-series 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. JUNOS software also provides a replay protection feature. 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. (For more information about replay protection, see “Replay Protection” on page 395.)

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.

In Phase 2, the peers also 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.

### 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.

### 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. The replay-protection feature 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.

### *Related Topics*

- Configuring an IKE IPsec Tunnel—Overview on page 399

## Configuring VPN Global Settings

Global VPN settings help you monitor and maintain the efficient operation of your VPN.

**Before You Begin**

For background information, read

- “Understanding IPsec Operational Modes” on page 380.
- “Understanding IPsec Security Protocols” on page 382.
- “Understanding IPsec Security Associations (SAs)” on page 384.
- “Understanding IPsec Key Management” on page 385.
- “Understanding IKE and IPsec Packets” on page 386.
- “Understanding IPsec Tunnel Negotiation” on page 391.

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.

To configure VPN global settings, use either the J-Web or the CLI configuration editor.

This topic covers:


- J-Web Configuration on page 396
- CLI Configuration on page 396
- Related Topics on page 397


## J-Web Configuration

You configure the bad SPI response on the IKE J-Web page, and VPN monitoring on the IPsec J-Web page.

1. In the J-Web user interface, select Configuration > View and Edit > Edit Configuration > Security > IKE.
2. Check the Respond Bad SPI check box, then click Configure..
3. Enter a value in the Max responses 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.
5. Select Configuration > View and Edit > Edit Configuration > Security > IPSec.
6. Check the VPN monitor options check box, and then click Configure.
7. Enter a value in the Interval field.
8. Enter a value in the Threshold field.
9. 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 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.

```
user@host# set security ike respond-bad-spi 5
user@host# set security ipsec vpn-monitor-options interval 15 threshold 15
```

### *Related Topics*

- Configuring VPN Global Settings—Quick Configuration on page 397
- Configuring an IKE IPsec Tunnel—Overview on page 399

## Configuring VPN Global Settings—Quick Configuration

You can use J-Web Quick Configuration to quickly configure VPN global settings.

**Before You Begin**

For background information, read

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Tunnel Negotiation" on page 391.
- "Understanding IPsec Security Protocols" on page 382.
- "Understanding IPsec Security Associations (SAs)" on page 384.
- "Understanding IPsec Key Management" on page 385.
- "Understanding IKE and IPsec Packets" on page 386.

Figure 89 on page 398 shows the VPN quick configuration page.

**Figure 89: VPN Global Settings**


To configure VPN global settings with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > Global Settings.
2. Fill in the options as shown in Table 64 on page 398, 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 64: 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. |

**Table 64: VPN Global Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| 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** | | |
| 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 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

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Tunnel Negotiation" on page 391.
- "Understanding IPsec Security Protocols" on page 382.
- "Understanding IPsec Security Associations (SAs)" on page 384.
- "Understanding IPsec Key Management" on page 385.
- "Understanding IKE and IPsec Packets" on page 386.

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 either the J-Web, or the J-Web Quick Configuration. 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
   b. Configure IKE policies (and reference the proposals)
   c. Configure IKE gateway (and reference the policy)

2. Phase 2
    1. Configure Phase 2 proposals
    2. Configure policies (and reference proposals)
    3. Configure IPSec Autokey IKE (and reference the policy and gateway)

### *Related Topics*

- Configuring an IKE Phase 1 Proposal on page 400
- Configuring an IKE Phase 1 Proposal—Quick Configuration on page 401

## Configuring an IKE Phase 1 Proposal

In Phase 1 proposal configuration, you must set the authentiation method and authentication and encryption algorithms that will be used to open a secure channel between participants.

| **Before You Begin** |
| --- |
| For background information, read<br>■ "Understanding IPsec Tunnel Negotiation" on page 391.<br>■ "Configuring an IKE IPsec Tunnel—Overview" on page 399. |

When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway.

To configure Phase 1 proposals, use the J-Web or CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure IKE:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ike, click Configure.
4. Next to Proposal, click Add new entry.
5. In the Name box, type **Ike_prop_1**.

6. In the Description box, type **new Ike proposal**.
7. From the Authentication method list, select pre-shared-keys.
8. From the Dh-group list, select group2.
9. From the Authentication algorithm list, select md5.
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*

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 SA.

```
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
```

### *Related Topics*

- Configuring an IKE Policy, Authentication, and Proposal on page 405
- Configuring an IKE Phase 1 Proposal—Quick Configuration on page 401
- Configuring an IPsec Phase 2 Proposal on page 416

## Configuring an IKE Phase 1 Proposal—Quick Configuration

You can use J-Web Quick Configuration to quickly configure an IKE Phase 1 proposal.

| **Before You Begin** |
| --- |
| For background information, read<br>■ "Understanding IPsec Tunnel Negotiation" on page 391.<br>■ "Configuring an IKE IPsec Tunnel—Overview" on page 399. |

Figure 90 on page 402 shows the quick configuration page where can you select an existing Phase 1 proposal, or click Add to create a new one.

**Figure 90: IKE Phase 1 Proposal Configuration**

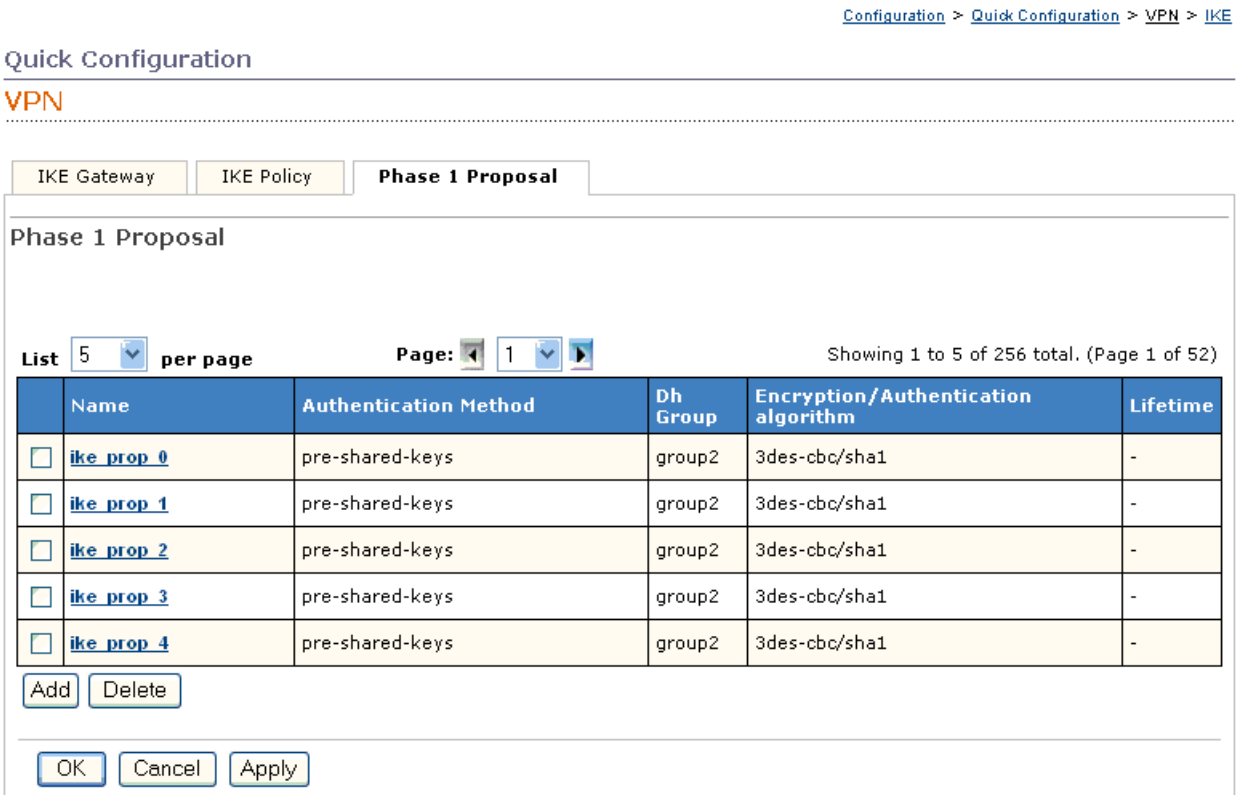


Figure 91 on page 403 shows the IKE Phase 1 Proposal Configuration Options page.

**Figure 91: IKE Phase 1 Proposal Configuration Options**

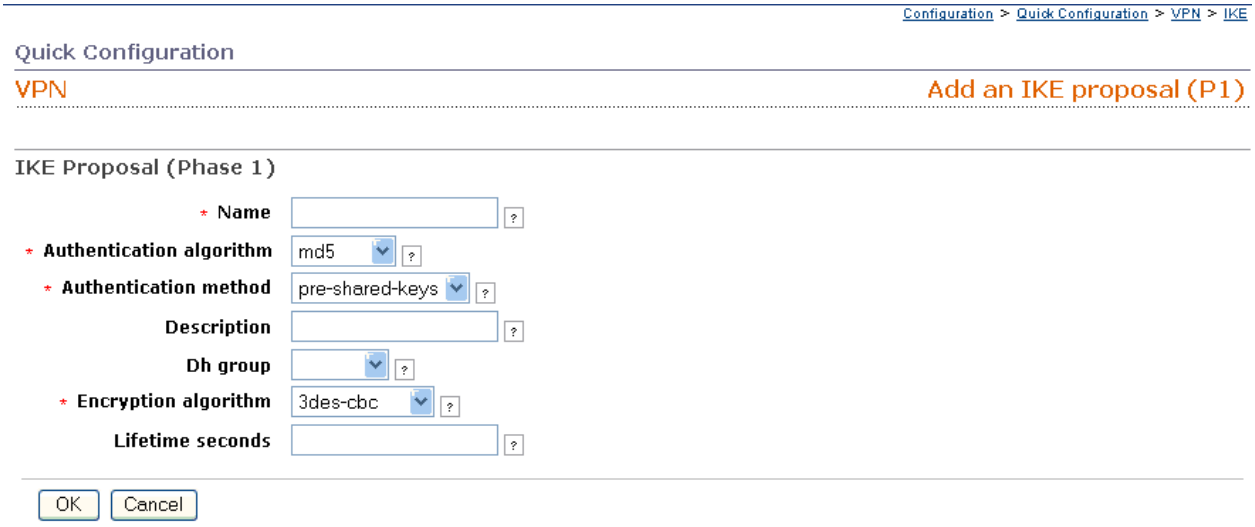


To configure an IKE proposal with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > IKE.
2. Select the Phase 1 Proposal tab if it is not selected.
3. To use an existing proposal, select it from among those listed 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.
4. To configure a new IKE policy, click Add.

   Figure 91 on page 403 shows the quick configuration page where you create a new IKE proposal.
5. Fill in the options as described in Table 65 on page 403.
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: Phase 1 Proposal Configuration Options**

| Field | Function | Action |
|---|---|---|
| **IKE Proposal (Phase 1)** | | |

**Table 65: Phase 1 Proposal Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Name | The name of the proposal. | Enter a name. |
| Authentication algorithm | The Authentication Header (AH) algorithm the device uses to verify the authenticity and integrity of a packet. Supported algorithms include the following:<br>■ md5—Produces a 128-bit digest.<br>■ sha1—Produces a 160-bit digest.<br>■ sha-256—Produces a 256-bit digest. | Select an algorithm. |
| Authentication method | The method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. Options include:<br>■ pre-shared-keys—Key for encryption and decryption that both participants must have before beginning tunnel negotiations.<br>■ 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 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 altorithm | Supported Internet Key Exchange (IKE) proposals include the following:<br>■ 3des-cbc—3DES-CBC encryption algorithm.<br>■ aes-128-cbc—AES-CBC 128-bit encryption algorithm.<br>■ aes-192-cbc—AES-CBC 192-bit encryption algorithm.<br>■ aes-256-cbc—AES-CBC 256-bit encryption algorithm.<br>■ des-cbc—DES-CBC encryption algorithm. | Select an encryption algorithm. |
| 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.<br>Default: 3,600 seconds.<br>Range: 180 through 86,400 seconds. |

## Configuring an IKE Policy, Authentication, and Proposal

In Phase 1 IKE 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.

**Before You Begin**

For background information, read

- "Configuring an IKE Phase 1 Proposal" on page 400.
- "Understanding IPsec Key Management" on page 385.

When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway.

To configure IKE policies, use the J-Web or CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure security Ike policy:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ike, click Configure.
4. Next to Policy, click Add new entry.
5. In the Name box, type **Ike_pol_1**.
6. From the Mode box, select main.
7. In the Description box, type **new Ike policy**.
8. Next to Proposals, click Add new entry.
9. In the Value keyword box, type **Ike_prop_1** and click OK.
10. Next to Pre shared key, click Configure.
11. From the key choice list, select Ascii text.
12. In the Ascii textbox, type **$9$UQiqf36A1RSTzRSreXxDik.Tzn/CuBI** and click OK.
13. To save and commit the configuration, click Commit.

### CLI Configuration

In this example, you create a policy called ike_pol_1, specify that participants exchange proposals in Main mode, and reference the proposal called ike_prop_1. You specifiy that the preshared key be of type ASCII, and enter the key.

```
user@host# set security ike policy ike_pol_1 mode main
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
```

### Related Topics

- Configuring an IKE Gateway and Peer Authentication on page 410
- Configuring an IKE IPsec Tunnel—Overview on page 399
- Configuring an IKE Policy, Authentication, and Proposal—Quick Configuration on page 406
- Configuring an IKE Phase 1 Proposal on page 400
- Configuring an IPsec Policy on page 420

## Configuring an IKE Policy, Authentication, and Proposal—Quick Configuration

You can use J-Web Quick Configuration to quickly configure an IKE policy, authentication, and to reference a proposal.

**Before You Begin**

For background information, read “Configuring an IKE Phase 1 Proposal—Quick Configuration” on page 401.

Figure 92 on page 407 shows the quick configuration page where you can select an existing policy, or click Add to create a new one.

**Figure 92: IKE Policy Configuration**

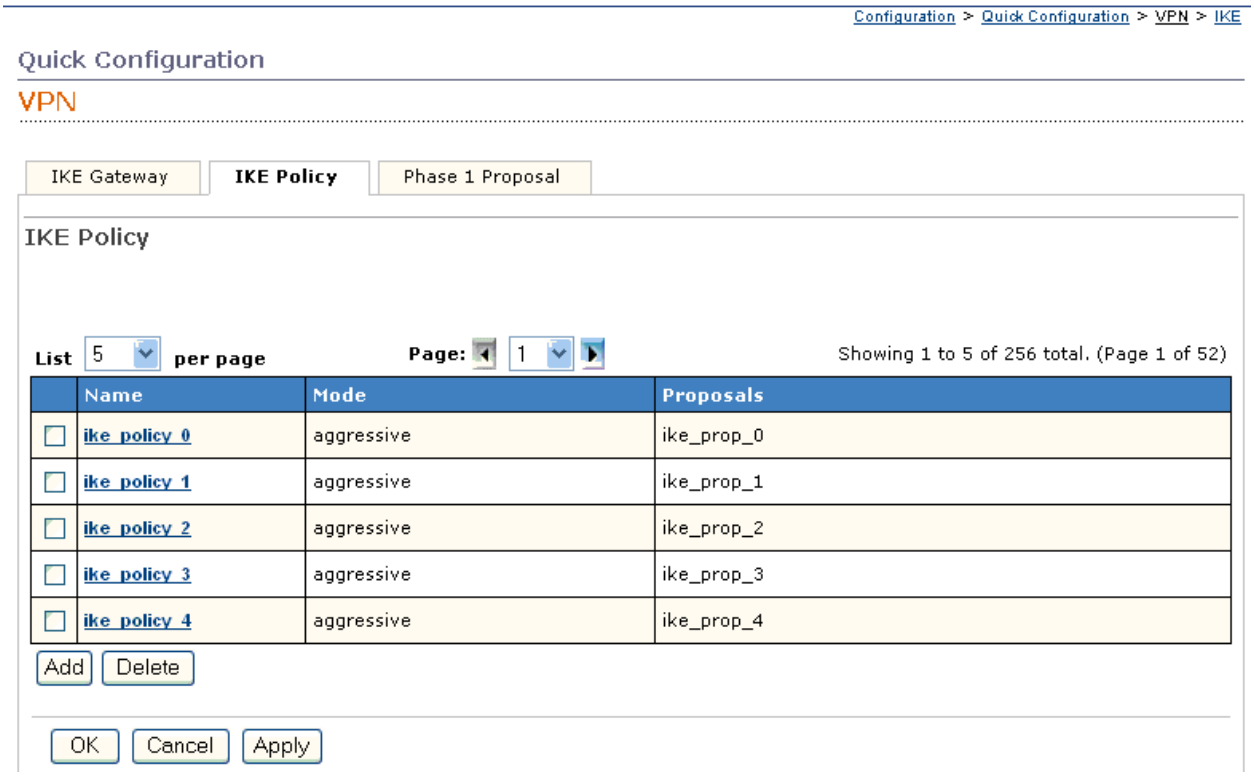


Figure 93 on page 408 shows the IKE Policy, Authentication, and Proposal Configuration page.

**Figure 93: IKE Policy, Authentication, and Proposal Configuration**

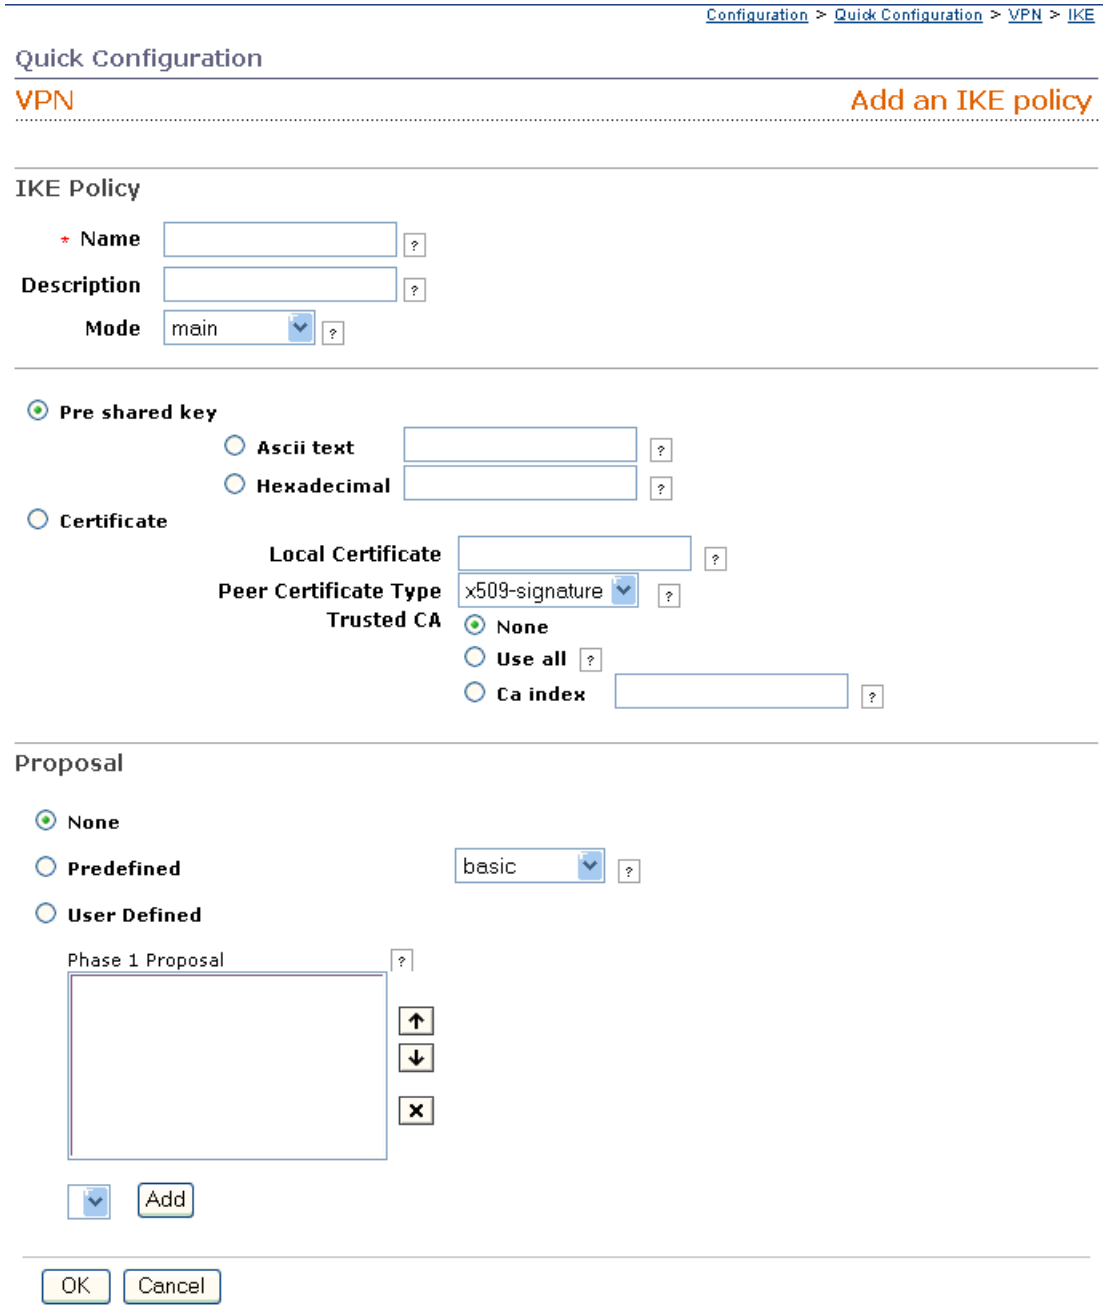


To configure an IKE policy with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > IKE.
2. Select IKE Policy tab if it is not selected.

3. To use an existing policy, select it from among those listed 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.
4. To configure a new IKE policy, authentication, and to reference a proposal, click Add.

   Figure 93 on page 408 shows the quick configuration page where you create a new IKE policy.
5. Fill in the options as described in Table 66 on page 409.
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: IKE Policy, Authentication, and Proposal Options**

| Field | Function | Action |
|---|---|---|
| **IKE Policy** | | |
| Name | Name of the policy. | Enter a name. |
| Descripton | Descripton of the policy. | Enter a description. |
| Mode | Use Main or Aggressive mode. | Select a mode. |
| Pre shared key | Use one of the following preshared key types:<br>■ ASCII text<br>■ Hexadecimal | Click Pre shared key, click the type of key, and enter the key in the appropriate format. |
| Certificate | Use certificates | Click Certificate. |
| 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 | Use a trusted certificate authority. | Click None or Use all, or click Ca index and select from the list. |
| **Proposal** | | |
| None | Do not use proposals | Click None. |

**Table 66: IKE Policy, Authentication, and Proposal Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Predefined | Use one of the following types of predefined Phase 1 proposals:<br>■ Basic<br>■ Compatible<br>■ 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. |

## Configuring an IKE Gateway and Peer Authentication

In Phase 1 gateway configuration, you must configure a gateway and reference the Phase 1 policy.

For background information, read "Configuring an IKE Policy, Authentication, and Proposal" on page 405.

When configuring Phase 1 of an IPsec tunnel using IKE, you first configure proposals, then policies, and finally you configure the gateway.

To configure a gateway for an IKE tunnel, use either the J-Web or the CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure gateway for IKE gateway:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ike, click Configure.
4. Next to Gateway, click Add new entry.
5. In the Gateway name box, type **Ike_gateway_1**.
6. In the Ike policy box, type **ike_pol_1**.
7. In the External interface box, **type ge-0/0/0**.

8. From the Remote identifier list, select Address.
9. Next to Address, click Add new entry.
10. In the Value box, type **1.1.1.2**.
11. Next to Dead peer detection, select the check box and click Configure.
12. In the Interval box, type **10**.
13. In the Threshold box, type **5** and click OK.
14. To save and commit the configuration, click Commit.

### *CLI Configuration*

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.

```
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
```

### *Related Topics*

- Configuring an IKE Gateway and Peer Authentication—Quick Configuration on page 411
- Configuring an IPsec Phase 2 Proposal on page 416
- Configuring an IKE IPsec Tunnel—Overview on page 399
- Understanding IPsec Tunnel Negotiation on page 391

## Configuring an IKE Gateway and Peer Authentication—Quick Configuration

You can use J-Web Quick Configuration to quickly configure an IKE Gateway.

| Before You Begin |
| --- |
| For background information, read<br>■ "Configuring an IKE Policy, Authentication, and Proposal—Quick Configuration" on page 406. |

Figure 94 on page 412 shows the quick configuration page where you can select an existing gateway, or click Add to create a new one.

**Figure 94: IKE Gateway Configuration**

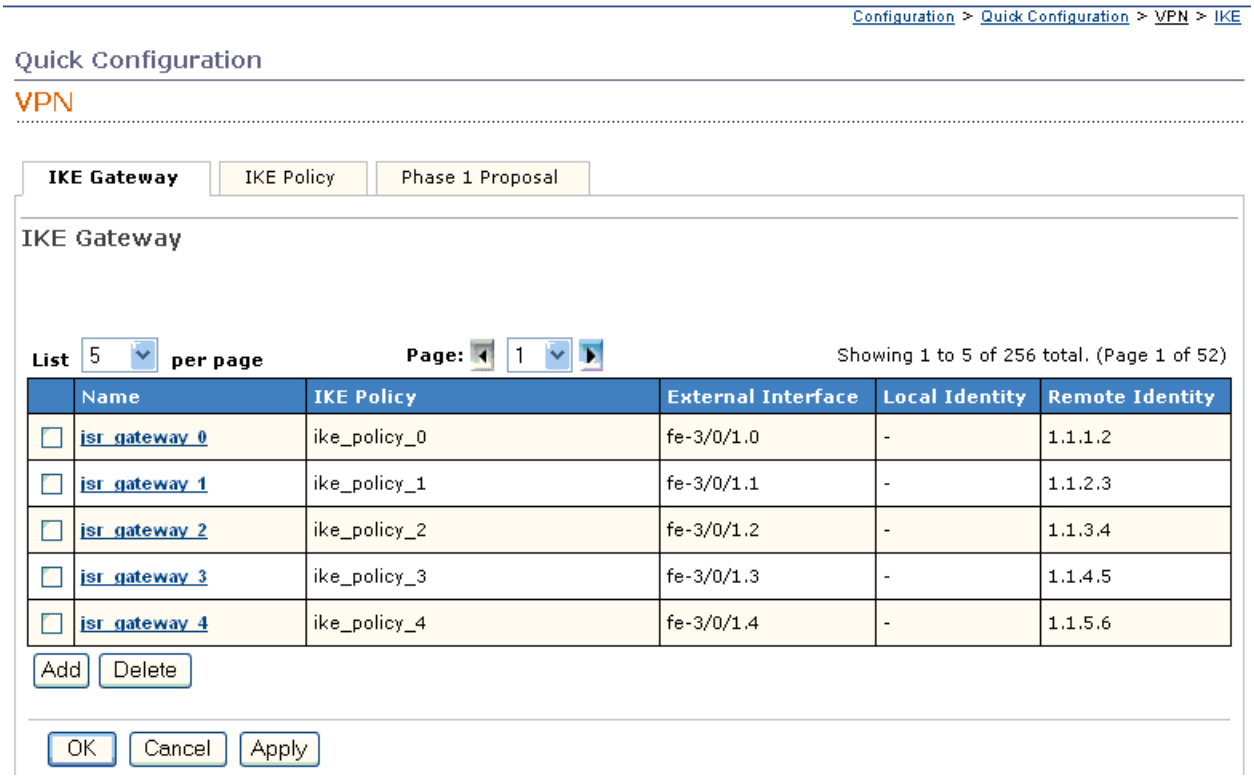


Figure 95 on page 413 shows the quick configuration page where you create a new IKE gateway.

**Figure 95: IKE Gateway and Peer Authentication Options**

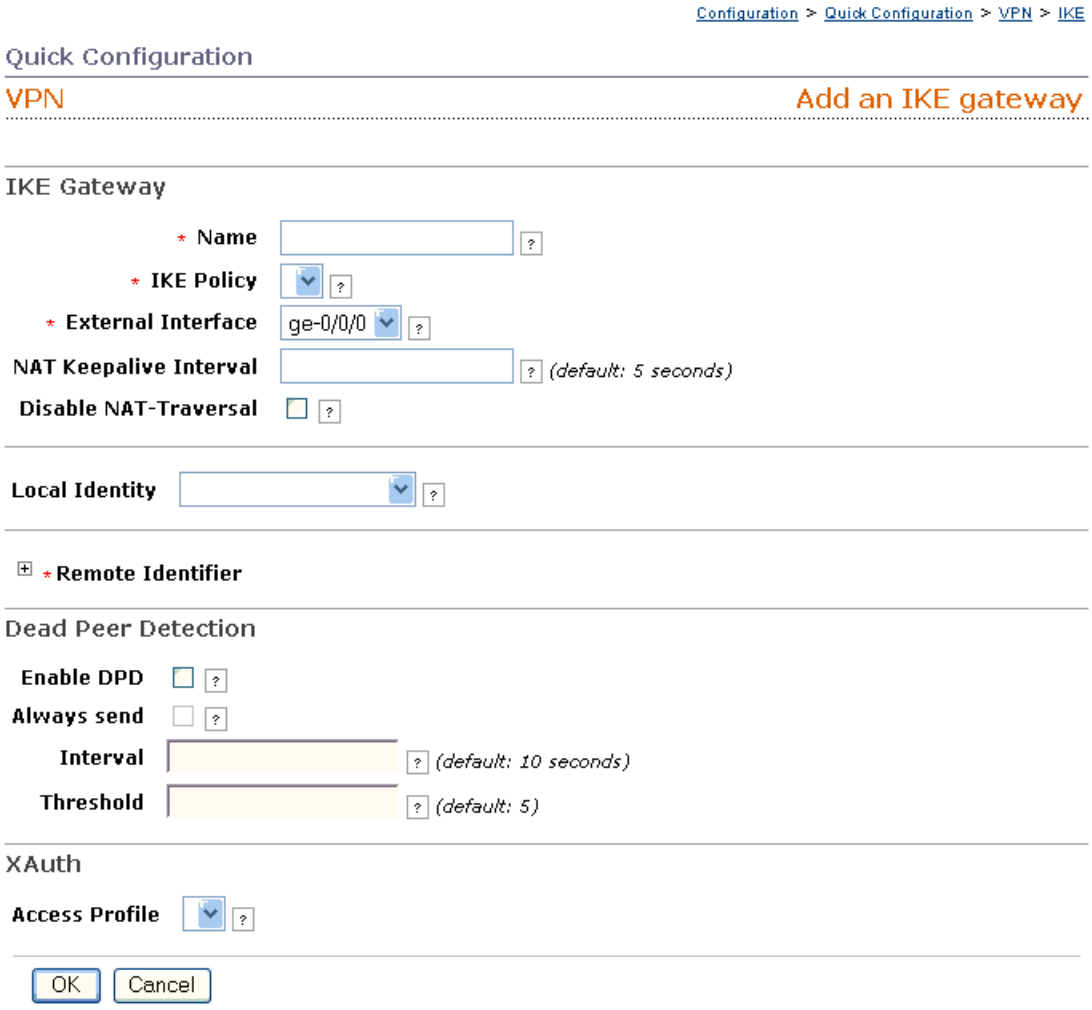


To configure an IKE gateway with Quick Configuration:

1. In the J-Web user interface, select Configuration > Quick Configuration > VPN > IKE.
2. Select the IKE Gateway tab if it is not selected.
3. To use an existing gateway, select one from among those listed 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
4. To configure a new IKE gateway, click Add.

5. Fill in the options as described in Table 67 on page 414.
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: IKE Gateway and Peer Authentication Options**

| Field | Function | Action |
|---|---|---|
| **IKE Gateway** | | |
| Name | Name of the destination peer gateway, specified as an alphanumeric string | Enter a name. |
| IKE Policy | The IKE policy to be used for communication with the destination peer gateway. | Select the IKE policy to use for the peer gateway |
| 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. |
| NAT Keepalive Interval | The interval at which 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. |
| Disable NAT Traversal | Disables UDP encapsulation of IPsec Encapsulating Security Payload (ESP) packets, otherwise known as Network Address Translation Traversal (NAT-T). NAT-T is enabled by default. | Click the check box to disable or enable. |
| 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 either of the following ways:<br>■ IP Address—IPv4 IP address to identify the dynamic peer.<br>■ Hostname—Fully qualified domain name (FQDN) to identify the dynamic peer.<br>■ User at Hostname—E-mail address to identify the dynamic peer.<br>■ 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. |
| Remote Identifier | Provides information about remote peer. | Click the check box to expand the field. |
| Address | Static address or hostname of remote peer. | Specify one primary and up to four backups. |

**Table 67: IKE Gateway and Peer Authentication Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Dynamic | Dynamic address of remote peer. | Select. |
| Connections limit | Maximum number of concurrent connections allowed. | Specify the maximum number of concurrent users that can be connected to the gateway. When the maximum number of connections is reached, no more dynamic VPN endpoint dial-up users attempting to access an IPsec VPN are allowed to begin Internet Key Exchange (IKE) negotiations. |
| IKE user type | ■ group-ike-id— E-mail address or fully qualified domain name (FQDN) shared for a group of remote access users so that each one does not need a separate IKE profile configured.<br>■ shared-ike-id—Email address shared for a large number of remote access users so that each one does not need a separate IKE profile configured. | Select the type of IKE user for a remote access connection. |
| Remote id type | Type of identifier for remote peer. | Select IP Address, Hostname, User at hostname, or Distinguished name. |
| IP Address | Use an IPv4 IP address to identify the dynamic peer. | Enter an IP address. |
| Hostname | Use a fully qualified domain name (FQDN) to identify the dynamic peer. | Select and enter the FQDN. |
| User at hostname | Use an e-mail address to to identify the dynamic peer. | Select and enter the remote identity as an e-mail address. |
| Distinguished name | Use a distinguished 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. | Select and specify a container or a wildcard. |
| container | The distinguished name of the remote peer. | Enter the distinguished name string exactly as it appears in the certificate. |
| wildcard | All or part of the distinguished name of the remote peer. | Enter all or parts of a distinguished name, in any order. |
| **Dead Peer Detection** | | |
| Enabled DPD | *RFC 3706* Dead Peer Detection. | Click the check box. |
| 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. |

**Table 67: IKE Gateway and Peer Authentication Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Interval | The amount of time that the peer waits for traffic from its destination peer before sending a dead peer detection (DPD) request packet. | Enter the interval at which to send DPD messages. Range: 1 through 60 seconds. |
| Threshold | The maximum number of unsuccessful dead peer detection (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 | Provides extended authentication (XAuth), in addition to IKE authentication for remote users trying to access a VPN tunnel. | Select a previously created access profile to reference for authentication information. |

## Configuring an IPsec Phase 2 Proposal

In Phase 2 proposal configuration, you must create a proposal, specify a security protocol, and select althentication and encryption algorithms for the traffic that will flow through the tunnel.

**Before You Begin**

For background information, read

- "Understanding IPsec Tunnel Negotiation" on page 391.
- "Configuring an IKE IPsec Tunnel—Overview" on page 399.

When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE).

To configure Phase 2 proposals, use either the J-Web or the CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure an IPSec proposal:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ipsec, click Configure.

4. Next to Proposal, click Add new entry.
5. In the Name box, type **Ipsec_prop_1**.
6. In the Description box, type **new ipsec proposal**.
7. From the Protocol list, select esp.
8. From the Authentication algorithm, select hmac-md5-96.
9. From the Encryption algorithm, select 3des-cbc.
10. In the Lifetime seconds box, type **1800** and click OK.
11. To save and commit the configuration, click Commit.

### *CLI Configuration*

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 algorithy and 3des-cbc as the encryption algorithm. You also specify that the security association (SA) terminate after 1,800 KBof data pass through it.

```
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
```

### *Related Topics*

- Configuring an IPsec Policy on page 420
- Configuring an IPsec Phase 2 Proposal—Quick Configuration on page 417
- Configuring an IKE Phase 1 Proposal on page 400

## Configuring an IPsec Phase 2 Proposal—Quick Configuration

You can use J-Web Quick Configuration to quickly configure IPsec phase 2 proposals.

| **Before You Begin** |
| --- |
| For background information, read “Configuring an IKE IPsec Tunnel—Overview” on page 399. |

Figure 96 on page 418 shows the quick configuration page where you can select an existing proposal, or click Add to create a new one.

**Figure 96: IPsec Phase 2 Proposal Configuration**

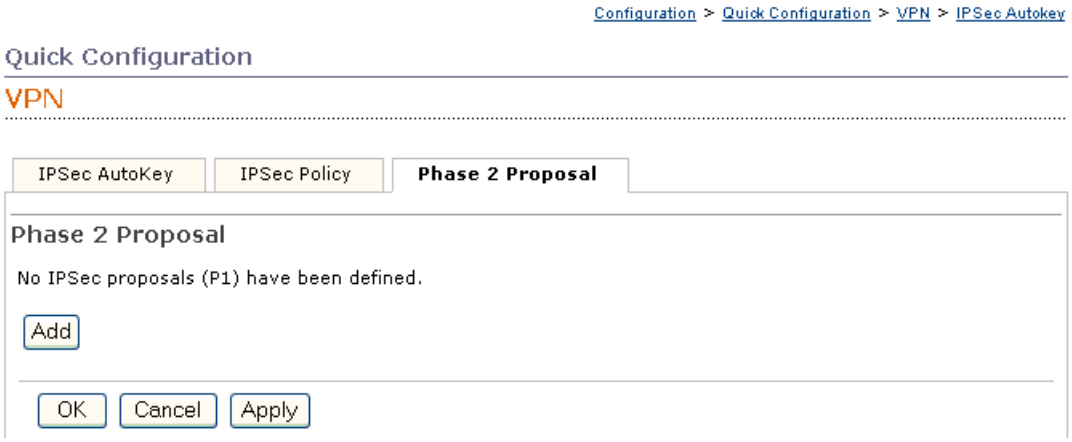


Figure 97 on page 418 shows the quick configuration page where you create a new proposal.

**Figure 97: IPsec Phase 2 Proposal Configuration**

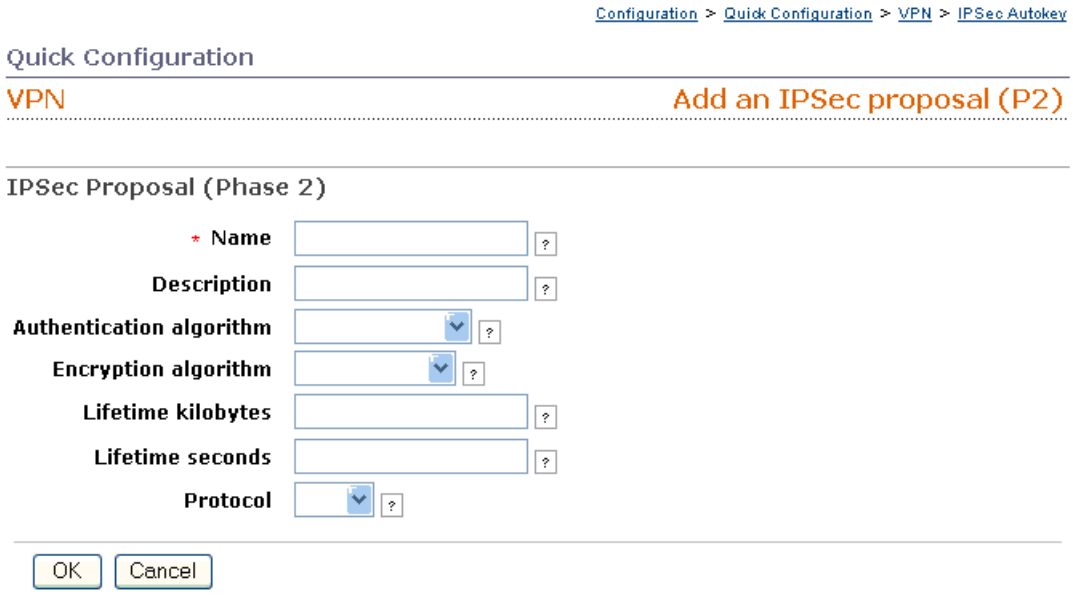


To configure an IPsec Phase 2 proposal with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > IPSec AutoKey.
2. Select the IPsec Phase 2 Proposal tab if it is not selected
3. To use an existing proposal, select one from from among those listed 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.
4. To configure a new IPsec phase 2 proposal, click Add.
5. Fill in the options as described in Table 68 on page 419.
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 68: IPsec Phase 2 Proposal Options**

| Field | Function | Action |
|---|---|---|
| **IPSec Proposal (Phase 2)** | | |
| Name | Description of the Phase 2 proposal. | Enter a name. |
| Description | Identify the proposal | Enter a text description. |
| Authentication algorithm | Hash algorithm that authenticates packet data. It can be one of the following:<br>■ hmac-md5-96—Produces a 128-bit digest.<br>■ hmac-sha1-96—Produces a 160-bit digest. | Select a hash algorithm. |
| Encryption algorithm | Configures an IKE encryption algorithm.<br>■ 3des-cbc—Has a block size of 24 bytes; the key size is 192 bits long.<br>■ des-cbc—Has a block size of 8 bytes; the key size is 48 bits long.<br>■ aes-128-cbc—AES 128-bit encryption algorithm.<br>■ aes-192-cbc—AES 192-bit encryption algorithm.<br>■ 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 have passed. | Enter a value from 64 through 1,048,576 bytes. |

**Table 68: IPsec Phase 2 Proposal 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. | Enter a value from 180 through 86,400 seconds. |
| Protocol | The type of security protocol. | Select a protocol for the proposal. |

## Configuring an IPsec Policy

In Phase 2 IPsec policy configuration, you must create a policy and reference a Phase 2 proposal.

| Before You Begin |
|---|
| For background information, read "Configuring an IPsec Phase 2 Proposal" on page 416. |

When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE).

To configure Phase policies, use either the J-Web or the CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure an IPSec policy:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ipsec, click Configure.
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 name box, type **Ipsec_prop_1** and click OK.
11. To save and commit the configuration, click Commit.

### *CLI Configuration*

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.

```
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
```

### *Related Topics*

- Configuring an IPsec Policy—Quick Configuration on page 421
- Configuring an IKE Policy, Authentication, and Proposal on page 405
- Configuring an IKE IPsec Tunnel—Overview on page 399

## Configuring an IPsec Policy—Quick Configuration

You can use J-Web Quick Configuration to quickly configure an IPsec policy.

| Before You Begin |
| --- |
| For background information, read "Configuring an IKE Phase 1 Proposal—Quick Configuration" on page 401. |

Figure 98 on page 422 shows the quick configuration page where you can select an existing policy, or click Add to create a new one.

**Figure 98: IPsec Policy Configuration**

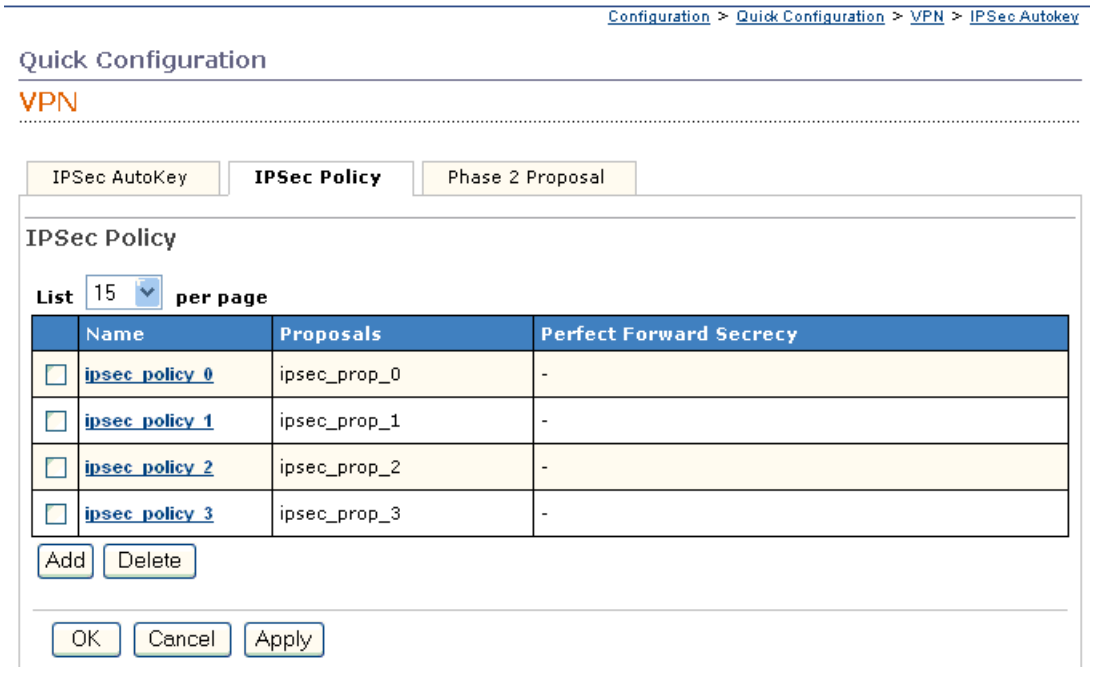


Figure 99 on page 423 shows the quick configuration page where you create a new IPsec policy.

**Figure 99: IPsec Policy Configuration**

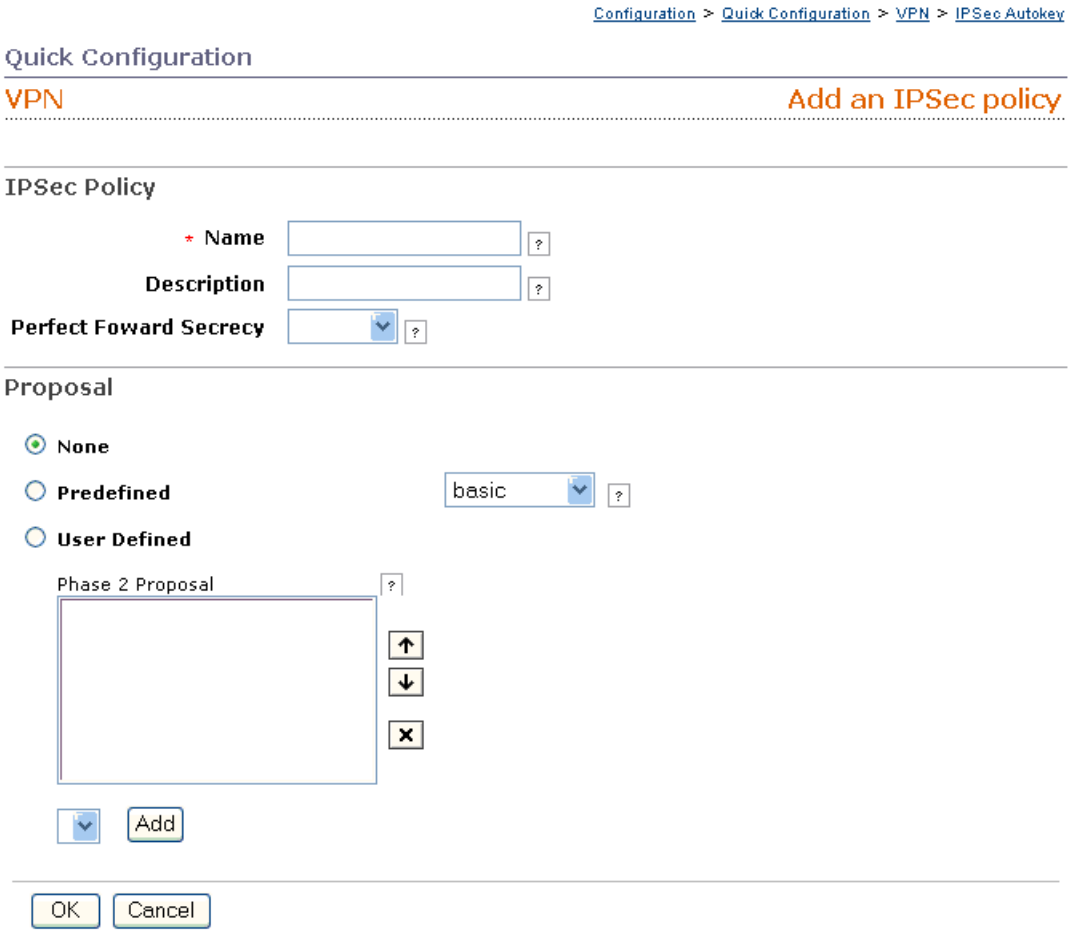


To configure an IPsec policy with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > IPsec Policy.
2. Select the IPsec Policy tab if it is not selected.
3. To use an existing policy, select one from among those listed 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.
4. To configure a new IPsec policy, click Add.
5. Fill in the options as described in Table 69 on page 424.
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 69: IPsec Policy Configuration Options**

| Field | Function | Action |
|---|---|---|
| **IPsec Policy** | | |
| Name | Name of the IPsec policy. | Enter a name. |
| Description | Description of the policy. | Enter a text description. |
| Perfect Forward Secrecy | The method the device uses to generate the encryption key. PFS generates each new encryption key independently from the previous key.<br>■ group1—Diffie-Hellman Group 1.<br>■ group2—Diffie-Hellman Group 2.<br>■ group5—Diffie-Hellman Group 5. | Select a method. |
| **Proposal** | | |
| None | Do not use a proposal. | Click None. |

**Table 69: IPsec Policy Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Predefined | A set of default Internet Key Exchange (IKE) proposals.<br>■ basic—Basic set of two IKE proposals:<br>  ■ Proposal 1—Preshared key, Data Encryption Standard (DES) encryption, and Diffie-Hellman Group 1 and Secure Hash Algorithm 1 (SHA-1) authentication<br>  ■ Proposal 2—Preshared key, DES encryption, and Diffie-Hellman Group 1 and MD5 authentication<br>■ compatible—Set of four commonly used IKE proposals:<br>  ■ Proposal 1—Preshared key, triple DES (3DES) encryption, and G2 and SHA-1 authentication<br>  ■ Proposal 2—Preshared key, 3DES, and Diffie-Hellman Group 2 and MD5 authentication<br>  ■ Proposal 3—Preshared key, DES encryption, and Diffie-Hellman Group 2 and SHA-1 authentication<br>  ■ Proposal 4—Preshared key, DES encryption, and Diffie-Hellman Group 2 and MD5 authentication<br>■ standard—Standard set of two IKE proposals:<br>  ■ Proposal 1— Preshared key, 3DES encryption, and Diffie-Hellman Group 2 and SHA-1 authentication<br>  ■ Proposal 2—Preshared key, Advanced Encryption Standard (AES) 128-bit encryption, and Diffie-Hellman Group 2 and SHA-1 authentication | Click Predefined, and select one of the following options:<br>■ basic<br>■ predefined<br>■ standard |
| User Defined | A list of proposals you previously defined. | Click User Defined and select proposals from the pop-up menu, then click Add. |

## Configuring IPsec AutoKey

In Phase 2 IPsec AutoKey (IKE) 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.

| Before You Begin |
|---|
| For background information, read "Configuring an IPsec Policy—Quick Configuration" on page 421. |

When configuring Phase 2 of an IPsec tunnel, you first configure proposals, then policies, and finally you configure IPsec AutoKey (IKE).

To configure AutoKey IKE, use either the J-Web or the CLI configuration editor.

This topic covers:



## J-Web Configuration

To configure an IPSec Autokey:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ipsec, click Configure.
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

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.

```
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
```

## Related Topics

- Configuring an IKE IPsec Tunnel—Overview on page 399
- Configuring VPN Global Settings on page 395

## Configuring IPsec Autokey—Quick Configuration

You can use J-Web Quick Configuration to quickly configure IPsec AutoKey.

**Before You Begin**

For background information, read "Configuring an IPsec Policy—Quick Configuration" on page 421.

Figure 100 on page 427 shows the quick configuration page, where you can select an existing AutoKey VPN, or click Add to create a new one.

**Figure 100: IPsec AutoKey Configuration**

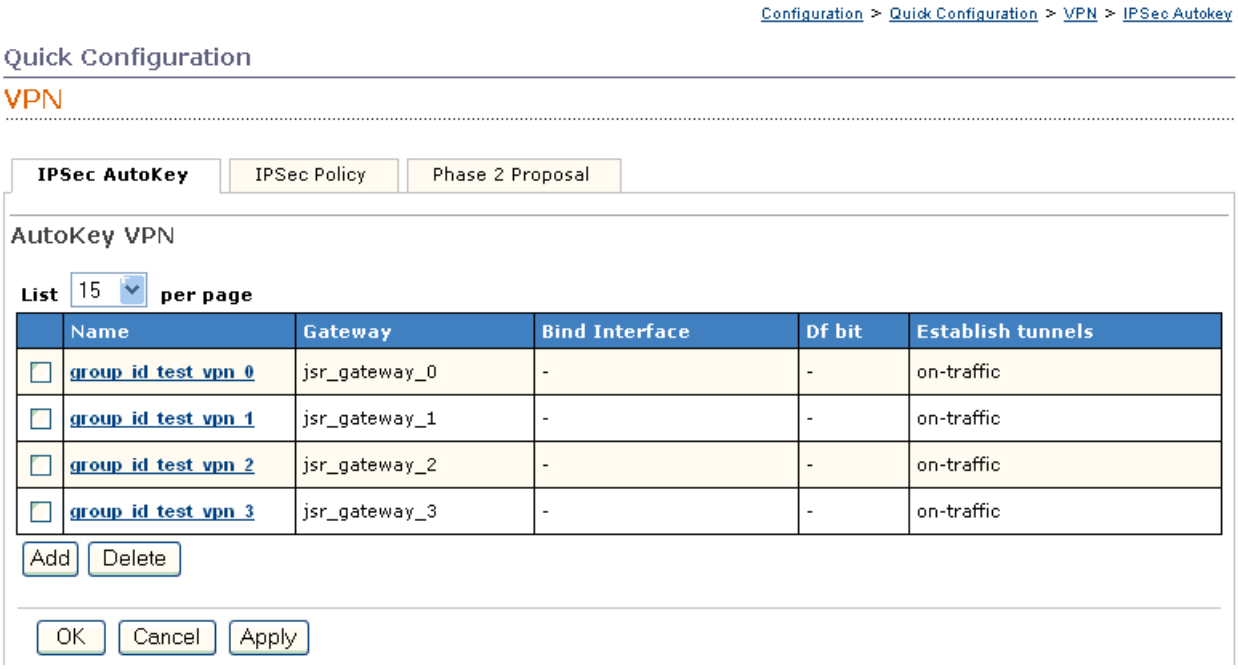


Figure 101 on page 428 shows the quick configuration page where you create a new IPsec AutoKey VPN.

**Figure 101: IPSec AutoKey Configuration Options**

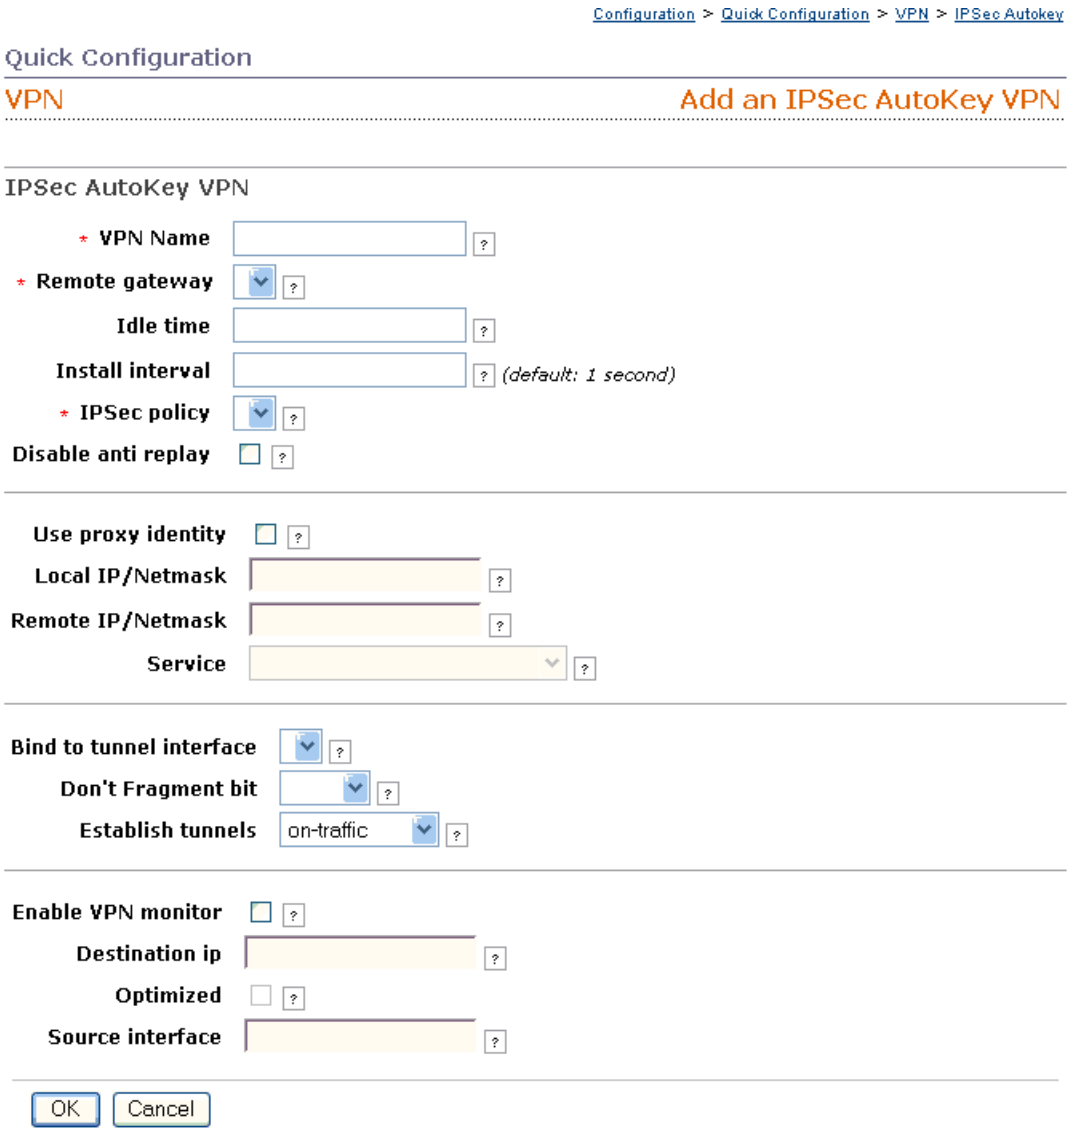


To configure an AutoKey VPN with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > IPSec AutoKey.
2. Select the IPSec AutoKey tab if it is not selected.
3. To use an existing IPsec AutoKey VPN, select one from among those listed 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

4. To configure a new ISec AutoKey VPN, click Add.
5. Fill in the options as described in Table 70 on page 429.
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 70: IPsec AutoKey Configuration Options**

| Field | Function | Action |
|---|---|---|
| **IPsec Autokey VPN** | | |
| VPN Name | Name of the IPsec tunne.l | Enter a name. |
| Remote gateway | Name of the remote gateway. | Select a name. |
| Idle time | The maximum amount of time to allow a security association (SA) to be idle before deleting it. | Specify a value between 60 and 999,999 seconds. |
| Install interval | The maximum number of seconds to allow the installation of a rekeyed outbound security association (SA) on the device. | Specify a value between 0 and 10 seconds. |
| IPsec policy | Associate a policy with this IPsec tunnel. | Select a policy. |
| Disable anti replay | Disable the anti-replay checking feature of IPsec. By default, anti-replay checking is enabled. | Click the check box. |
| Use proxy identity | Optionally, specify the IPsec proxy identify to use in IKE negotiations. The default behavior is to use the identities taken from the firewall policies. | Click the check box. |
| Local IP/Netmask | The local IP address and subnet mask for the proxy identity. | Enter an IP address and subnet mask. |
| Remote IP/Netmask | The remote IP address and subnet mask for the proxy identity. | Enter an IP address and subnet mask. |
| Service | The service (port and protocol combination) to protect. | Select a service. |
| Bind to tunnel interface | The tunnel interface to which the route-based virtual private network (VPN) is bound. | Select an interface. |
| Don't fragment bits | Specifies how the device handles the Don't Fragment (DF) bit in the outer header.<br>■ clear—Clear (disable) the DF bit from the outer header. This is the default.<br>■ copy—Copy the DF bit to the outer header.<br>■ set—Set (enable) the DF bit in the outer header. | Choose an option. |

**Table 70: IPsec AutoKey Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Establish tunnels | Specifies when IKE is activated.<br>■ immediately—IKE is activated immediately after VPN configuration and configuration changes are committed.<br>■ on-traffic—IKE is activated only when data traffic flows and must be negotiated. | Choose an option. |
| Enable VPN monitor | Allows for monitoring of the VPN. | Click the check box. |
| 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. |

## Configuring an IPsec Manual Key VPN

With manual key IPsec, participants agree beforehand on the security association (SA) paramenters, then configure them individually. This, in effect, establishes the tunnel. The devices then encrypt and authenticate any traffic matching a security policy, or when a route involves the tunnel in route-based VPNs, and forward it to the gatway.

**Before You Begin**

For background information, read

- "Understanding IPsec Operational Modes" on page 380.
- "Understanding IPsec Security Protocols" on page 382.
- "Understanding IPsec Security Associations (SAs)" on page 384.
- "Understanding IPsec Key Management" on page 385.
- "Understanding IKE and IPsec Packets" on page 386.
- "Understanding IPsec Tunnel Negotiation" on page 391.

To configure , use either the J-Web or JUNOS CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure an IPSec Manual Key VPN:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Ipsec, click Configure.
4. Next to Vpn, click Add new entry.
5. In the Name box, type **manual_0**.
6. From the Negotiation list, select Manual.
7. Next to Manual, click Configure.
8. In the Gateway box, type **1.1.1.2**.
9. In the External interface box, type **ge-0/0/2.0**.
10. From the Protocol list, select esp.
11. In the Spi box, type **16100**.
12. Next to Authentication, click Configure.
13. From the Algorithm list, select hmac-md5-96.
14. Next to Key, click Configure.
15. From the Key choice list, select Ascii text.
16. In the Ascii text box, type **$9$NrbgJiHmQz6Vwi.5T9CO1REclKvLdb28XVY24DjHqmfz6/Ct** and click OK.
17. Next to Encryption, click Configure.
18. From the Algorithm box, select 3des-cbc.
19. Next to Key, click Configure.
20. From the Key choice list, select Ascii text.
21. In the Ascii text box, type **$9$WoPxVYDjq.5FjHz6/AIRdbs2oGmfz6CujH1hrlXx24aJHmTz69A069K8XNY2n/CuIc8LN-bs**, and click OK.
22. To save and commit the configuration, click Commit.

### CLI Configuration

In this example, you set the gateway address to 1.1.1.2 and the external (outgoing) interface, to ge-0/0/2.0. You specify ESP as the security protocol, set 16,100 as the Security Parameter Index, and se hmac-md5-96 as your authentication algorithm and enter an ASCII key. You then specify 3DES-CBC encription and also enter an ASCII key.

```
user@host# set security ipsec vpn manual_0 manual gateway 1.1.1.2
user@host# set security ipsec vpn manual_0 manual external-interface ge-0/0/2.0
user@host# set security ipsec vpn manual_0 manual protocol esp
user@host# set security ipsec vpn manual_0 manual spi 16100
user@host# set security ipsec vpn manual_0 manual authentication algorithm
  hmac-md5-96
user@host# set security ipsec vpn manual_0 manual authentication key ascii-text
  "$9$NrbgJiHmQz6Vwi.5T9CO1REclKvLdb28XVY24DjHqmfz6/Ct"
user@host# set security ipsec vpn manual_0 manual encryption algorithm 3des-cbc
user@host# set security ipsec vpn manual_0 manual encryption key ascii-text
  "$9$WoPxVYDjq.5FjHz6/AlRdbs2oGmfz6CujH1hrlXx24aJHmTz69A069K8XNY2n/Culc8LN-bs"
```

### Related Topics

- Configuring an IPsec Manual Key VPN—Quick Configuration on page 432
- Configuring an IKE IPsec Tunnel—Overview on page 399
- Configuring VPN Global Settings on page 395
- Configuring VPN Global Settings—Quick Configuration on page 397

## Configuring an IPsec Manual Key VPN—Quick Configuration

You can use J-Web Quick Configuration to quickly configure an IPsec manual key VPN

**Before You Begin**

For background information, read

- “Understanding IPsec Operational Modes” on page 380.
- “Understanding IPsec Security Protocols” on page 382.
- “Understanding IPsec Security Associations (SAs)” on page 384.
- “Understanding IPsec Key Management” on page 385.
- “Understanding IKE and IPsec Packets” on page 386.
- “Understanding IPsec Tunnel Negotiation” on page 391.

Figure 102 on page 433 shows the quick configuration page where you can select an existing manual key VPN, or click Add to create a new one

**Figure 102: IPsec Manual Key VPN**


Figure 103 on page 434 shows the quick configuration page where you create a new IPsec manual key VPN.

**Figure 103: IPsec Manual Key VPN Configuration**


To configure a manual key with Quick Configuration:

1. Select Configuration > Quick Configuration > VPN > IPSec Manual Key.
2. To use an existing manual key VPN, select one from among those listed 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
3. To create a new manual key VPN, click Add.
4. Fill in the VPN options as described in Table 71 on page 435.
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 71: IPsec Manual Key VPN Configuration Options**

| Field | Function | Action |
|---|---|---|
| **IPSec Manual Key VPN** | | |
| VPN Name | Name of the VPN. | Enter a name. |
| Remote gateway | Name of the remote gateway. | Select a name. |
| External Interface | Outgoing interface. | Select an interface. |
| Protocol | Security protocol for this VPN. | Select a protocol. |
| Spi | The security parameter index. An arbitrary value that uniquely identifies which security association (SA) to use at the receiving host (the destination address in the packet). | Enter a value from 256 through 16,639. |
| **Authentication** | | |
| Algorithm | Hash algorithm that authenticates packet data. It can be one of the following:<br>■ hmac-md5-96—Produces a 128-bit digest.<br>■ hmac-sha1-96—Produces a 160-bit digest. | Select an algorithm. |

**Table 71: IPsec Manual Key VPN Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Key | Type of authentication. It can be one of the following:<br>■ None<br>■ ascii-text—ASCII text key. For hmac-md5-96, the key is 16 ASCII characters; for hmac-sha1-96, the key is 20 ASCII characters.<br>■ hexadecimal—Hexadecimal key. For hmac-md5-96, the key is 32 hexadecimal characters; for hmac-sha1-96, the key is 40 hexadecimal characters. | Select none, or select the type of key and enter it in the appropriate format. |
| **Encryption** | | |
| Algorithm | Supported Internet Key Exchange (IKE) proposals include the following:<br>■ 3des-cbc—3DES-CBC encryption algorithm.<br>■ aes-128-cbc—AES-CBC 128-bit encryption algorithm.<br>■ aes-192-cbc—AES-CBC 192-bit encryption algorithm.<br>■ aes-256-cbc—AES-CBC 256-bit encryption algorithm.<br>■ des-cbc—DES-CBC encryption algorithm | Select an encryption algorithm. |
| Key | Type of encryption key. It can be one of the following:<br>■ None<br>■ ascii-text—ASCII text key. For hmac-md5-96, the key is 16 ASCII characters; for hmac-sha1-96, the key is 20 ASCII characters.<br>■ hexadecimal—Hexadecimal key. For hmac-md5-96, the key is 32 hexadecimal characters; for hmac-sha1-96, the key is 40 hexadecimal characters. | Select none, or select the type of key and enter it in the appropriate format. |
| Bind to tunnel interface | The tunnel interface to which the route-based virtual private network (VPN) is bound. | Select an interface. |
| Don't Fragment bit | Specifies how the device handles the Don't Fragment (DF) bit in the outer header.<br>■ clear—Clear (disable) the DF bit from the outer header. This is the default.<br>■ copy—Copy the DF bit to the outer header.<br>■ set—Set (enable) the DF bit in the outer header. | Choose an option. |
| Enable VPN monitor | Allows for monitoring of the VPN. | Click the check box. |

**Table 71: IPsec Manual Key VPN Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| Destination ip | IP address of the destination peer. | Enter an IP address |
| Optimized | Specify 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 monitor messages | Enter a source interface name. |

Chapter 15

# 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.

This section includes:

## 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 440
- Understanding Self-Signed Certificates on page 446
- Understanding Certificate Revocation Lists on page 443
- Understanding Public Key Infrastructure on page 443
- Using Digital Certificates on page 448
- Using Automatically Generated Self-Signed Certificates on page 465
- Manually Generating Self-Signed Certificates on page 466

## 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 446.

| Before You Begin |
| --- |
| For background information, read “Understanding Public Key Cryptography” on page 440. |

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*

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 104 on page 442 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 104 on page 442 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 443). 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 104: 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.

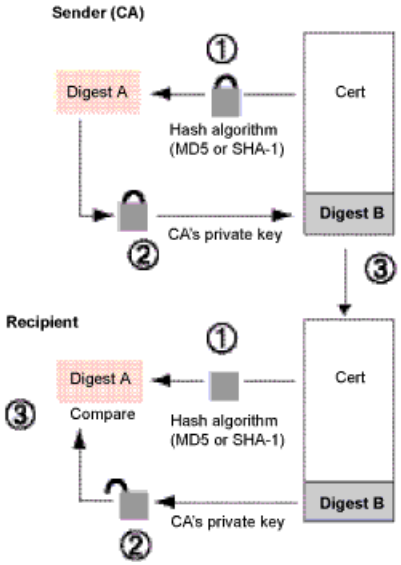


## *Internet Key Exchange*

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 448
- Generating a Public-Private Key Pair on page 450
- Generating a Local Certificate Request Manually on page 456
- Loading CA and Local Certificates Manually on page 458
- Verifying Certificate Validity on page 462
- Understanding Certificate Revocation Lists on page 443

## 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 440
- Understanding Certificates on page 440

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 463
- Deleting a Loaded CRL on page 469
- Understanding Public Key Infrastructure on page 443

## 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 105 on page 444.

**Before You Begin**

For background information, read

- Understanding Public Key Cryptography on page 440
- Understanding Certificates on page 440
- Understanding Certificate Revocation Lists on page 443

This topic covers:



## PKI Hierarchy for a Single CA Domain or Across Domains

Figure 105 on page 444 shows the structure of a single-domain certificate authority.

**Figure 105: PKI Hierarchy of Trust—CA Domain**

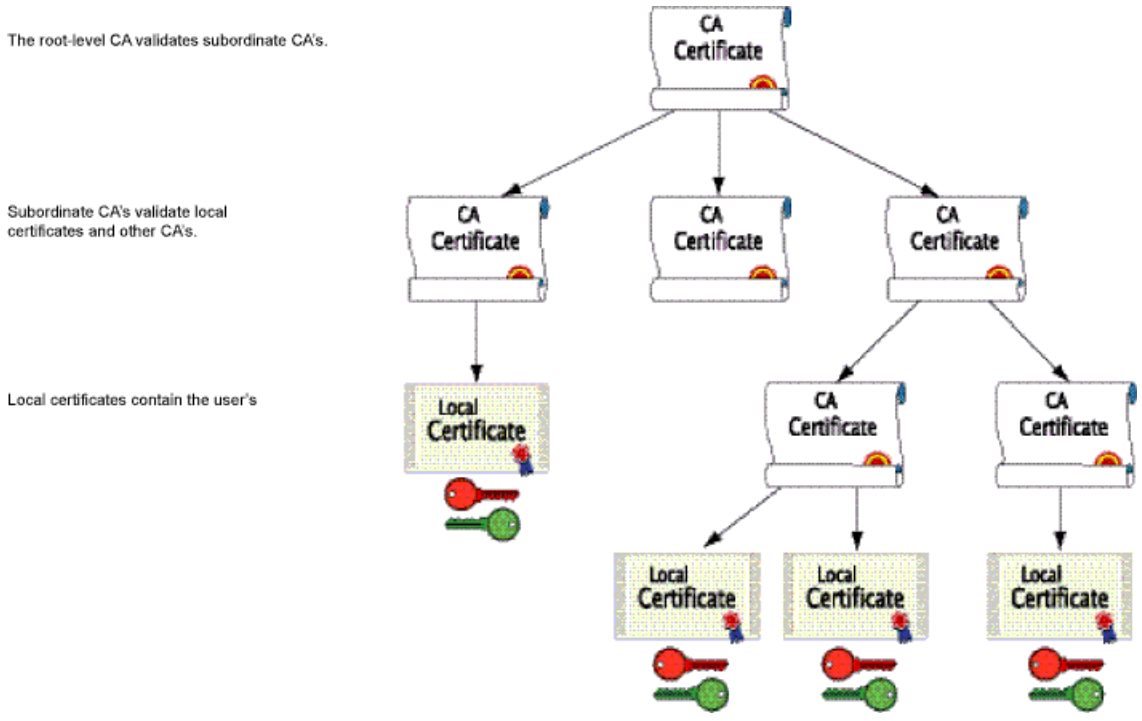


If certificates are used solely within an organization, that organization can have its own CA domain within which a company CA issues and validates certificates for its employees. If that organization later wants its employees to exchange their certificates

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 106 on page 445.

**Figure 106: Cross-Certification**

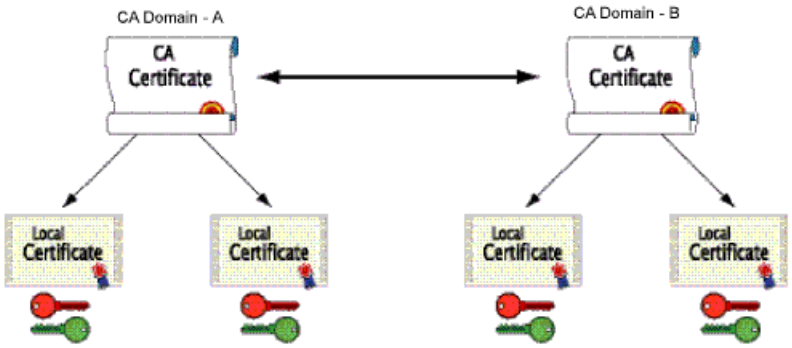


Users in the CA domain A can use their certificates and key pairs with users in CA domain B because the CA's have cross-certified each other.

## 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 450
- Generating a Local Certificate Request Manually on page 456

## 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 440
- Understanding Certificates on page 440

**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*

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 447
- Understanding Manually Generated Self-Signed Certificates on page 448
- Using Automatically Generated Self-Signed Certificates on page 465
- Manually Generating Self-Signed Certificates on page 466
- Deleting Certificates on page 468
- Understanding Public Key Cryptography on page 440

## 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 446
- Understanding Certificates on page 440

### *Related Topics*

- Using Automatically Generated Self-Signed Certificates on page 465
- Understanding Manually Generated Self-Signed Certificates on page 448
- Manually Generating Self-Signed Certificates on page 466
- Deleting Certificates on page 468
- Understanding Public Key Cryptography on page 440

## 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 446
- Understanding Certificates on page 440
- Understanding Public Key Cryptography on page 440

You generate this kind of named local certificate using the command-line interface (CLI).

### *Related Topics*

- Manually Generating Self-Signed Certificates on page 466
- Understanding Automatically Generated Self-Signed Certificates on page 447
- Using Automatically Generated Self-Signed Certificates on page 465
- Deleting Certificates on page 468

## 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 440
- Understanding Certificates on page 440
- Understanding Certificate Revocation Lists on page 443
- Understanding Public Key Infrastructure on page 443

This topic covers:



## *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 450.
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 451.
3. Enroll the CA certificate onto the device. See “Enrolling a CA Certificate Online” on page 453.
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 454.
5. Configure automatic re-enrollment. See “Understanding SecurID User Authentication” on page 168.

## *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 450.
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 451.

3. Generate a certificate request using the key pair, and manually copy that request and paste it into the appropriate field at the CA Web site to obtain a personal certificate (also known as a local certificate). See “Generating a Local Certificate Request Manually” on page 456.
4. Load the certificate onto the device. See “Loading CA and Local Certificates Manually” on page 458.
5. Configure automatic re-enrollment. See “Understanding SecurID User Authentication” on page 168.
6. If necessary, load the certificate's CRL on the device. See “Manually Loading a CRL onto the Device” on page 461.

### *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 462.
- To verify manually with CRLs, see “Checking Certificate Validity Using CRLs” on page 463.

### *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 440
- Understanding Certificates on page 440
- Understanding Certificate Revocation Lists on page 443
- Using Digital Certificates on page 448

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

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 451.

### Related Topics



## 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<br>■ Public Key Cryptography for Certificates on page 439<br>■ Understanding Certificates on page 440<br>■ Understanding Certificate Revocation Lists on page 443<br>■ Using Digital Certificates on page 448 |

This topic covers:



## 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

## 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 450.
2. Configure a CA profile. See "Configuring a Certificate Authority Profile" on page 451.
3. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

This topic covers:



### 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 454.

### Related Topics

## 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 450.
2. Configure a CA profile. See "Configuring a Certificate Authority Profile" on page 451.
3. Enroll a CA certificate. See "Enrolling a CA Certificate Online" on page 453.
4. For background information, read:
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

This topic covers:



### 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 email 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 459

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 168.

## *Related Topics*

- Using Digital Certificates on page 448
- Generating a Local Certificate Request Manually on page 456
- Loading CA and Local Certificates Manually on page 458
- Verifying Certificate Validity on page 462
- Checking Certificate Validity Using CRLs on page 463
- Deleting a Loaded CRL on page 469

## 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 450.
2. Create a CA profile. See "Configuring a Certificate Authority Profile" on page 451.
3. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

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*

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 Web site 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 458.

## Related Topics

- Loading CA and Local Certificates Manually on page 458
- Re-enrolling Local Certificates Automatically on page 459
- Verifying Certificate Validity on page 462
- Checking Certificate Validity Using CRLs on page 463
- Deleting Certificates on page 468
- Deleting a Loaded CRL on page 469

## 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 450.
2. Create a CA profile. See “Configuring a Certificate Authority Profile” on page 451.
3. Generate a certificate request. See “Generating a Local Certificate Request Manually” on page 456 .
4. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

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*

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 459.

### *Related Topics*

- Using Digital Certificates on page 448
- Generating a Public-Private Key Pair on page 450
- Generating a Local Certificate Request Manually on page 456
- Deleting Certificates on page 468
- Re-enrolling Local Certificates Automatically on page 459
- Verifying Certificate Validity on page 462
- Checking Certificate Validity Using CRLs on page 463
- Deleting a Loaded CRL on page 469

## 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 449 or "Obtaining Digital Certificates Manually" on page 449.
2. For background information, read:
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

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

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

## 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 450.
2. Generate a certificate request. See "Generating a Local Certificate Request Manually" on page 456.
3. Configure a CA profile. See "Configuring a Certificate Authority Profile" on page 451.
4. Load your certificate onto the device. See "Loading CA and Local Certificates Manually" on page 458.
5. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

This topic covers:



### 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

## 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 449 or "Obtaining Digital Certificates Manually" on page 449.
2. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

You verify certificates from the CLI in operational mode.

This topic covers:



### *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 448
- Generating a Public-Private Key Pair on page 450
- Generating a Local Certificate Request Manually on page 456
- Loading CA and Local Certificates Manually on page 458
- Re-enrolling Local Certificates Automatically on page 459
- Checking Certificate Validity Using CRLs on page 463
- Deleting Certificates on page 468
- Deleting a Loaded CRL on page 469

## 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 449 or "Obtaining Digital Certificates Manually" on page 449.
2. For background information, read:
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

**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

To configure a certificate authority profile.

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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

### 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 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 446
- Understanding Manually Generated Self-Signed Certificates on page 448
- Understanding Certificates on page 440

This topic covers:



### *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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 on page 74.

### *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 448
- Manually Generating Self-Signed Certificates on page 466
- Understanding Public Key Infrastructure on page 443

## 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 446
- Understanding Automatically Generated Self-Signed Certificates on page 447
- Understanding Certificates on page 440

This topic covers:



### 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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 on page 74.

### CLI Configuration

Use the following CLI command to manually generate a self-signed certificate created and signed by the user whose email 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

## 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 449 or "Obtaining Digital Certificates Manually" on page 449.
2. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

This topic covers:



### *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 448
- Generating a Public-Private Key Pair on page 450

- Generating a Local Certificate Request Manually on page 456
- Loading CA and Local Certificates Manually on page 458
- Re-enrolling Local Certificates Automatically on page 459
- Checking Certificate Validity Using CRLs on page 463
- Deleting a Loaded CRL on page 469
- Understanding Automatically Generated Self-Signed Certificates on page 447

## 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 449 or "Obtaining Digital Certificates Manually" on page 449.
2. For background information, read
   - Public Key Cryptography for Certificates on page 439
   - Understanding Certificates on page 440
   - Understanding Certificate Revocation Lists on page 443
   - Using Digital Certificates on page 448

This topic covers:

- CLI Operation on page 469
- Related Topics on page 469

### *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 448
- Generating a Public-Private Key Pair on page 450
- Generating a Local Certificate Request Manually on page 456
- Loading CA and Local Certificates Manually on page 458
- Re-enrolling Local Certificates Automatically on page 459

- Checking Certificate Validity Using CRLs on page 463
- Deleting Certificates on page 468

Chapter 16

# 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.

JUNOS software supports the ALGs described in the following sections.

This section includes:

## 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*



## 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. For SRX-series devices, SIP, MGCP, FTP, and TFTP ALGs are disabled by default.

| **Before You Begin** |
| --- |
| For background information, read “Understanding Application Layer Gateways” on page 473 |

**Figure 107: Quick Configuration Page for General ALGs**

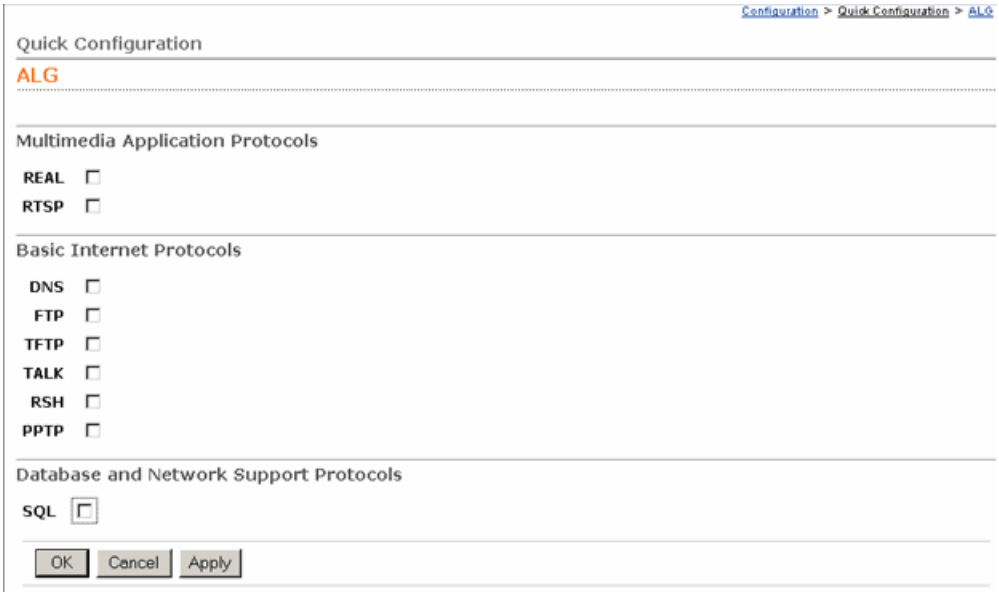


To enable or disable an ALG with Quick Configuration:

1. In the J-Web user interface, select Configuration > Quick Configuration > ALG > General ALG.

   Figure 107 on page 474 shows the General ALG page.

2. Select the check box next to an ALG, described in Table 72 on page 474, 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 72: General Configuration Options**

| Field | Function | Action |
|---|---|---|
| **Multimedia Application Protocols** | | |
| REAL | Provides an ALG for the RealAudio and RealVideo Protocol. The REAL ALG processes Progressive Networks Audio (PNA) packets over the TCP connection and looks for the control commands in the packet where the port number is embedded. It performs NAT and opens gates for the UDP data connection. | Select the check box to enable the ALG. |
| RTSP | Provides an ALG for the Real-Time Streaming Protocol. | Select the check box to enable the ALG. |

**Table 72: General Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| **Multimedia Application Protocols** | | |
| **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. |
| 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. |

**Table 72: General Configuration Options** *(continued)*

| Field | Function | Action |
|---|---|---|
| **Multimedia Application Protocols** | | |
| 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** | | |
| 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 471.

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 108 on page 477.)

**Figure 108: H.323 Protocol**

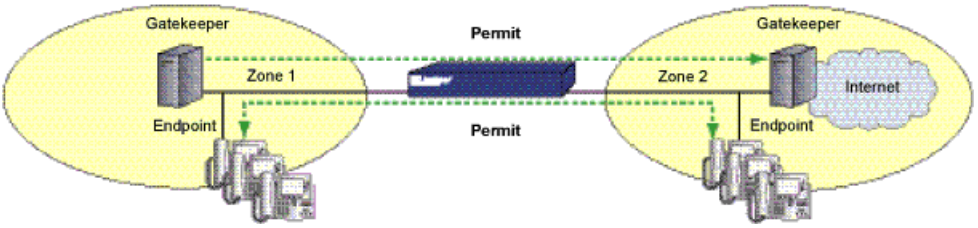


**NOTE:** The illustrations use IP phones for illustrative purposes, although it is possible to make configurations for other hosts that use VoIP, such as Microsoft NetMeeting multimedia devices.

## *Related Topics*

- Understanding the SIP ALG on page 499
- Understanding the SCCP ALG on page 561
- Understanding the MGCP ALG on page 578
- Setting H.323 Endpoint Registration Timeout on page 480
- Setting H.323 Media Source Port Range on page 481
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 482
- Allowing Unknown H.323 Message Types on page 483
- Verifying the H.323 Configuration on page 485
- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 487
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 490

- Using NAT and the H.323 ALG to Enable Incoming Calls on page 498
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 496

## 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 476 |

To configure the H.323 ALG with Quick Configuration:

**Figure 109: H.323 ALG Configurationy**

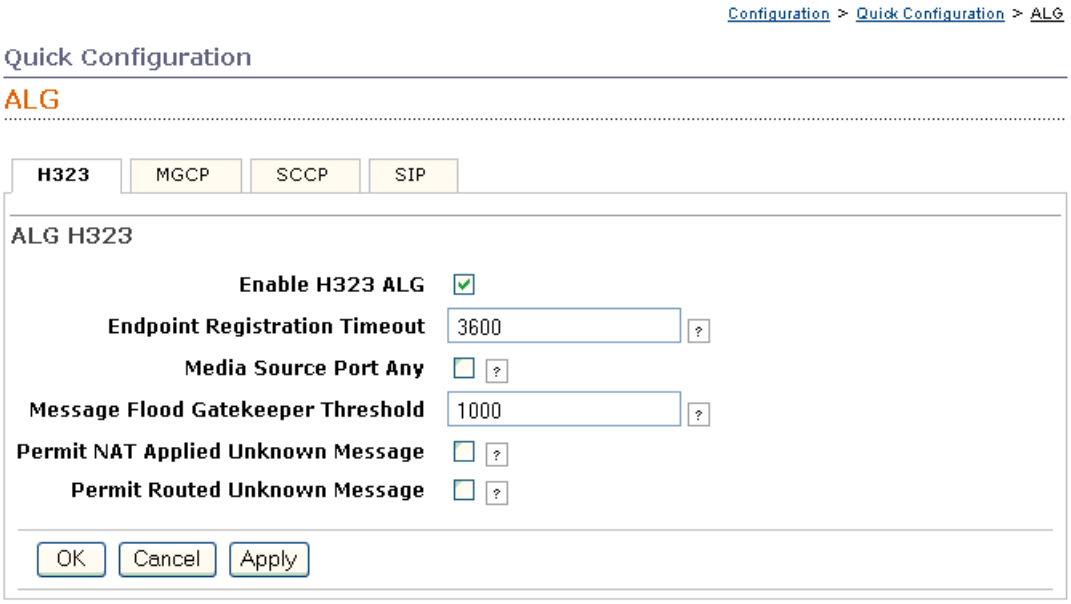


To configure the H.323 ALG with Quick Configuration:

1. In the J-Web user interface, select Configuration > Quick Configuration > ALG > VoIP ALG.

   Figure 109 on page 478 shows the H.323 ALG page.

2. Select the H.323 tab if it is not selected.
3. Fill in the parameter settings as described in Table 73 on page 479 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 73: H.323 Configuration Options**

| | | |
|---|---|---|
| Enable H323 ALG | Enable or disable the H.323 ALG | Click the check box. |
| 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 J-series 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 J-series 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.<br><br>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 476.

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

To specify endpoint registration timeout using the J-Web configuration editor:

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to ALG, click Configure or Edit.
4. Next to H323, click **Configure** or Edit.
5. In the Endpoint-registration-timeout box, type 5000 and click **OK**.
6. If you are finished configuring the J-series device, commit the configuration.
7. To check the configuration, see "Verifying the H.323 Configuration" on page 485.

### 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 J-series device, commit the configuration.

To check the configuration, see “Verifying the H.323 Configuration” on page 485.

### *Related Topics*

- Allowing Unknown H.323 Message Types on page 483
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 482
- Setting H.323 Media Source Port Range on page 481
- Verifying the H.323 Configuration on page 485

## Setting H.323 Media Source Port Range

The media source port feature enables you to configure the J-series 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 476. |

By default, the J-series 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 481
- CLI Configuration on page 482
- Related Topics on page 482

### *J-Web Configuration*

To configure the J-series device to open a narrow gate for media traffic by disabling the media source port feature using the J-Web configuration editor:

1. Select Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to ALG, click Configure or Edit.
4. Next to H323, click **Configure** or Edit.
5. Select the Media source port any check box and click **OK**.

6. If you are finished configuring the J-series device, commit the configuration.
7. To check the configuration, see “Verifying the H.323 Configuration” on page 485.

### CLI Configuration

In this example, you configure the J-series 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 J-series device, commit the configuration.

To check the configuration, see “Verifying the H.323 Configuration” on page 485.

### Related Topics

- Setting H.323 Endpoint Registration Timeout on page 480
- Allowing Unknown H.323 Message Types on page 483
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 482
- Verifying the H.323 Configuration on page 485

## 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 476. |

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 483
- CLI Configuration on page 483
- Related Topics on page 483

### *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 **Configuration > View and Edit >** Edit **Configuration**.

   The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to ALG, click Configure or Edit.
4. Next to H323, 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 J-series device, commit the configuration.
9. To check the configuration, see "Verifying the H.323 Configuration" on page 485.

### *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 J-series device, commit the configuration.

To check the configuration, see "Verifying the H.323 Configuration" on page 485.

### *Related Topics*

- Setting H.323 Endpoint Registration Timeout on page 480
- Setting H.323 Media Source Port Range on page 481
- Allowing Unknown H.323 Message Types on page 483
- Verifying the H.323 Configuration on page 485

## 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 J-series 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 476.

This feature enables you to specify how unidentified H.323 messages are handled by the J-series 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 J-series 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*

To configure the J-series device to allow unknown H.323 message types in both route and NAT modes using the J-Web configuration editor:

1. Select **Configuration > View and Edit >** Edit **Configuration**.

   The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to ALG, click Configure or Edit.
4. Next to H323, 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 J-series device, commit the configuration.
9. To check the configuration, see "Verifying the H.323 Configuration" on page 485.

### *CLI Configuration*

In this example, you configure the J-series 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 J-series device, commit the configuration.

To check the configuration, see "Verifying the H.323 Configuration" on page 485.

### *Related Topics*

- Setting H.323 Endpoint Registration Timeout on page 480
- Setting H.323 Media Source Port Range on page 481
- Configuring H.323 Denial of Service (DoS) Attack Protection on page 482
- Verifying the H.323 Configuration on page 485

## 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 476

## 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 476. |

In this example, the J-series device can be in either route or NAT mode. See Figure 110 on page 487.

**Figure 110: 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 487
- CLI Configuration on page 489
- Related Topics on page 490

### *J-Web Configuration*

To configure an address book, to configure a 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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 **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 **Configuration > View and Edit > Edit Configuration**.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 J-series 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 J-series device, commit the configuration.

### *Related Topics*

- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 490
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 496
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 498
- Understanding the SIP ALG on page 499
- Setting SCCP Inactive Media Timeout on page 569
- Configuring a Media Gateway in Subscribers' Homes on page 594

## Passing H.323 ALG Traffic to a Gatekeeper in the External Zone

Because route mode does not require address mapping of any kind, a J-series 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 476. |

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 J-series device can be in transparent or route mode. See Figure 111 on page 491.

**Figure 111: 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 491
- CLI Configuration on page 494
- Related Topics on page 495


## *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 **Configuration > View and Edit > Edit Configuration**.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
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 **Configuration > View and Edit > Edit Configuration**.

   The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration.

   The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Select the Policies 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 J-series 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 J-series device, commit the configuration.

## *Related Topics*

- Passing H.323 ALG Traffic to a Gatekeeper in the Internal Zone on page 487
- Using NAT and the H.323 ALG to Enable Incoming Calls on page 498
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 496
- Understanding the SIP ALG on page 499

- Setting SCCP Inactive Media Timeout on page 569
- Configuring a Media Gateway in Subscribers' Homes on page 594

## 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 J-series 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 476. |

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 112 on page 496.

**Figure 112: Network Address Translation—Outgoing Calls**

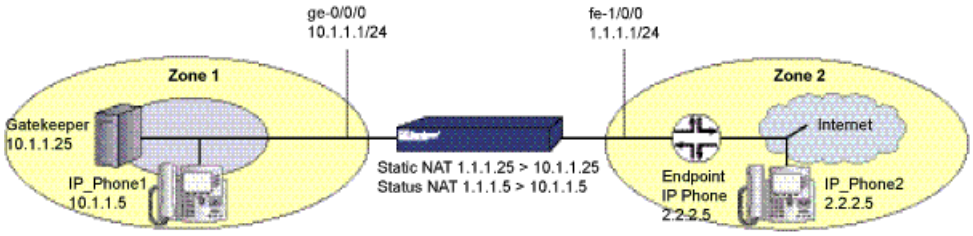


You can use either J-Web or the CLI configuration editor to enable outgoing calls.

This topic covers:

- CLI Configuration on page 496
- Related Topics on page 497

### *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 487
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 490
- Using NAT and the H.323 ALG to Enable Outgoing Calls on page 496
- Understanding the SIP ALG on page 499
- Setting SCCP Inactive Media Timeout on page 569
- Configuring a Media Gateway in Subscribers' Homes on page 594

## Using NAT and the H.323 ALG to Enable Incoming Calls

In this example, you configure the J-series 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 113 on page 498.

| Before You Begin |
| --- |
| For background information, read "Understanding the H.323 ALG" on page 476. |

**Figure 113: Network Address Translation—Incoming Calls**

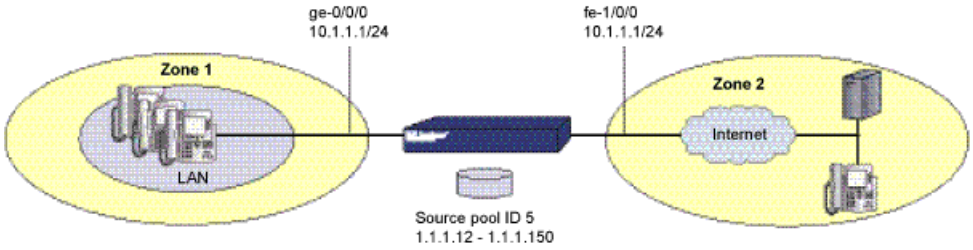


Interface NAT is when the source pool uses the same address as an interface IP address. You can use such address entries as destination addresses in policies, together with H.323, SIP, or other VoIP protocols, to support incoming calls.

To configure incoming calls using NAT, use either the J-Web or CLI Configuration editor.

This topic covers:



### *CLI Configuration*

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 487
- Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 490
- Understanding the SIP ALG on page 499
- Setting SCCP Inactive Media Timeout on page 569
- Configuring a Media Gateway in Subscribers' Homes on page 594

## 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 471.

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*

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 499.

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 J-series 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 J-series 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 504 and "Classes of SIP Responses" on page 524). You can allow SIP transactions to traverse the Juniper Networks firewall by creating a static policy that permits SIP service. This policy enables the J-series 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 J-series 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 J-series 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 J-series 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 J-series 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 114 on page 504 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 114: SIP ALG Call Setup**

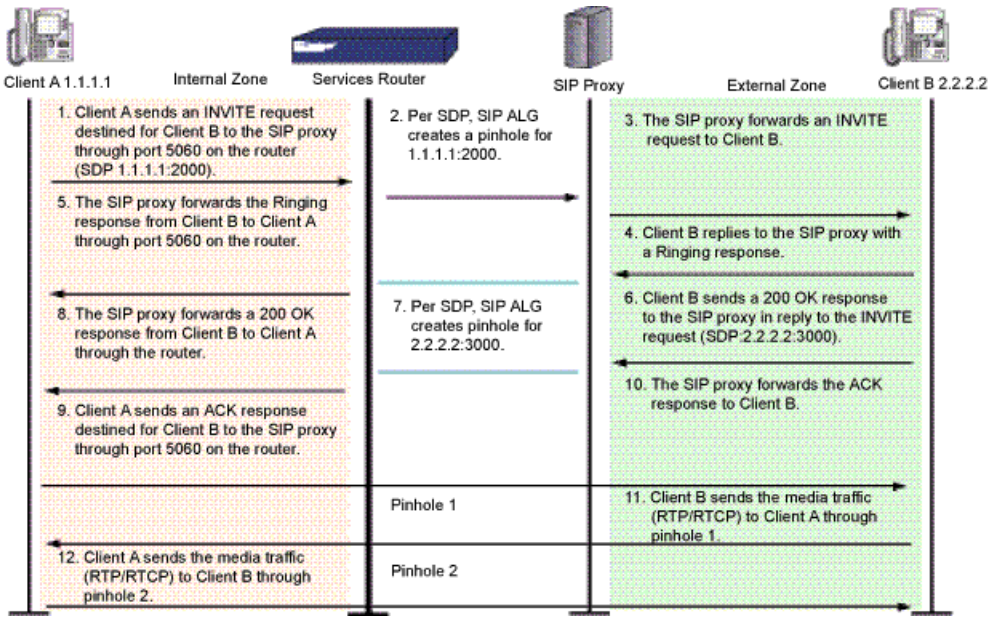


**NOTE:** The SIP ALG does not create pinholes for RTP and RTCP traffic when the destination IP address is 0.0.0.0, which indicates that the session is on hold. To put a session on hold during a telephone communication, for example, user A sends user B a SIP message in which the destination IP address 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 J-series 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 499.

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 508

## 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 499 |

Figure 115 on page 506 shows the SIP ALG configuration page.

**Figure 115: SIP ALG Configuration**

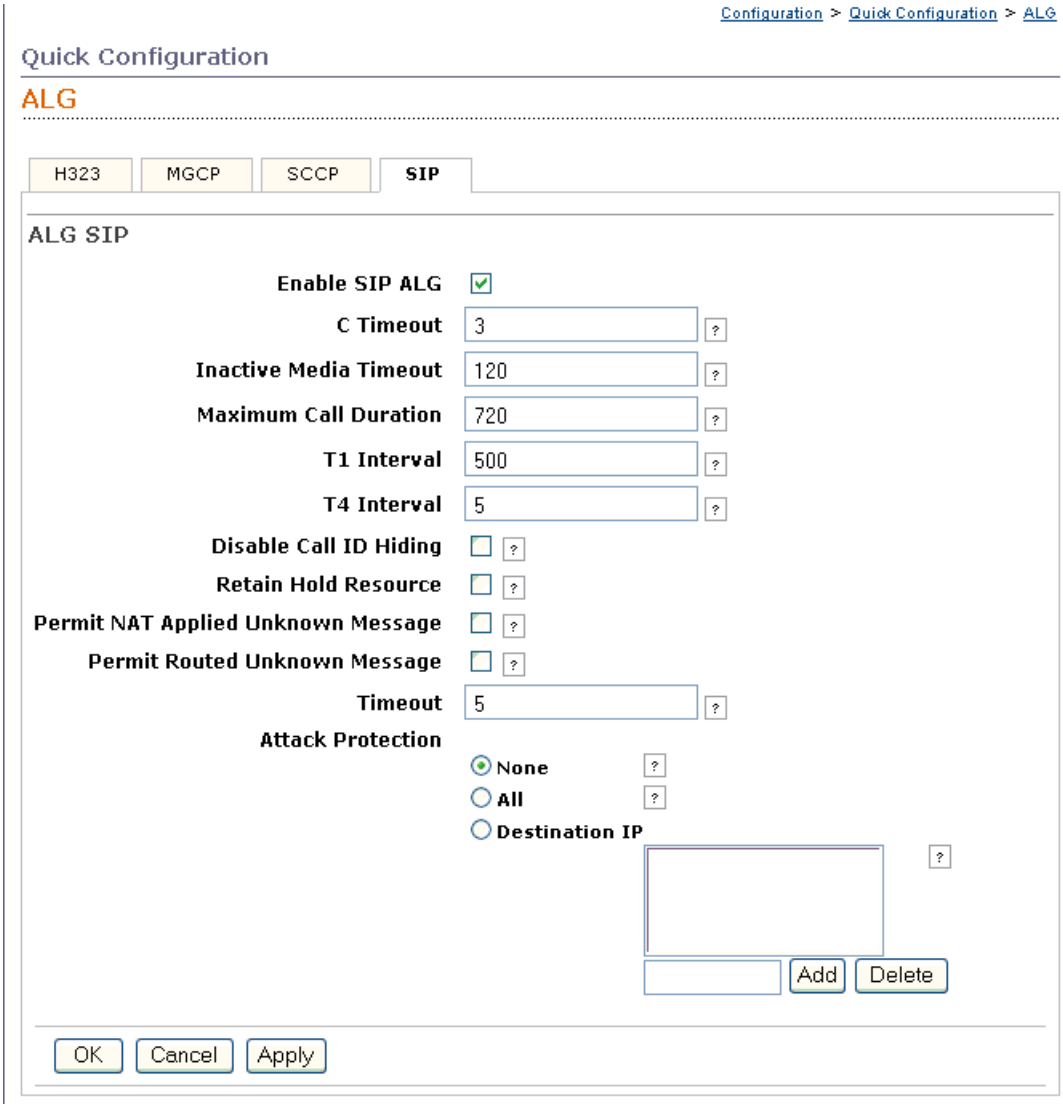


To configure the SIP ALG with Quick Configuration:

1. In the J-Web user interface, select Configuration > Quick Configuration > ALG > VoIP ALG.
2. Select the SIP tab if it is not selected.
3. Fill in the parameter settings as described in Table 74 on page 507 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 74: 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 J-series 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 74: SIP ALG Configuraiton Options** *(continued)*

| | | |
|---|---|---|
| Permit NAT Applied Unknown Message | Specifies how unidentified SIP messages are handled by the J-series 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.<br><br>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 499
- SIP ALG Request Methods Overview on page 504

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 J-series 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 J-series 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 509

## 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<br>■ Understanding the SIP ALG on page 499<br>■ Understanding SIP ALG Call Duration and Timeouts on page 508 |

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 510
- CLI Configuration on page 510
- Related Topics on page 511


## J-Web Configuration

1. Select Configuration > View and Edit > Edit Configuration > Security > ALG > SIP.
2. Enter a value for any of the following parameters:
3. C timeout
4. Inactive media timeout
5. Maximum call duration
6. T1 interval
7. T4 interval
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 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 508
- Configuring SIP Denial of Service (DoS) Attack Protection on page 511
- Allowing Unknown SIP Message Types on page 512

## 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 J-series 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 524), the ALG stores the source IP address of the request and the IP address of the proxy server in a table. Subsequently, the J-series 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 J-series 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 499
- SIP ALG Request Methods Overview on page 504

To configure DoS attack protection, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 511
- CLI Configuration on page 512
- Related Topics on page 512

### *J-Web Configuration*

1. In the J-Web user interface, select Configuration > View and Edit > Edit Configuration > Security > ALG > SIP.
2. Click Application Screen.
3. Click Protect.
4. Click the Deny check box.
5. 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 J-series 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 J-series 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 511
- Allowing Unknown SIP Message Types on page 512

## Allowing Unknown SIP Message Types

To accomodate 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 J-series 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 499
- SIP ALG Request Methods Overview on page 504

This feature enables you to specify how unidentified SIP messages are handled by the J-series 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 J-series 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*

1. Select Configuration > View and Edit > Edit Configuration > Security Edit > ALG Configure > SIP Configure > Application screen configure > Unknown message Configure.
2. Click one of the following check boxes:
   - To allow unknown messge types in NAT mode, click the Permit nat applied Yes check box.
   - To allow unknown messge types in route mode, click the Permit routed Yes check box.
3. 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 J-series 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*

## 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 J-series 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 499
- SIP ALG Request Methods Overview on page 504

To disable SIP call hiding, use either the J-Web or the CLI configuration editor.

This topic covers:



### J-Web Configuration

1. Select Configuration > View and Edit > Edit Configuration > Security Edit > ALG Configure > SIP Configure.
2. Click the Disable call id hiding Yes check box.
3. 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 508
- Setting SIP Call Duration and Inactive Media Timeout on page 509

- Configuring SIP Denial of Service (DoS) Attack Protection on page 511

## 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 J-series 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 499
- SIP ALG Request Methods Overview on page 504

To retain SIP hold resources, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 515
- CLI Configuration on page 516
- Related Topics on page 516

### *J-Web Configuration*

1. Select Configuration > View and Edit > Edit Configuration > Security > ALG > SIP.
2. Click the Retain hold resources check box.
3. 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 accomodate proprietary SIP call flows:

```
user@host# set security alg sip retain-hold-resource
```

### Related Topics

- Understanding the SIP ALG on page 499
- SIP ALG Request Methods Overview on page 504
- SDP Session Descriptions on page 502
- Pinhole Creation on page 502

## 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 499
- SIP ALG Request Methods Overview on page 504

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 J-series 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 J-series 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*

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 J-series 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 J-series 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 J-series 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 J-series 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 J-series 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 J-series 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 75 on page 520 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 75: 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 |

**Table 75: Requesting Messages with NAT Table continued on next page** *(continued)*

| | | |
|---|---|---|
| Outbound Request | To: | None |
| (from private to public) | 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 |
| 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 502.

### *SIP NAT Scenario*

Figure 116 on page 523 and Figure 117 on page 524 show a SIP call INVITE and 200 OK. In Figure 116 on page 523, 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 J-series 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 117 on page 524, 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 116: SIP NAT Scenario 1**

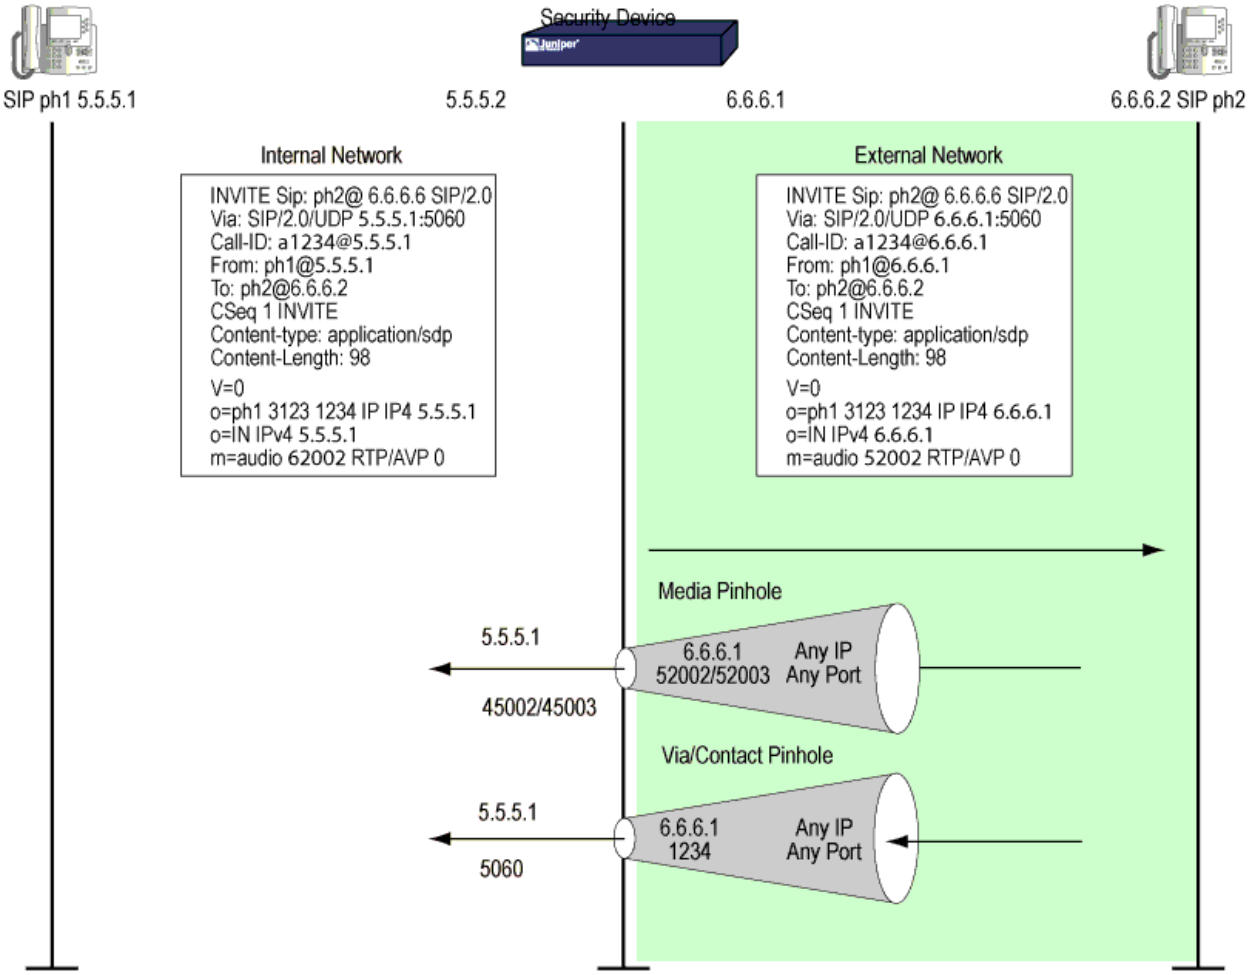

**Figure 117: SIP NAT Scenario 2**

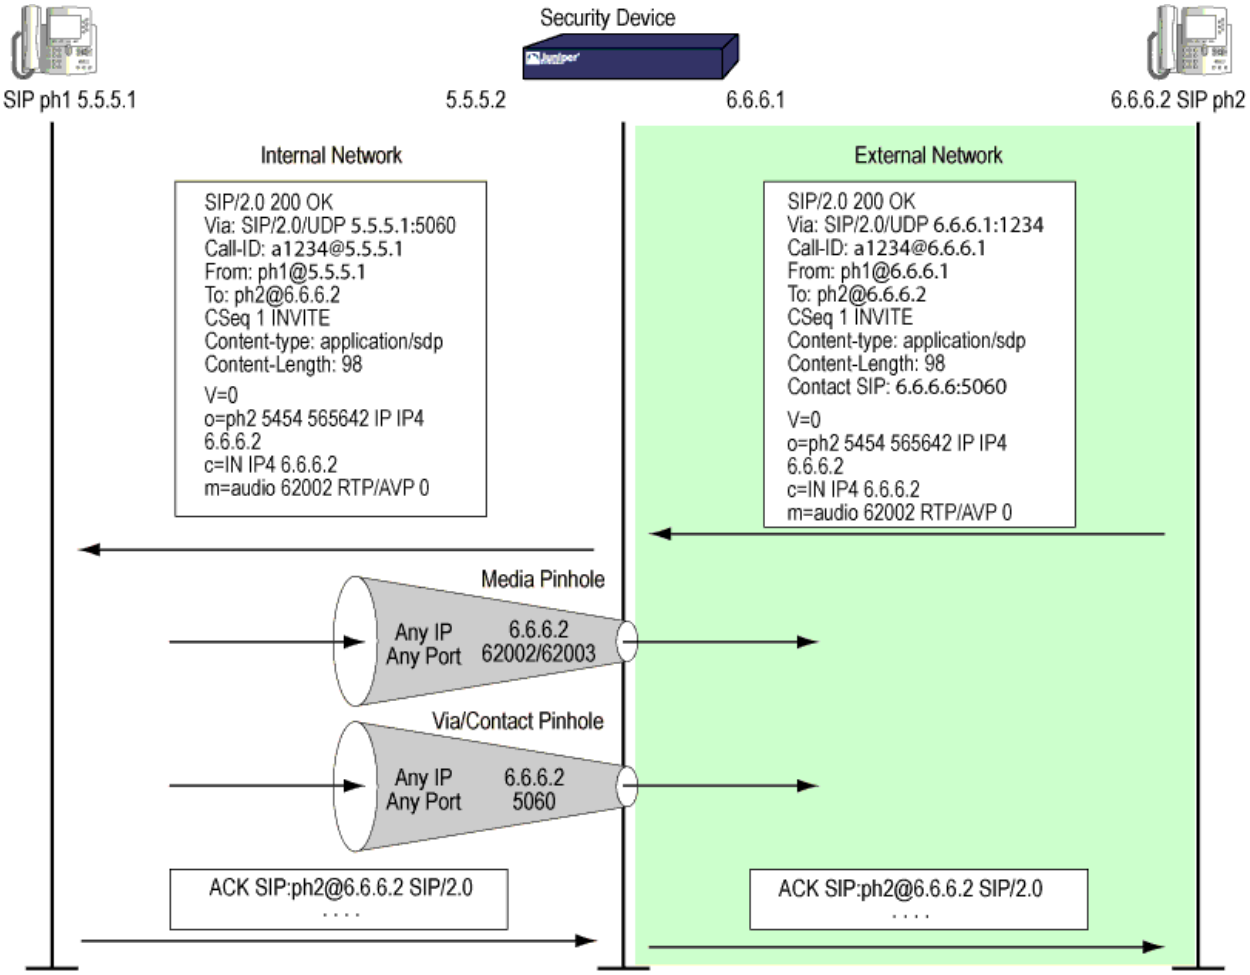


## *Classes of SIP Responses*

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 76 on page 525 provides a complete list of current SIP responses, all of which are supported on J-series devices.

**Table 76: 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 55
- Understanding Security Zones on page 49

## 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 118 on page 527. 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 499
- SIP ALG Request Methods Overview on page 504
- Understanding SIP with Network Address Translation (NAT) on page 516

The J-series 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 J-series 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 J-series 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 118 on page 527.

**Figure 118: Using the SIP Registrar**

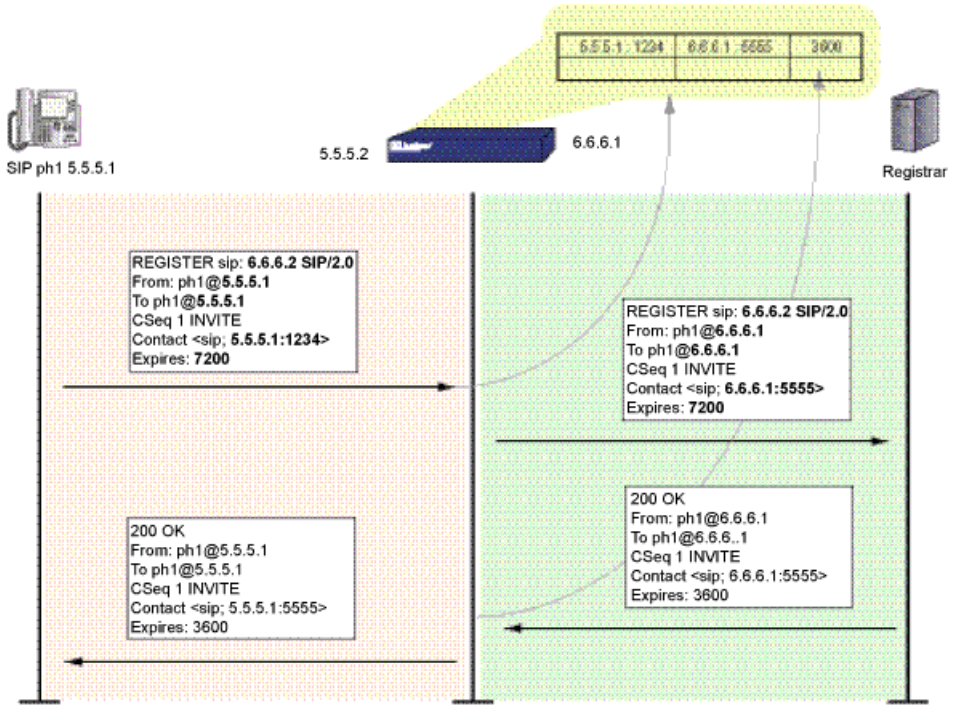


## *Related Topics*

- Configuring Interface Source NAT for Incoming SIP Calls on page 528
- Configuring a Source NAT Pool for Incoming SIP Calls on page 530
- Configuring Static NAT for Incoming SIP Calls on page 535
- Configuring the SIP Proxy in the Private Zone on page 540
- Configuring the SIP Proxy in the Public Zone on page 542
- Configuring a Three-Zone SIP Scenario on page 547
- Incoming Calls on page 517
- Outgoing Calls on page 517

## 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 499
- SIP ALG Request Methods Overview on page 504
- Understanding SIP with Network Address Translation (NAT) on page 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525

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 119 on page 528.

**Figure 119: Source NAT for Incoming Calls**

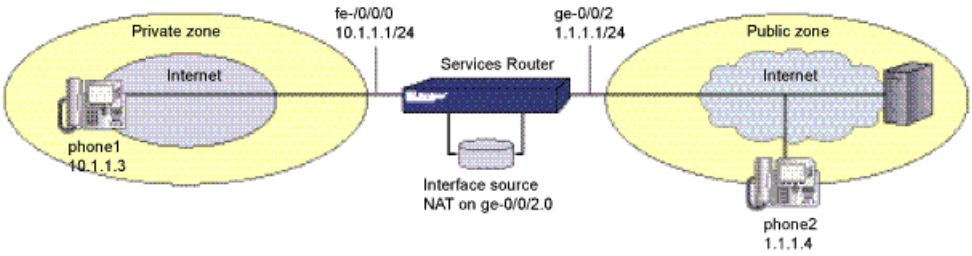


To configure interface source NAT for incoming calls, use either the J-Web or the CLI configuration editor.

This topic covers:

- CLI Configuration on page 528
- Related Topics on page 529

### 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 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525
- Configuring Static NAT for Incoming SIP Calls on page 535
- Configuring a Source NAT Pool for Incoming SIP Calls on page 530
- Configuring the SIP Proxy in the Private Zone on page 540
- Configuring the SIP Proxy in the Public Zone on page 542
- Configuring a Three-Zone SIP Scenario on page 547

## 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 499
- SIP ALG Request Methods Overview on page 504
- Understanding SIP with Network Address Translation (NAT) on page 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525

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 120 on page 530.

**Figure 120: Source NAT Pool for Incoming Calls**

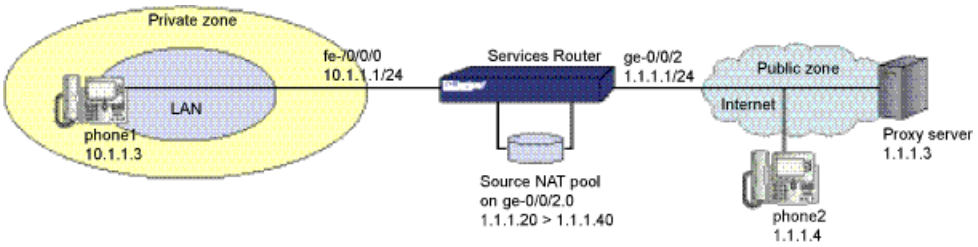


To configure a source NAT pool for incoming calls, use either the J-Web or CLI configuration editor.

This topic covers:

- J-Web Configuration on page 530
- CLI Configuration on page 534
- Related Topics on page 535

### *J-Web Configuration*

To configure interfaces:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Source Nat, click Configure.
4. Next to Address persistent, select the check box and click OK.
5. Next to Interface, click Add new entry.
6. In the Name box, type **ge-0/0/2.0**.
7. Next to Source nat, click Configure.
8. Next to Pool, click Add new entry.
9. In the Name box, type **sip-pool**.
10. Next to Address range, click Add new entry.
11. Next to High box, type **1.1.1.60** and next to Low box, type **1.1.1.20** and click OK.
12. Next to Allow incoming, select the check box and click OK.
13. To save and commit the configuration, click Commit.

To configure policies:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525
- Configuring Interface Source NAT for Incoming SIP Calls on page 528
- Configuring Static NAT for Incoming SIP Calls on page 535
- Configuring the SIP Proxy in the Private Zone on page 540
- Configuring the SIP Proxy in the Public Zone on page 542
- Configuring a Three-Zone SIP Scenario on page 547

## 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 499
- "SIP ALG Request Methods Overview" on page 504
- "Understanding SIP with Network Address Translation (NAT)" on page 516
- "Understanding Incoming SIP Call Support Using the SIP Registrar" on page 525

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 528 and "Configuring a Source NAT Pool for Incoming SIP Calls" on page 530, 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 121 on page 536.

**Figure 121: Static NAT for Incoming Calls**

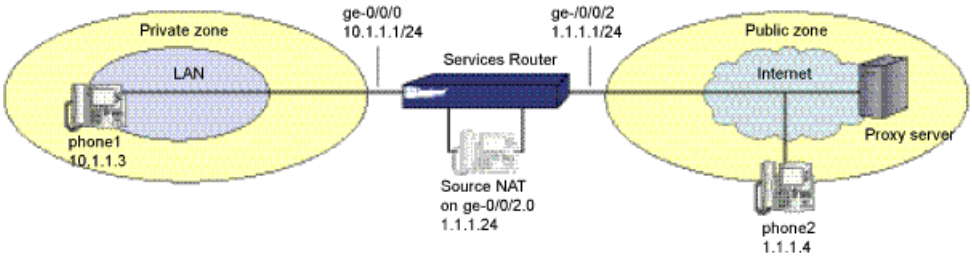


To configure static NAT for incoming calls, use either the J-Web or CLI configuration editor.

This topic covers:



## *J-Web Configuration*

To configure interfaces:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Interfaces, click Configure or Edit.
3. Next to Interfaces, 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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525
- Configuring Interface Source NAT for Incoming SIP Calls on page 528
- Configuring a Source NAT Pool for Incoming SIP Calls on page 530
- Configuring the SIP Proxy in the Private Zone on page 540
- Configuring a Three-Zone SIP Scenario on page 547

## 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 499
- SIP ALG Request Methods Overview on page 504
- Understanding SIP with Network Address Translation (NAT) on page 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525

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 122 on page 540.

**Figure 122: Proxy in the Private Zone**

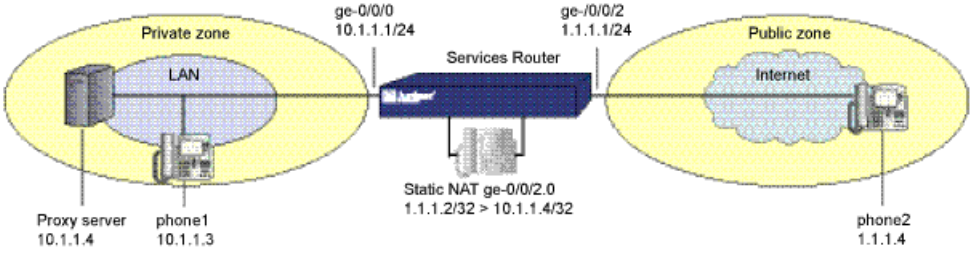


Use either the J-Web or CLI configuration editor.

This topic covers:



### *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 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525
- Configuring a Source NAT Pool for Incoming SIP Calls on page 530
- Configuring Static NAT for Incoming SIP Calls on page 535
- Configuring the SIP Proxy in the Private Zone on page 540
- Configuring a Three-Zone SIP Scenario on page 547
- Configuring Static NAT for Incoming SIP Calls on page 535

## 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 499
- SIP ALG Request Methods Overview on page 504
- Understanding SIP with Network Address Translation (NAT) on page 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525

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 123 on page 542.

**Figure 123: Proxy in the Public Zone**

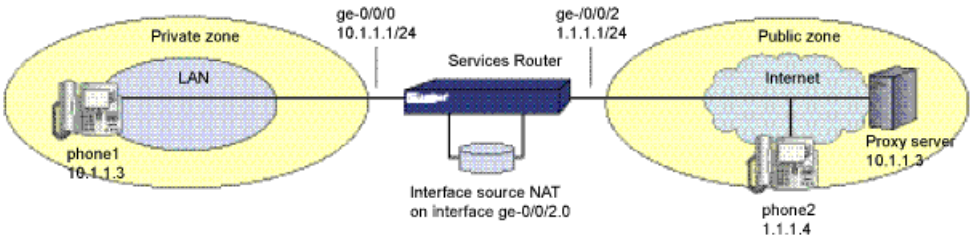


To configure the SIP proxy in the public zone, use either the J-Web or the CLI configuration editor.

This topic covers:



### *J-Web Configuration*

To configure interfaces:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 interface Source-Nat:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Nat, click Configure.
4. Next to Source nat, 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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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** 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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 516
- "Understanding Incoming SIP Call Support Using the SIP Registrar" on page 525
- "Configuring Interface Source NAT for Incoming SIP Calls" on page 528
- "Configuring a Source NAT Pool for Incoming SIP Calls" on page 530
- "Configuring Static NAT for Incoming SIP Calls" on page 535
- "Configuring the SIP Proxy in the Private Zone" on page 540
- "Configuring a Three-Zone SIP Scenario" on page 547

## 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 499.
- "SIP ALG Request Methods Overview" on page 504.
- "Understanding SIP with Network Address Translation (NAT)" on page 516.
- "Understanding Incoming SIP Call Support Using the SIP Registrar" on page 525.

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 124 on page 548 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 124: Three-Zone, Proxy in the DMZ**

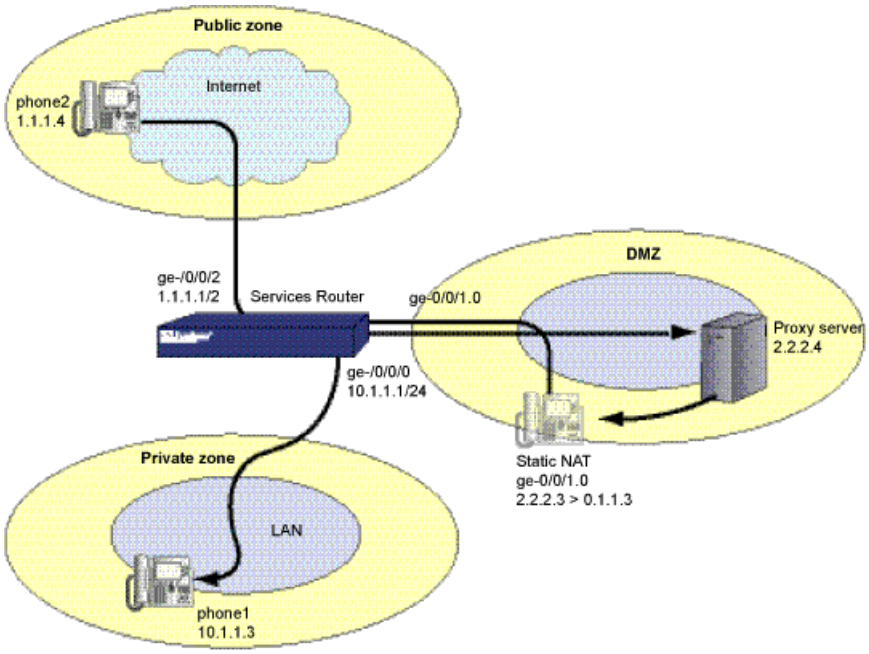


To configure a three-zone SIP scenario, use either the J-Web or CLI configuration editor.

This topic covers:



## *J-Web Configuration*

To configure interfaces:

1. Select **Configuration > View and Edit > Edit Configuration**. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 516
- Understanding Incoming SIP Call Support Using the SIP Registrar on page 525
- Configuring Interface Source NAT for Incoming SIP Calls on page 528
- Configuring a Source NAT Pool for Incoming SIP Calls on page 530
- Configuring Static NAT for Incoming SIP Calls on page 535

- Configuring the SIP Proxy in the Private Zone on page 540
- Configuring the SIP Proxy in the Public Zone on page 542

## Verifying the SIP Configuration

To verify the SIP configuration, perform these tasks:

- "Verifying SIP Calls" on page 557
- "Verifying SIP Call Detail" on page 557
- "Verifying SIP Transactions" on page 558
- "Verifying SIP Counters" on page 559
- "Verifying the Rate of SIP Messages" on page 560

### *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 499
- Verifying SIP Calls on page 557
- Verifying SIP Call Detail on page 557
- Verifying SIP Transactions on page 558
- Verifying SIP Counters on page 559
- Verifying the Rate of SIP Messages on page 560

### *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 499
- Verifying the SIP ALG on page 556
- Verifying SIP Call Detail on page 557
- Verifying SIP Transactions on page 558
- Verifying SIP Counters on page 559
- Verifying the Rate of SIP Messages on page 560


### *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 499
- Verifying the SIP ALG on page 556
- Verifying SIP Calls on page 557
- Verifying SIP Transactions on page 558
- Verifying SIP Counters on page 559
- Verifying the Rate of SIP Messages on page 560

## *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 499
- Verifying the SIP ALG on page 556
- Verifying SIP Calls on page 557
- Verifying SIP Call Detail on page 557
- Verifying SIP Counters on page 559
- Verifying the Rate of SIP Messages on page 560


### *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 499
- Verifying the SIP ALG on page 556
- Verifying SIP Calls on page 557
- Verifying SIP Call Detail on page 557
- Related Topics on page 559
- Verifying the Rate of SIP Messages on page 560


## *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 499
- Verifying the SIP ALG on page 556
- Verifying SIP Calls on page 557
- Verifying SIP Call Detail on page 557
- Verifying SIP Transactions on page 558
- Verifying SIP Counters on page 559


## 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 471
- Understanding NAT on page 276
- Pinhole Creation on page 502


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

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 563
- "CallManager" on page 563
- "Cluster" on page 563

### 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

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 J-series device allows this.

Table 77 on page 564 lists SCCP messages and indicates messages that are of interest to the J-series device.

**Table 77: SCCP Registration Messages**

| | | |
|---|---|---|
| RegisterMessage | | b |
| IPortMessage | | b |
| | RegisterAckMessage | b |
| | CapabilititsRequest | |
| CapabilitiesResMessage | | |
| ButtonTemplateReqMessage | | |
| | ButtonTemplateResMessage | |
| SoftKeyTemplateReqMessage | | |
| | SoftKeyTemplateResMessage | |

**Table 77: SCCP Registration Messages** *(continued)*

| | | |
|---|---|---|
| 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 J-series 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 125 on page 566 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 125: Call Setup and Teardown**

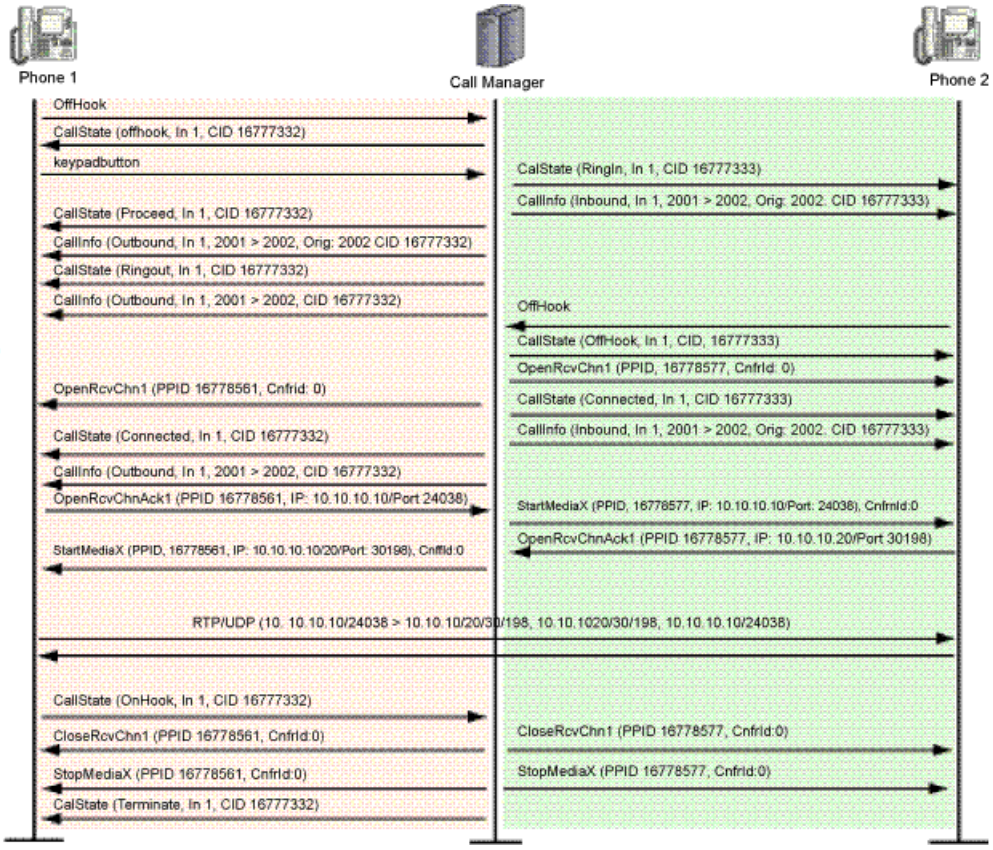


## *SCCP Messages*

Table 78 on page 566, Table 79 on page 566, Table 80 on page 567, and Table 81 on page 567 list the SCCP call message IDs in the four intervals allowed by the J-series device.

**Table 78: 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 79: CallManager to Station Messages**

| | |
|---|---|
| #define STATION_START_MEDIA_TRANSMISSION | 0x00000001 |
| #define STATION_STOP_MEDIA_TRANSMISSION | 0x00000002 |

**Table 79: 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 80: 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 81: 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 478
- Understanding the SIP ALG on page 499
- Understanding the MGCP ALG on page 578

## 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 561 |

Figure 126 on page 568 shows the SCCP ALG configuration page.

**Figure 126: SCCP ALG Configuration**

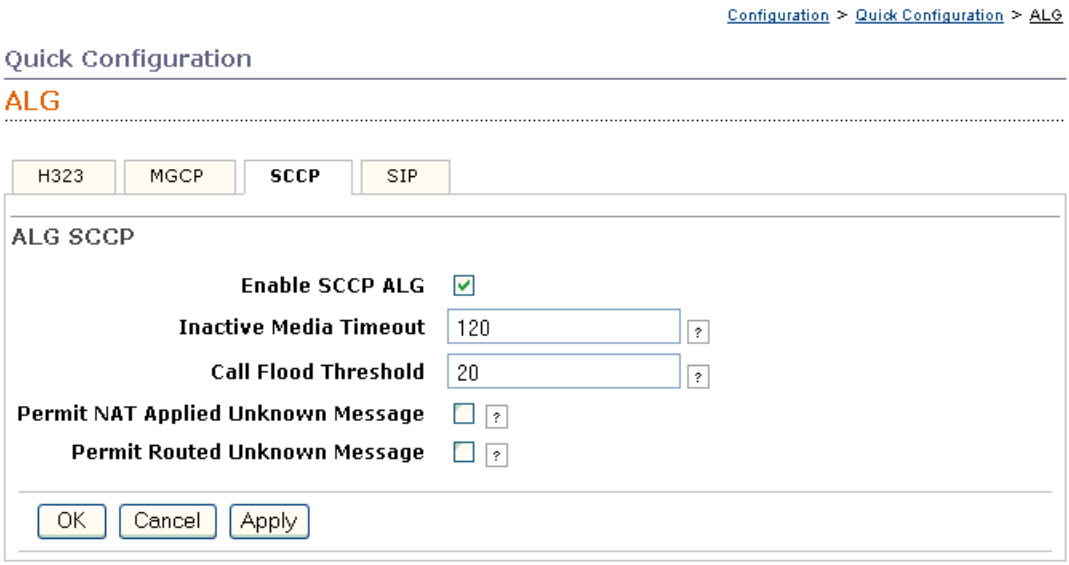


To configure the SCCP ALG with Quick Configuration:

1. In the J-Web user interface, select Configuration > Quick Configuration > ALG > VoIP ALG.
2. Select the SCCP tab if it is not selected.
3. Enter your parameter settings as described in Table 82 on page 568 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 82: 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 82: SCCP ALG Configuration Options** *(continued)*

| | | |
|---|---|---|
| Permit NAT Applied Unknown Message | Specifies how unidentified SCCP messages are handled by the J-series 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.<br><br>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 561.

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*

1. Select Configuration > View and Edit > Edit Configuration > Security Edit > ALG Configure > SCCP Configure.
2. Enter a value in the Inactive media timeout field,
3. 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*



## 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 J-series 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 561. |

This feature enables you to specify how unidentified SCCP messages are handled by the J-series 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

1. Select Configuration > View and Edit > Edit Configuration > Security Edit > ALG Configure > SCCP Configure > Application Screen Configure > Unknown message Configure
2. 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.
3. 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 561
- Setting SCCP Inactive Media Timeout on page 569
- Configuring SCCP Denial of Service (DoS) Attack Protection on page 572
- Verifying the SCCP Configuration on page 575

## 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 561. |

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 572
- CLI Configuration on page 573
- Related Topics on page 573

### *J-Web Configuration*

To configure the Services Router to drop any calls exceeding 500 per second per client:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
2. Next to Security, click Configure or Edit.
3. Next to Alg, click Configure.
4. Next to Sccp, click Configure.
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 J-series 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 561
- Setting SCCP Inactive Media Timeout on page 569
- Allowing Unknown SCCP Message Types on page 570
- Verifying the SCCP Configuration on page 575

## 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<br>■ Understanding the SCCP ALG on page 561<br>■ Configuring the SCCP ALG—Quick Configuration on page 567 |

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 127 on page 574.

**Figure 127: Call Manager/TFTP Server in the Private Zone**

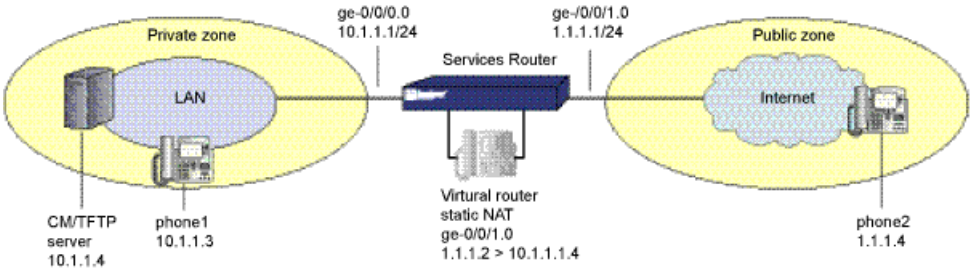


Use either the J-Web or CLI configuration editor.

This topic covers:

■ CLI Configuration on page 574

## *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 575
- "Verifying SCCP Call Details" on page 576
- "Verifying SCCP Counters" on page 577

### *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 561
- Verifying SCCP Calls on page 575
- Verifying SCCP Call Details on page 576
- Verifying SCCP Counters on page 577

### *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 561
- Verifying the SCCP ALG on page 575
- Verifying SCCP Call Details on page 576
- Verifying SCCP Counters on page 577

### *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 561
- Verifying the SCCP ALG on page 575
- Verifying SCCP Calls on page 575
- Verifying SCCP Counters on page 577

### *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 561
- Verifying the SCCP ALG on page 575
- Verifying SCCP Calls on page 575
- Verifying SCCP Call Details on page 576

## 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 471
- Understanding NAT on page 276
- Pinhole Creation on page 502

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*

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

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 83 on page 581 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 83: 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 83: MGCP Commands** *(continued)*

| | | | |
|---|---|---|---|
| DLCX | DeleteConnection—used by a call agent to instruct a gateway to delete an existing connection.<br><br>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<br><br>DLCX 9210 aaln/1@rgw-25.att.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8<br><br>Example 2: MG -> MGC<br><br>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:<br><br>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 83: 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 476
- Understanding the SIP ALG on page 499
- Configuring a Media Gateway in Subscribers' Homes on page 594

## 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 578 |

Figure 128 on page 584 shows the MGCP ALG configuration page.

**Figure 128: MGCP ALG Configuration**

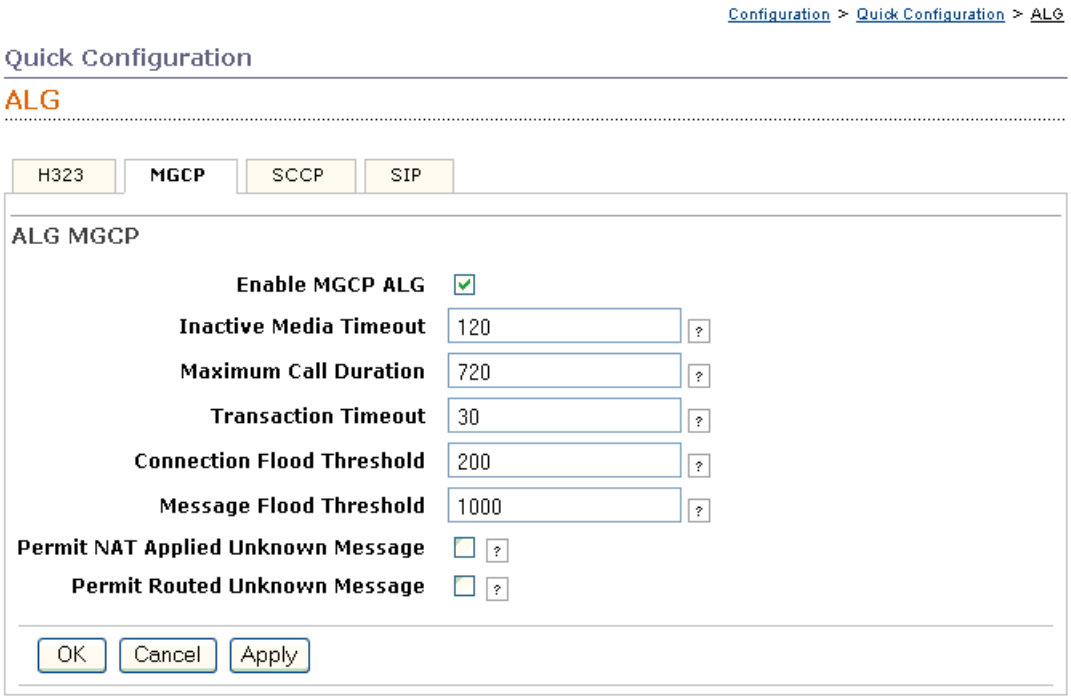

To configure the MGCP ALG with Quick Configuration:

1. In the J-Web user interface, select Configuration > Quick Configuration > ALG > VoIP ALG.
2. Select the MGCP tab if it is not selected.
3. Fill in the options as described in Table 84 on page 585 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 84: 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 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. |
| 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 84: MGCP Configuration Options** *(continued)*

| | | |
|---|---|---|
| Permit NAT Applied Unknown Message | Specifies how unidentified SIP 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 SIP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped.<br><br>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. |
| 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 MGCP ALG Call Duration and Timeouts

The call duration feature gives you control over SIP call activity and helps you to manage network resources.

**Before You Begin**

For background information, read “Understanding the MGCP ALG” on page 578.

Typically a call ends when one of the clients sends a BYE or CANCEL request. The MGPC 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 Juniper Networks 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) 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 SIP 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 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.

### *Related Topics*

- Understanding the MGCP ALG on page 578
- Setting MGCP Call Duration on page 587
- Setting MGCP Inactive Media Timeout on page 589
- Setting the MGCP Transaction Timeout on page 590
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 591
- Allowing Unknown MGCP Message Types on page 592

## 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 578
- Understanding MGCP ALG Call Duration and Timeouts on page 586

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 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.

Use either the J-Web or CLI configuration editor.

This topic covers:



### J-Web Configuration

1. Select Configuration > View and Edit > Edit Configuration > Security > ALG > MGCP.
2. Enter a value in the Maximum call duration box.
3. 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

## 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 578
- Understanding MGCP ALG Call Duration and Timeouts on page 586
- Setting the MGCP Transaction Timeout on page 590

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 SIP 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 589
- CLI Configuration on page 589
- Related Topics on page 590

### *J-Web Configuration*

1. Select Configuration > View and Edit > Edit Configuration > Security > ALG > MGCP.
2. Enter a value in the Inactive media timeout box.
3. 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 578
- Understanding MGCP ALG Call Duration and Timeouts on page 586
- Setting MGCP Call Duration on page 587
- Setting the MGCP Transaction Timeout on page 590
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 591
- Allowing Unknown MGCP Message Types on page 592

## 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 578.
- "Understanding MGCP ALG Call Duration and Timeouts" on page 586.
- "Setting MGCP Inactive Media Timeout" on page 589.

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*

1. Select Configuration > View and Edit > Edit Configuration > Security > ALG > MGCP.
2. Enter a value in the Transaction timeout box.
3. 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 578
- Understanding MGCP ALG Call Duration and Timeouts on page 586
- Setting MGCP Call Duration on page 587
- Setting MGCP Inactive Media Timeout on page 589
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 591
- Allowing Unknown MGCP Message Types on page 592

## 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 578.
- "Understanding MGCP ALG Call Duration and Timeouts" on page 586.

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

1. Select Configuration > View and Edit > Edit Configuration > Security > ALG > MGCP.
2. Click Application Screen.
3. Under Connection flood, enter a value in the Threshold box.
4. Under Message flood, enter a value in the Threshold box.
5. 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 sip application-screen message-flood threshold 10000
user@host# set security alg sip application-screen connection-flood threshold 4000
```

### Related Topics



## 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 SIP 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 578.

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*

1. Select Configuration > View and Edit > Edit Configuration > Security Edit > ALG Configure > MGCP Configure > Application screen configure > Unknown message Configure.
2. Click one of the following check boxes:
   - To allow unknown messge types in NAT mode, click the Permit nat applied Yes check box.
   - To allow unknown messge types in route mode, click the Permit routed Yes check box.
3. 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 578
- Understanding SIP ALG Call Duration and Timeouts on page 508
- Setting MGCP Call Duration on page 587
- Setting MGCP Inactive Media Timeout on page 589
- Configuring MGCP Denial of Service (DoS) Attack Protection on page 591

## 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 578. |

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 129 on page 595.

Use either the J-Web or CLI configuration editor.

**Figure 129: Media Gateway in Subscribers' Home**

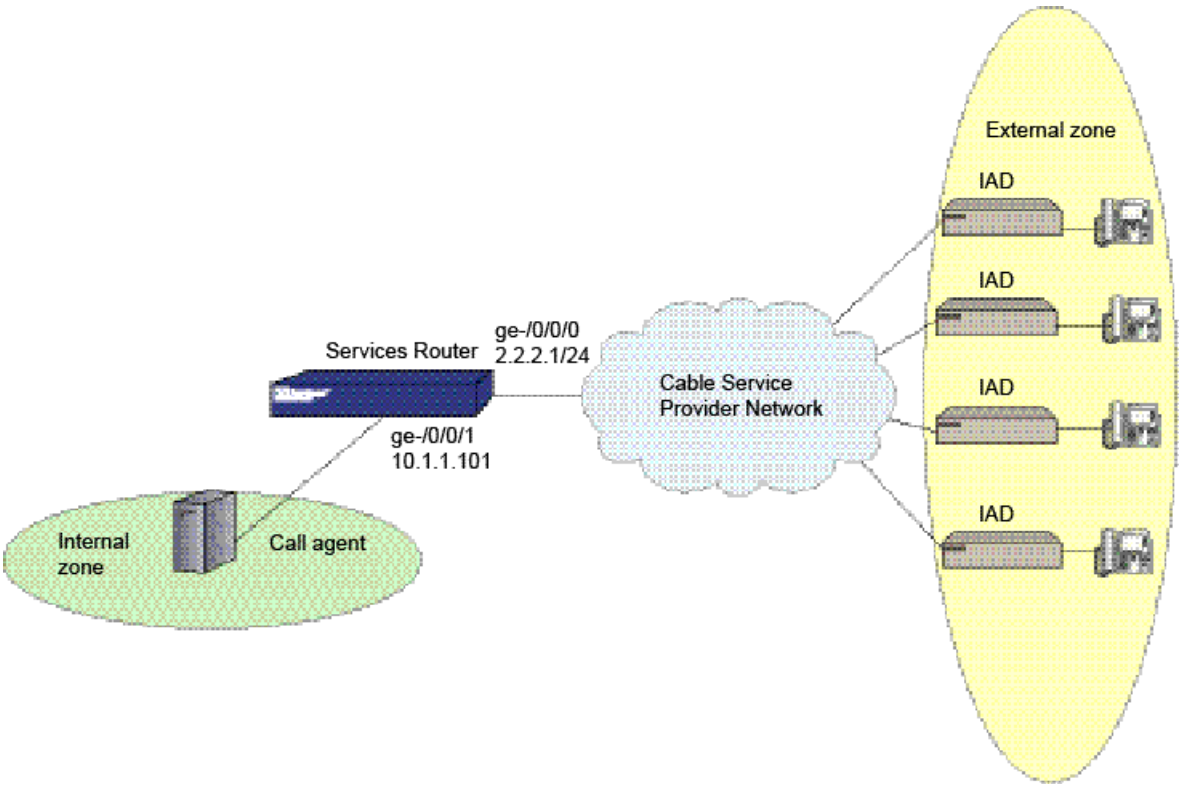


This topic covers:

- J-Web Configuration on page 595
- CLI Configuration on page 600
- Related Topics on page 601

## *J-Web Configuration*

To configure zones:

1. Select Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 Configuration > View and Edit > Edit Configuration. The Configuration page appears.
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 **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 578
- Configuring Three-Zone ISP-Hosted Service Using Source and Static NAT on page 601

## 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 578. |

In this example, (see Figure 130 on page 602) 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 85 on page 602, showing the corrsponding IP address, interface, and zone.

**Table 85: 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 130: Three-Zone ISP-Hosted Service**

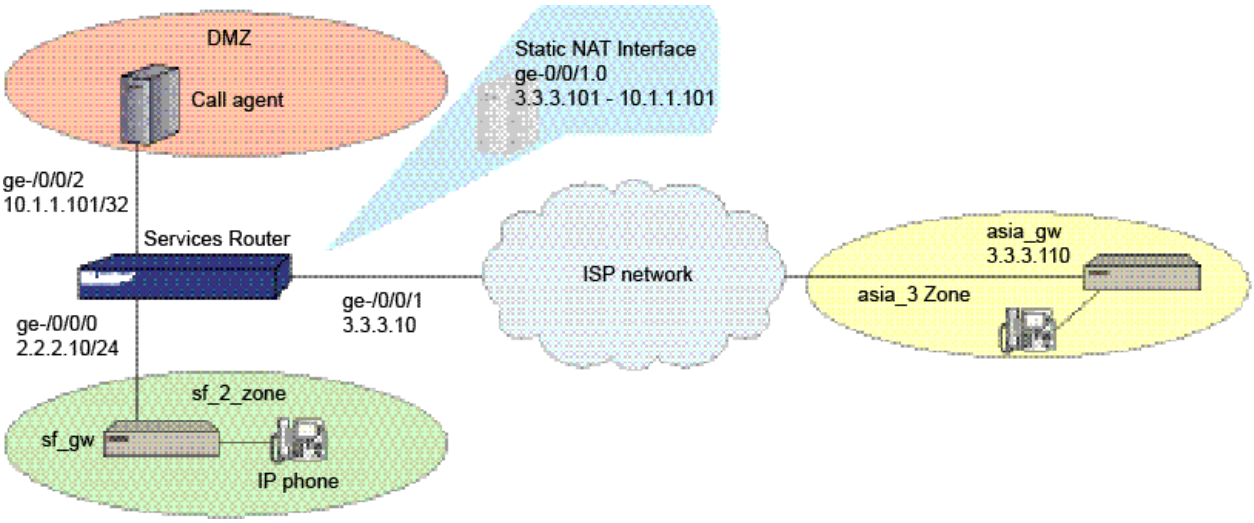


To configure a three-zone ISP-hosted service using source and static NAT, use either the J-Web or CLI configuration editor.

This topic covers:

- CLI Configuration on page 602
- Related Topics on page 605

## 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 578
- Configuring a Media Gateway in Subscribers' Homes on page 594

## Verifying the MGCP Configuration

To verify the MGCP configuration, perform these tasks:

- Verifying MGCP Calls on page 605
- Verifying MGCP Endpoints on page 606
- Verifying MGCP Counters on page 607

### *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 578
- Verifying MGCP Calls on page 605
- Verifying MGCP Endpoints on page 606
- Verifying MGCP Counters on page 607

### *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 578
- Verifying the MGCP ALG on page 605
- Verifying MGCP Endpoints on page 606
- Verifying MGCP Counters on page 607

## *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 578
- Verifying the MGCP ALG on page 605
- Verifying MGCP Calls on page 605
- Verifying MGCP Counters on page 607

### *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 578
- Verifying the MGCP ALG on page 605
- Verifying MGCP Calls on page 605
- Verifying MGCP Endpoints on page 606

## 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 471
- Understanding NAT on page 276

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 NFS or Microsoft Outlook to work behind the J-series device. The RPC ALG functionality is enabled by default.

This topic covers:

- Sun RPC ALG on page 608
- Microsoft RPC ALG on page 610
- Related Topics on page 611

### *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 609
- "Sun RPC Services" on page 609
- "CustomizingSun RPC Services" on page 610


## 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 86 on page 609 lists predefined Sun RPC services.

**Table 86: Predefined Sun RPC Services**

| | |
|---|---|
| junos-sun-rpc-portmap-tcp | junos-sun-rpc-portmap |

**Table 86: Predefined Sun RPC Services** *(continued)*

| |
|---|
| junos-sun-rpc-portmap-udp |

### CustomizingSun RPC Services

Because Sun RPC services use dynamically negotiated ports, you cannot use regular fixed TCP/UDP ports to permit Sun RPC services in a security policy. Instead, you must specify a Sun RPC program number. For example, NFS uses two program numbers: 100003 and 100227. The corresponding TCP/UDP ports are dynamic. To permit the program numbers, you use a set applications *application-name* term *term-name* rpc-program-number *number* statement for each number. The ALG maps the program numbers into dynamically negotiated TCP/UDP ports and permits or denies the service based on a policy you configure.

## *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 608), 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 device 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 Services 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 87 on page 611 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 87: 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 499
- Understanding the SCCP ALG on page 561
- Understanding the MGCP ALG on page 578
- Disabling and Enabling RPC ALG on page 611
- Verifying the RPC ALG Tables on page 612

## 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 608. |

This topic covers:

- J-Web Configuration on page 611
- CLI Configuration on page 612
- Related Topics on page 612

### *J-Web Configuration*

The Sun RPC ALG and MS 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 **Configuration > View and Edit > Edit Configuration > Security > ALG**

To disable or re-enable the Mirosoft RPC ALG:

1. Select **Configuration > View and Edit > Edit Configuration > Security > ALG**

### CLI Configuration

The Sun RPC ALG and MS 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 Mirosoft 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 608
- Verifying the RPC ALG Tables on page 612

## Verifying the RPC ALG Tables

To verify the RPC ALG, perform these tasks:

- Display the MS UUID mapping table

### Display the Sun RPC Port Mapping Table

**Purpose** Display the Sun RPC port map 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 MS RPC UUID Mapping Table*

**Purpose** Display the MS RPC UUID-to-Object ID (OID) map table. The MS RPC ALG monitors packets on TCP port 135. (You can display the MS RPC port map 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:

- MS RPC UUID
- MS RPC object ID

### *Related Topics*

- Understanding the RPC ALG on page 608
- Disabling and Enabling RPC ALG on page 611

Chapter 17

# 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.

This section includes:



## System Requirements for NetScreen-Remote Client Installation

You can install the NetScreen-Remote client in the environment described in Table 88 on page 615.

**Table 88: System Requirements**

| | |
|---|---|
| PC-compatible computer | Pentium processor or its equivalent |
| Operating system | Microsoft Windows 2000 Professional<br>Microsoft Windows XP Professional or Home Edition. |
| Minimum RAM | 64 MB RAM for Windows 2000 or Windows XP. |

**Table 88: System Requirements** *(continued)*

| | |
|---|---|
| Available hard disk space | Minimum 5 MB, maximum 35 MB |
| Software installation | CD-ROM drive, network drive, or Web site |
| Communications protocol | IPsec and IKE L2TP with Windows 2000 (*optional*)<br>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<br>Native Microsoft Dial-up Networking<br>PPPoE drivers<br>Compatible with America Online (AOL) 6.0 or later |
| Network connections | Ethernet<br>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 616
- Completing NetScreen-Remote Client Installation on page 618

### *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 617
- Starting Installation from a Network Share Drive on page 617
- Starting Installation from a Web Site on page 617

### Starting Installation from a CD-ROM

Figure 131 on page 617 shows the Install menu.

**Figure 131: Select Install**

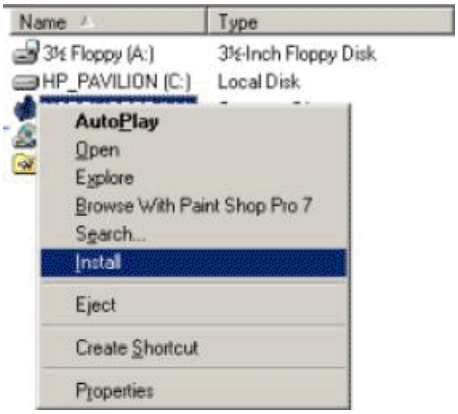


To install the NetScreen-Remote Client from a CD-ROM:

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 131 on page 617 appears.
5. To install the NetScreen-Remote client, select **Install** from the menu.
6. Go on to "Completing NetScreen-Remote Client Installation" on page 618.

### 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 618.

### Starting Installation from a Web Site

To install the NetScreen-Remote Client from a Web site:

1. Locate the NetScreen-Remote client files on the Juniper Networks Web site 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 618.

### *Completing NetScreen-Remote Client Installation*

Figure 132 on page 618 shows ther Welcome Screen page.

**Figure 132: Welcome Screen**

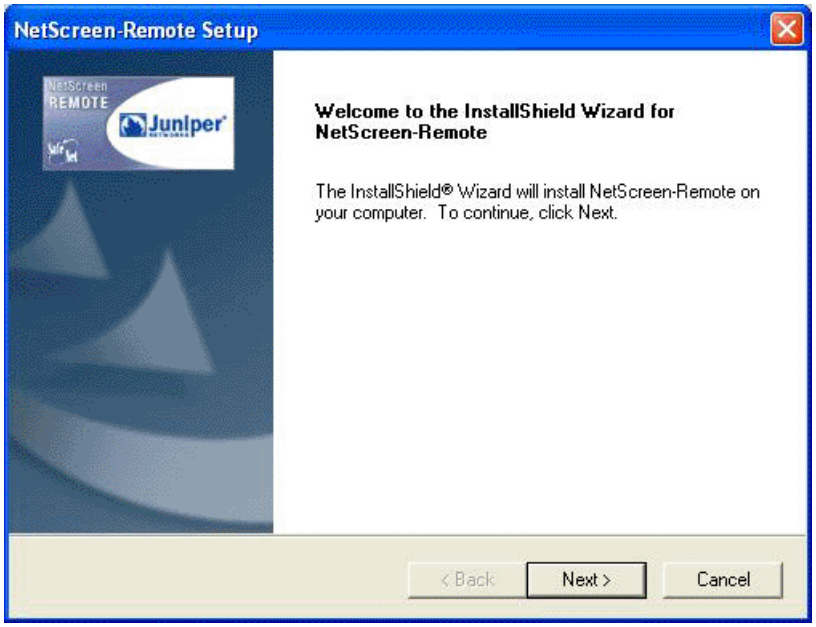


Figure 133 on page 619 shows the License Agreement page.

**Figure 133: License Agreement**

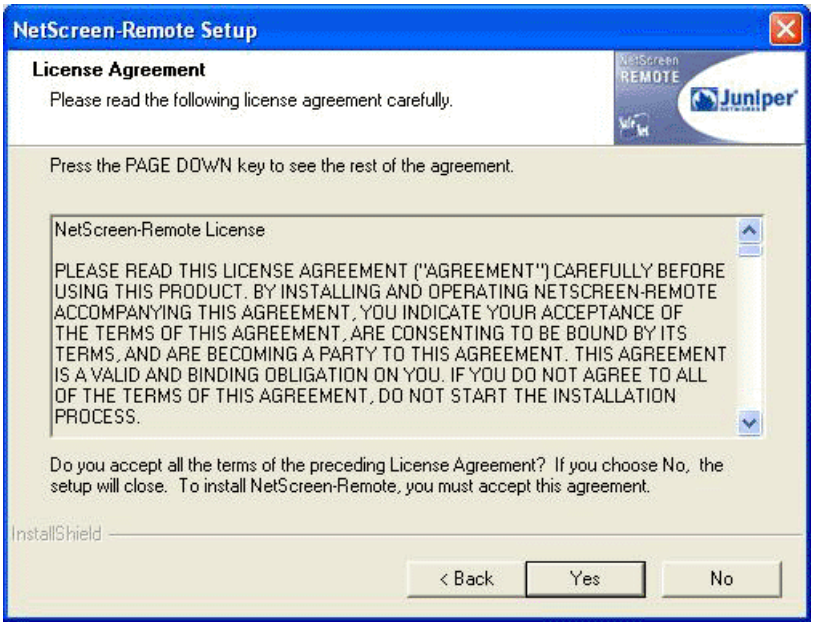


Figure 134 on page 619 shows the Installation Setup Type page.

**Figure 134: Installation Setup Type**

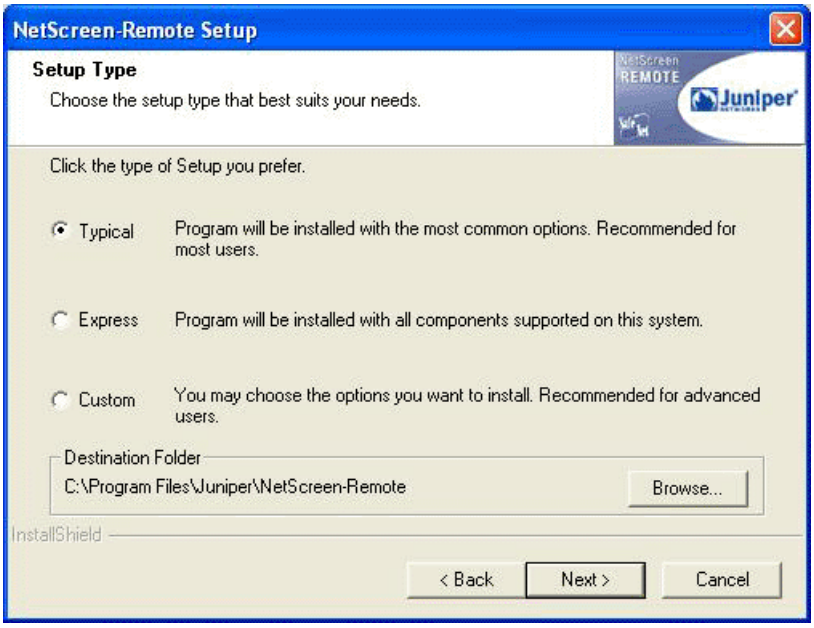


Figure 135 on page 620 shows the Start Copying Files page.

**Figure 135: Start Copying Files**

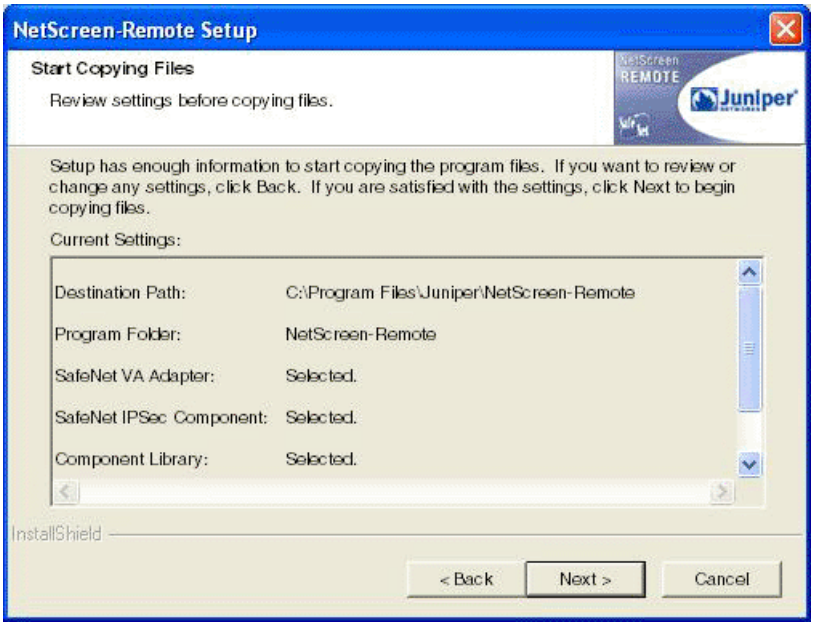


After you have finished initial installation (see "Starting NetScreen-Remote Client Installation" on page 616), complete the NetScreen-Remote client installation:

1. When the InstallShield Wizard starts as shown in Figure 132 on page 618, click **Next**.
2. On the Software License Agreement page (Figure 133 on page 619), read the agreement and click **Yes** to continue.
3. In the Setup Type dialog box (Figure 134 on page 619), 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 135 on page 620), 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 J-series 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 616. |

This topic covers:



### *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 621.
2. Configure a tunnel interface to identify the VPN. For more information, see "Configuring a Tunnel Interface" on page 622.
3. Configure the eXtended Authentication (XAuth) settings. For more information, see "Configuring an Access Profile for XAuth" on page 623.
4. Configure an IKE gateway and phase 1 proposal. For more information, see "Configuring an IKE Gateway" on page 623.
5. Configure firewall policies to permit or deny traffic. For more information, see "Configuring Policies" on page 624.

### *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 179 and "Security Policies Overview" on page 69.

To configure a security zone:

1. Select Configuration > Quick Configuration > 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 53.

## 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 621. Then you can bind the VPN to the tunnel interface in the VPN configuration.

Figure 136 on page 622 shows the Interface List page.

**Figure 136: Interface List**

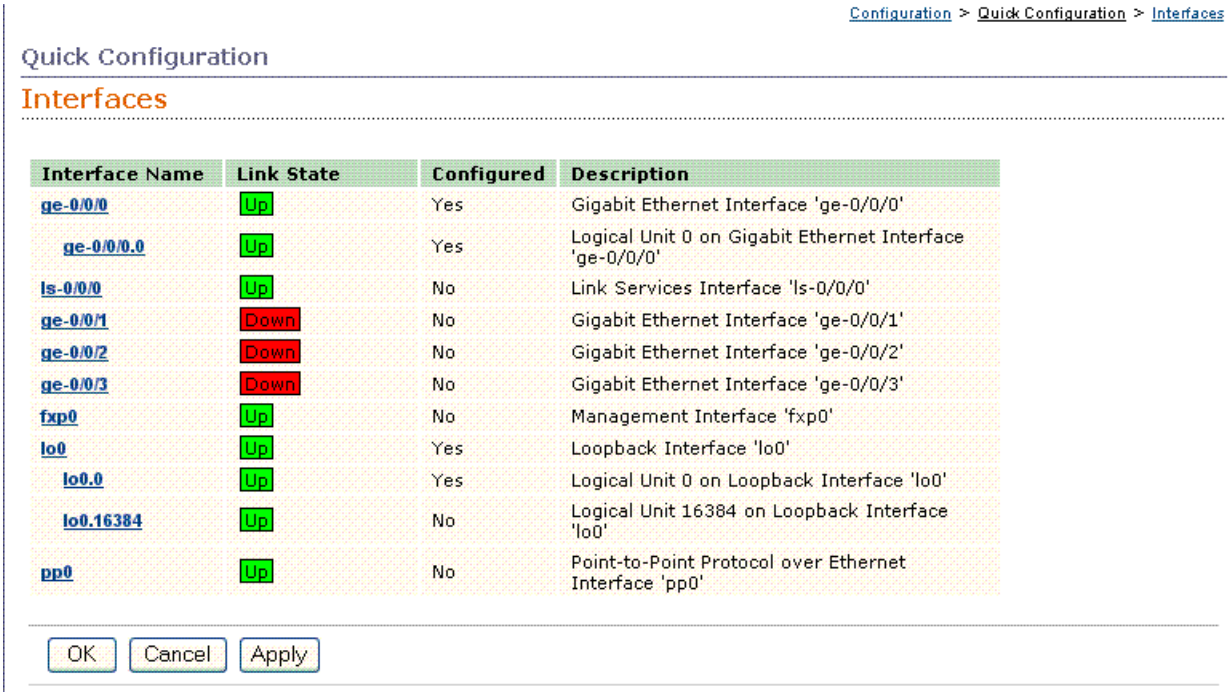


To configure a tunnel interface to be used for the VPN:

1. Select Configuration > Quick Configuration > Interfaces. A list of the network interfaces available on the router appears, as shown in Figure 136 on page 622. 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 65.

   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. See "Configuring an IKE Phase 1 Proposal—Quick Configuration" on page 401.
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. See "Configuring an IKE Policy, Authentication, and Proposal—Quick Configuration" on page 406.
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 email 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 623)

   ```
   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
   ```

Alternatively, for J-Web Quick Configuration, see “Configuring an IKE Gateway and Peer Authentication—Quick Configuration” on page 411.

**NOTE:** Make sure you select the access profile that you set using the CLI statements in “Configuring an Access Profile for XAuth” on page 623.

### *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 Configuration > Quick Configuration > Security Policies > 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 53.

## 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 621. |

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 625
2. “Creating the Preshared Key” on page 628
3. “Defining the IPsec Protocols” on page 629

## *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 137 on page 625 shows the NetScreen-Remote Client Icon in the Task Bar page.

**Figure 137: NetScreen-Remote Client Icon**

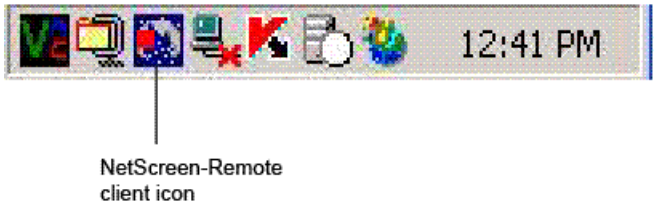


Figure 138 on page 625 shows the Security Policy Editor page.

**Figure 138: Security Policy Editor**

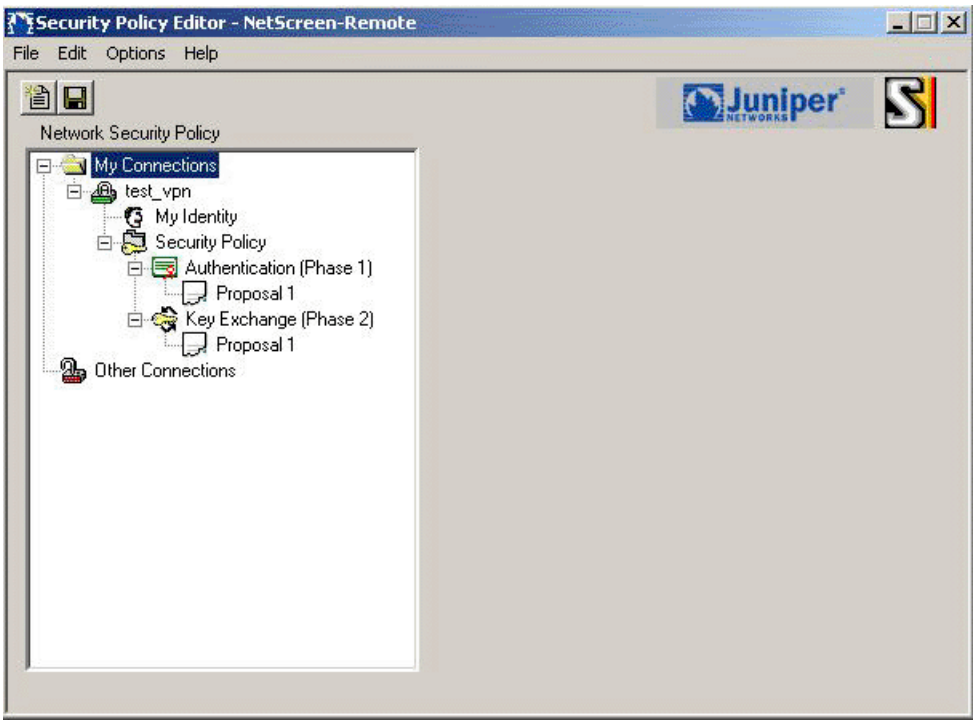


New Connection appears in the Network Security Policy list, as shown in Figure 139 on page 626.

**Figure 139: Configure Connection**

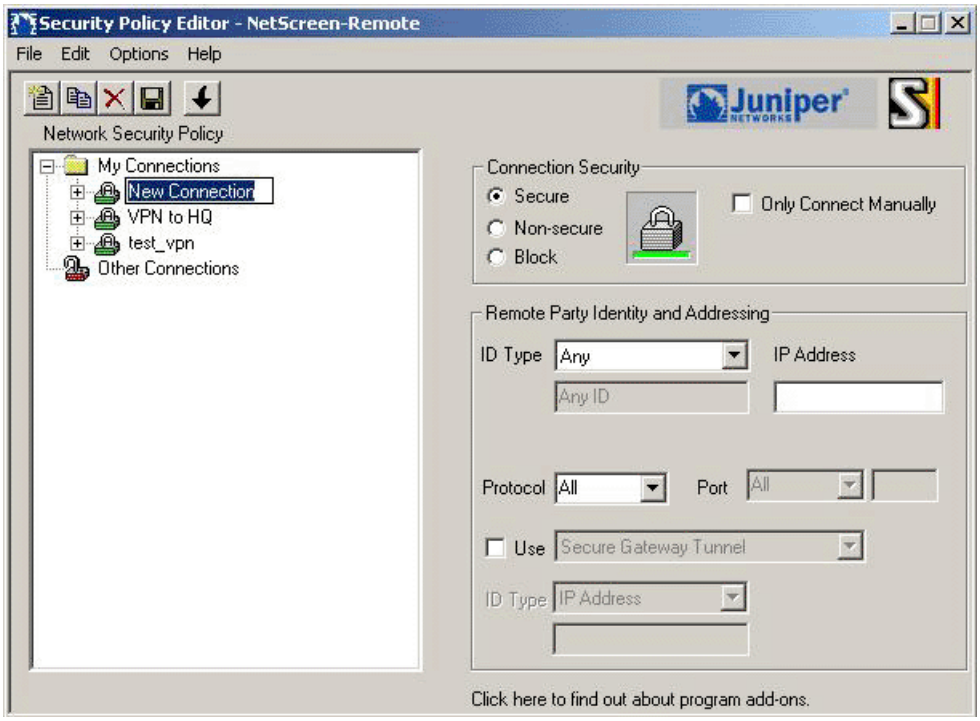


Figure 140 on page 627 shows the Configuring the New Connection page.

**Figure 140: Configuring the New Connection**

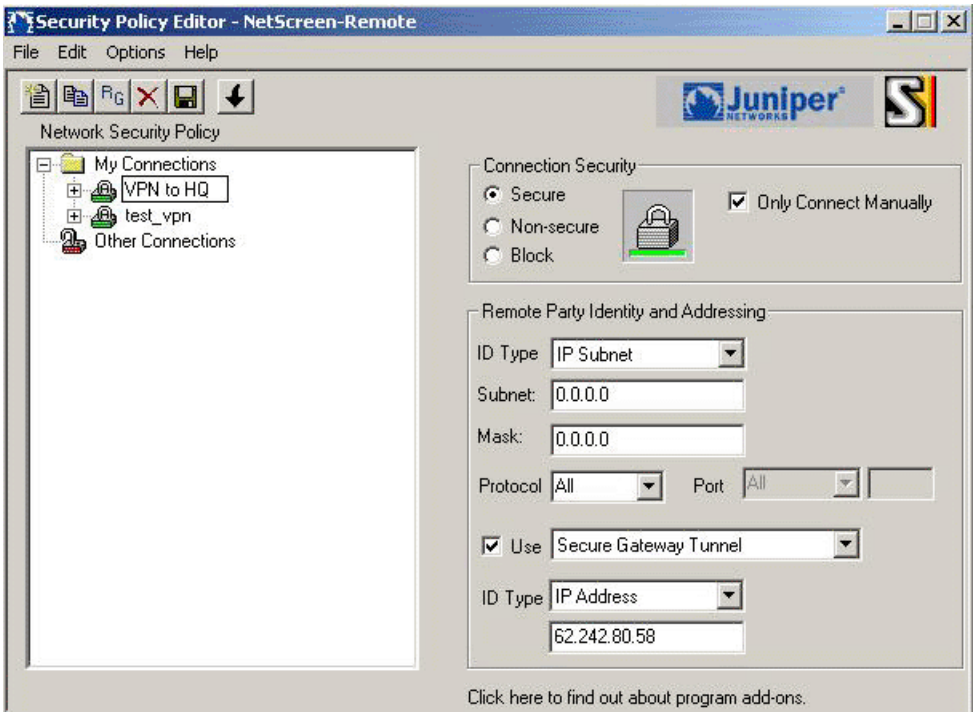


1. Double-click the NetScreen-Remote icon in the Windows taskbar (Figure 137 on page 625). The Security Policy Editor screen appears (Figure 138 on page 625).
2. Click the New Connection icon to create a new connection.
3. Give the new connection a unique name—for example, VPN to HQ.
4. In the Connection 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 router 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 140 on page 627.

## *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 141 on page 628 shows the My Identity and Internet Interface page.

**Figure 141: My Identity and Internet Interface**

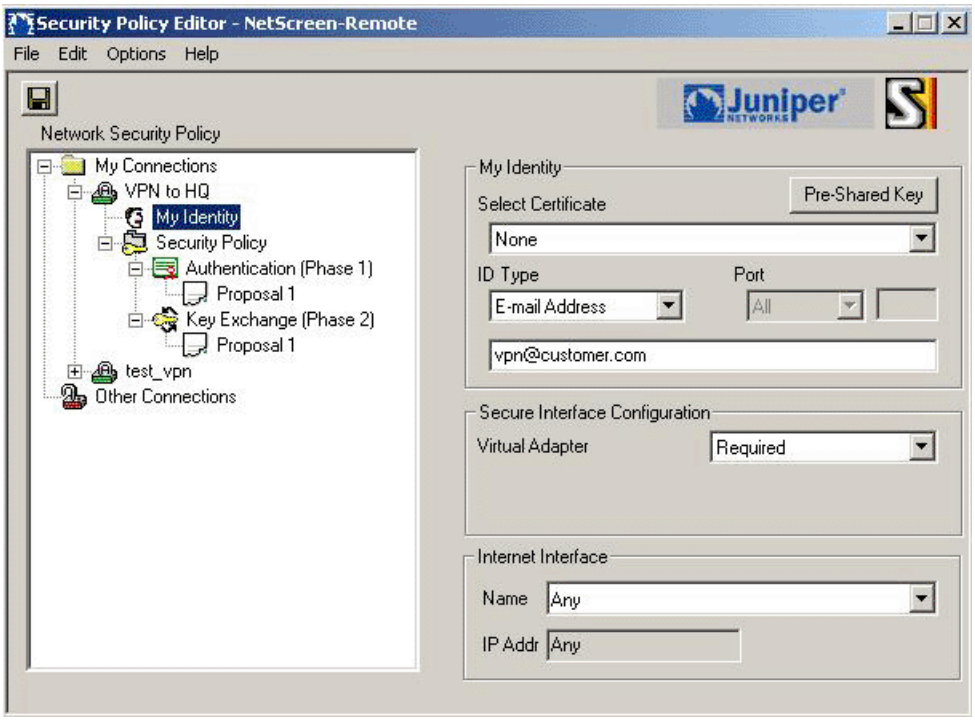


Figure 142 on page 629 shows the Per-Shared Key Dialog Box page.

**Figure 142: 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 141 on page 628).
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 142 on page 629).
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 143 on page 630.

**Figure 143: Security Policy**

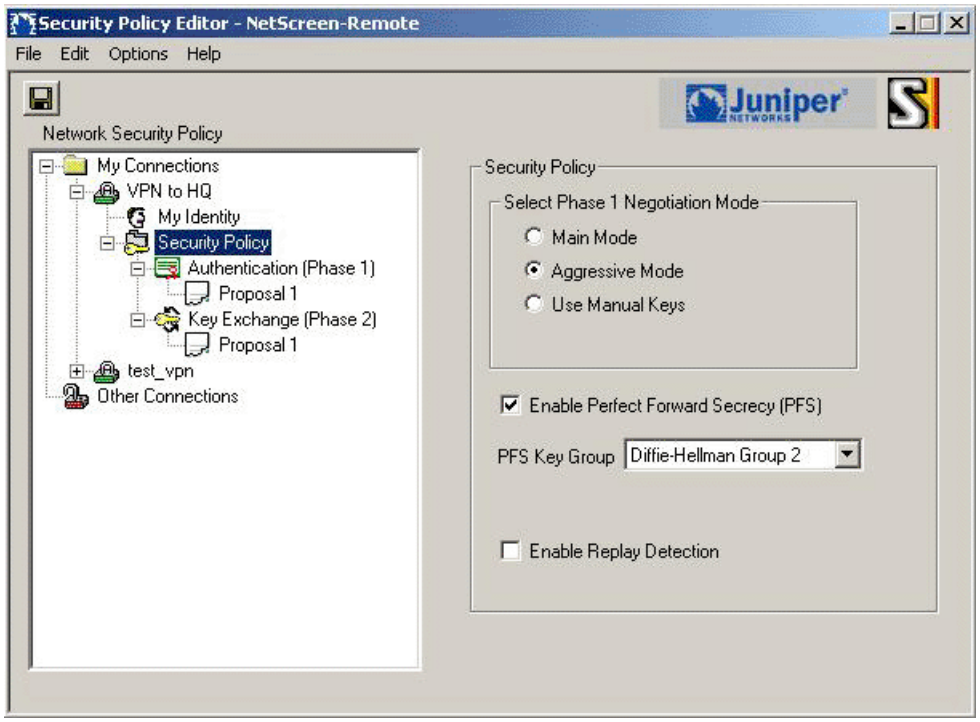


Figure 144 on page 631 shows the Algorithims Area page.

**Figure 144: Algorithms Area**

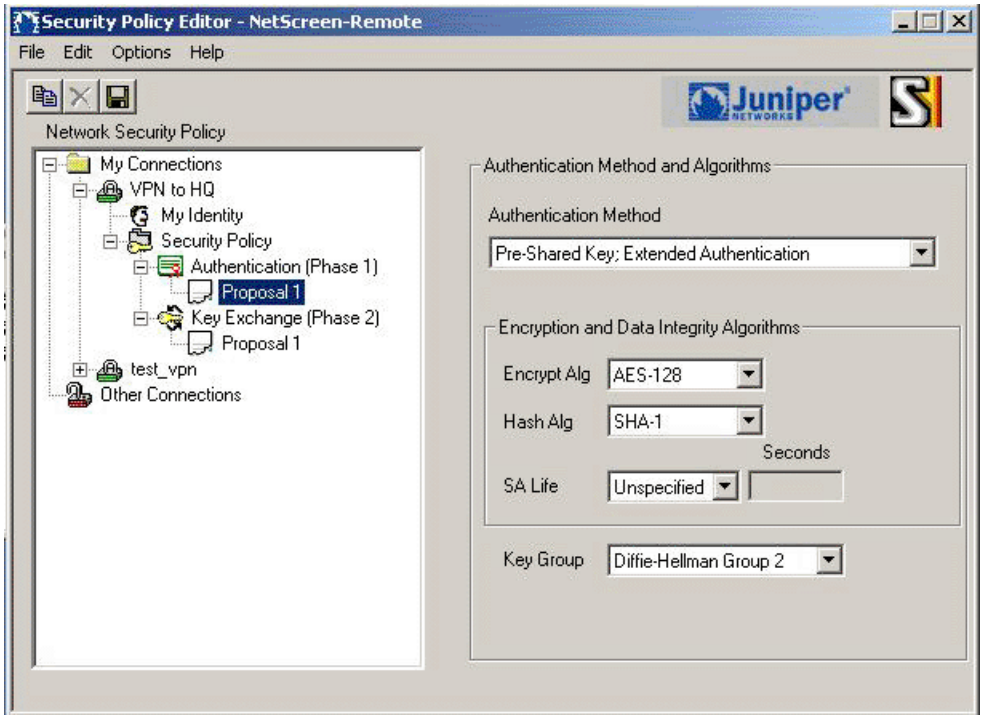

**Figure 145: IPsec Protocols Area**

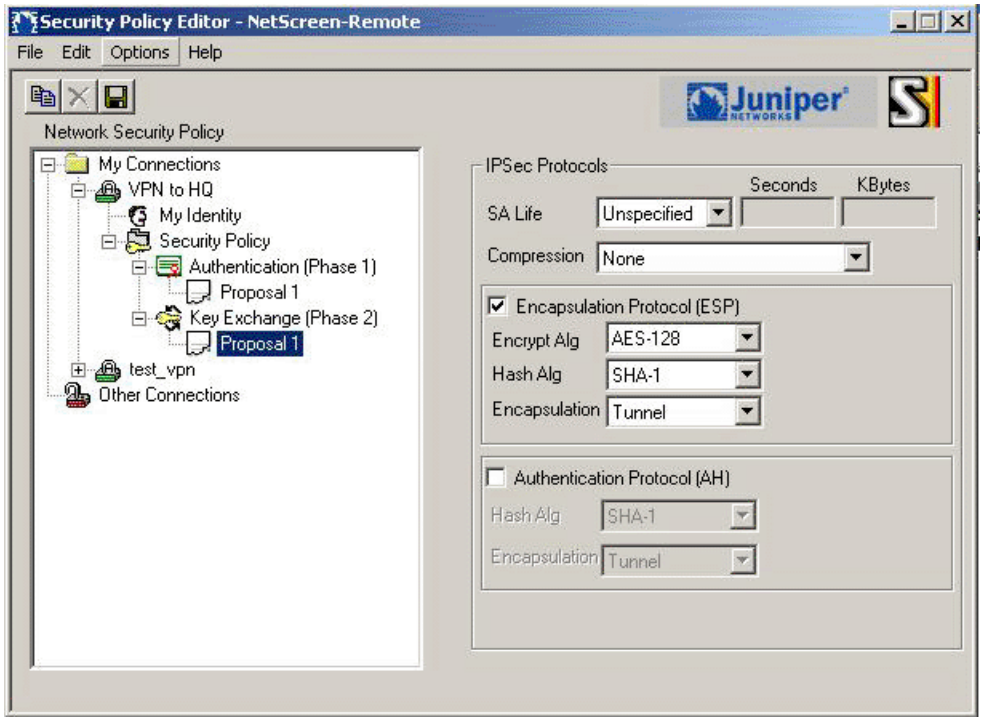


To define the Internet Protocol Security (IPsec) protocols for securing the VPN tunnel:

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 141 on page 628.
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 router 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 623 and “Configuring an IKE Gateway” on page 623 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 89 on page 633 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 145 on page 632.
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 89: 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 624. |

Figure 146 on page 634 shows the Login page.

**Figure 146: Login**


Figure 147 on page 634 shows the Successfully Connected page.

**Figure 147: 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 146 on page 634).
2. Type the username and password configured for the XAuth user. See "Configuring an Access Profile for XAuth" on page 623 and "Configuring an IKE Gateway" on page 623 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 147 on page 634 appears.
4. Click OK.

Part 3

# Intrusion Detection and Prevention Features

Chapter 18

# 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.

This topic covers:

## 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 90 on page 640.

**Table 90: 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. |
| **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 652).
- Create an IDP policy starting with one of the predefined templates provided by Juniper Networks (see "Using Predefined Policy Templates" on page 706).
- 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 697).
- 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 19
- IDP Policies Overview on page 640

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:



## 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 91 on page 642 summarizes the options that you can configure in the IPS-rulebase rules.

**Table 91: 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 644. |
| **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 645. |
| **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 677. |
| **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 647. |
| **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 649. |
| **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 651. |

## 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 92 on page 643 summarizes the options that you can configure in the exempt-rulebase rules.

**Table 92: 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 643
- Defining Rules for an IPS Rulebase on page 652
- Defining Rules for an Exempt Rulebase on page 656

## 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 640
- Understanding IDP Policy Rulebases on page 641

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 644)
- Attack Objects and Groups (see “Understanding IDP Rule Objects” on page 645)
- Actions (see “Understanding IDP Rule Actions” on page 647)
- IP Actions (see “Understanding IDP Rule IP Actions” on page 649)
- Notifications (see “Understanding IDP Rule Notifications” on page 651)

### *Related Topics*

- Defining Rules for an IPS Rulebase on page 652
- Defining Rules for an Exempt Rulebase on page 656
- Using Predefined Policy Templates on page 706
- Inserting a Rule in the Rulebase on page 667
- Deactivating and Reactivating Rules in a Rulebase on page 669

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 645
- Understanding IDP Rule Actions on page 647
- Understanding IDP Rule IP Actions on page 649
- Understanding IDP Rule Notifications on page 651

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 four 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 644
- Understanding IDP Rule Actions on page 647
- Understanding IDP Rule IP Actions on page 649
- Understanding IDP Rule Notifications on page 651

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

Table 93 on page 648 shows the actions you can specify for IDP rules:

**Table 93: 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.<br><br>**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.<br><br>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. |
| **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 93: 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.<br><br>**NOTE:** This action is supported only for IPS rulebases.<br><br>Recommended —A list of all attack objects that Juniper Networks considers to be serious threats, organized into categories.<br><br>■ Attack type groups attack objects by type (anomaly or signature). Within each type, attack objects are grouped by severity.<br>■ Category groups attack objects by predefined categories. Within each category, attack objects are grouped by severity.<br>■ 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.<br>■ 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 644
- Understanding IDP Rule Objects on page 645
- Understanding IDP Rule IP Actions on page 649
- Understanding IDP Rule Notifications on page 651

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 94 on page 650 summarizes the types IP actions supported by IDP rules:

**Table 94: 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 644
- Understanding IDP Rule Objects on page 645
- Understanding IDP Rule Actions on page 647
- Understanding IDP Rule Notifications on page 651

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 644
- Understanding IDP Rule Objects on page 645
- Understanding IDP Rule Actions on page 647
- Understanding IDP Rule IP Actions on page 649

## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Policy Rules on page 643
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 51.
5. Enable IDP in security policies. See "Enabling IDP in a Security Policy" on page 673.

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*

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 **Configuration > View and Edit > Edit Configuration**. The Configuration page appears.
   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 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 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 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 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 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 656
- Using Predefined Policy Templates on page 706
- Inserting a Rule in the Rulebase on page 667
- Deactivating and Reactivating Rules in a Rulebase on page 669

## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Policy Rules on page 643
2. Create rules in the IPS rulebase. See “Defining Rules for an IPS Rulebase” on page 652.

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*

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 **Configuration > View and Edit > Edit Configuration**. The Configuration page appears.
   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**.
4. 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**.

5. 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.

6. 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 706
- Inserting a Rule in the Rulebase on page 667
- Deactivating and Reactivating Rules in a Rulebase on page 669

## IDP Policies—Quick Configuration

This topic contains:



### *Configuring IDP Policies—Quick Configuration*

You can use J-Web Quick Configuration to quickly configure an IDP policy.

**Figure 148: Quick Configuration page for IDP Policies**


To configure an IDP policy with Quick Configuration:

1. Select Configuration > Quick Configuration > Security Policies > IDP.

   Figure 148 on page 659 shows the Quick Configuration page for IDP policies.
2. Fill in the information as described in Table 95 on page 660.
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 95: 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.<br><br>For more information, see "Configuring a Security Package Update—Quick Configuration" on page 714. |
| 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.<br><br>For more information, see "Adding a New IDP Policy—Quick Configuration" on page 660. |
| 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 149: Quick Configuration Page for Adding a New IDP Policy**

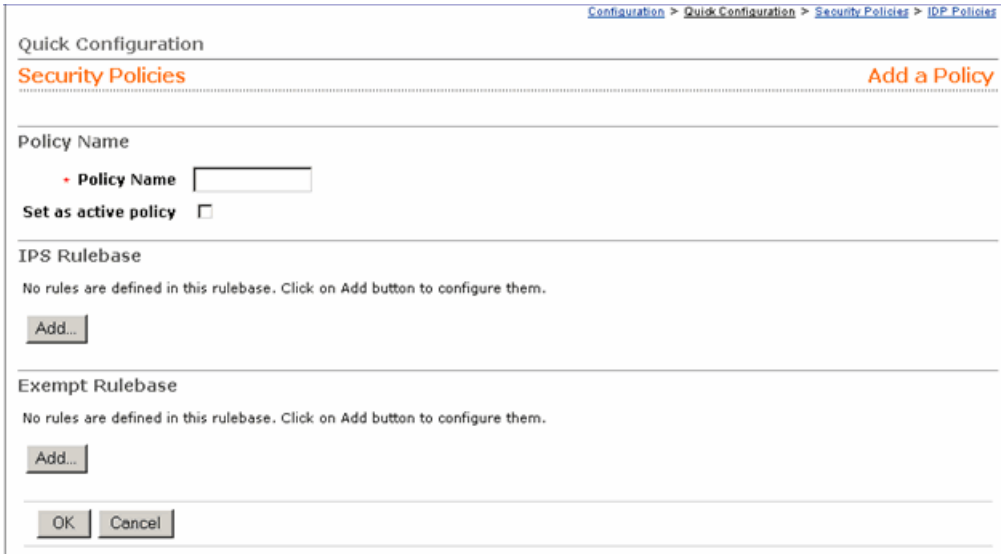


To configure a new IDP policy with Quick Configuration:

1. Select Configuration > Quick Configuration > Security Policies > IDP.
2. From the IDP Policies page, click Add to add a new IDP policy.

   Figure 149 on page 661 shows the Quick Configuration page for adding a new IDP policy.
3. Fill in the information as described in Table 96 on page 661.
4. 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 96: 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. |

**Table 96: Adding a New IDP Policy Quick Configuration Page Summary** *(continued)*

| Field | Function | Actions |
|---|---|---|
| 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.<br><br>For more information, see "Adding an IPS Rulebase—Quick Configuration" on page 662. |
| 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.<br><br>For more information, see "Adding an Exempt Rulebase—Quick Configuration" on page 665. |

## *Adding an IPS Rulebase—Quick Configuration*

You can use J-Web Quick Configuration to quickly configure and add an IPS rulebase.

**Figure 150: Quick Configuration Page for Adding an IPS Rulebase**


To configure an IPS rulebase with Quick Configuration:

1. Select Configuration > Quick Configuration > Security Policies > IDP.
2. In the Policy Name text box, type a policy name.
3. Under IPS Rulebase, click Add to add an IPS rulebase.

   Figure 150 on page 662 shows the Quick Configuration page for IPS rulebase.
4. Fill in the information as described in Table 97 on page 663.
5. 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 97: 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:<br>■ To match the from-zone and source address/address sets to the rule, click the left arrow.<br>■ 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:<br>■ To match the to-zone and destination addresses/address sets to the rule, click the left arrow.<br>■ 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. |

**Table 97: Adding an IPS Rulebase Quick Configuration Page Summary** *(continued)*

| Field | Function | Actions |
|---|---|---|
| **Applications and Application Sets** | | |
| Matched | Specifies the type of network traffic you want the device to monitor for attacks. | |
| 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:<br>■ To match the rule to the applications/application sets, click the left arrow.<br>■ 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. | Type a valid predefined attack name and do either one the following:<br>■ To add a predefined attack, type it next to the Add button, and click Add.<br>■ To remove a predefined attack, select it in the Predefined Attacks box, and click Delete. |
| Predefined Attack Groups | Specifies predefined attack groups that are used to match the traffic against known attack objects. | Enter a valid predefined attack group name and do either one the following:<br>■ To add a predefined attack group, type it next to the Add button, and click Add.<br>■ To remove a predefined attack group, select it in the Predefined Attack groups box, and click Delete. |
| 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:<br>■ To match a custom attack to the rule, click the left arrow.<br>■ 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 97: 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. |
| 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 151: Quick Configuration Page for Adding an Exempt Rulebase**

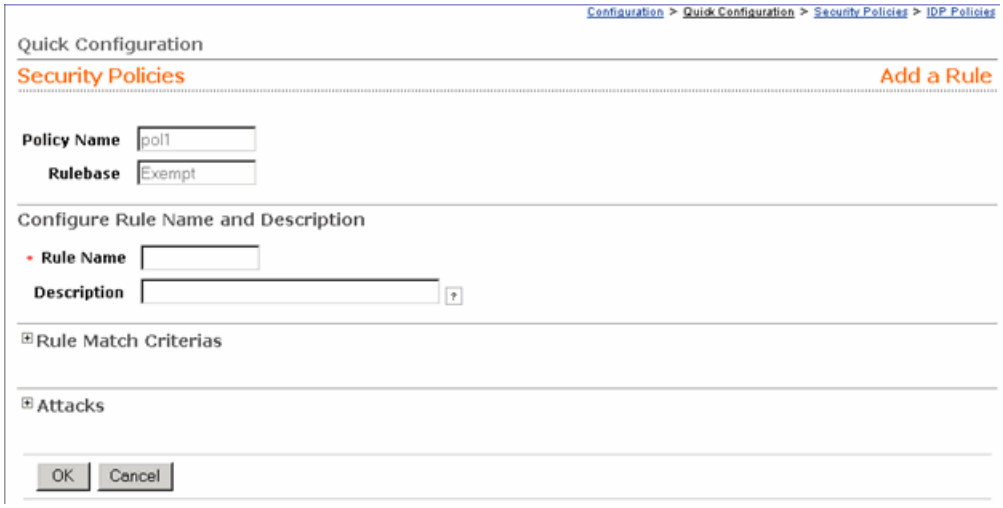


To configure an exempt rulebase with Quick Configuration:

1. Select Configuration > Quick Configuration > Security Policies > IDP.
2. In the Policy Name text box, type a policy name.

3. Under Exempt Rulebase, click Add to add an exempt rulebase.

   Figure 151 on page 665 shows the Quick Configuration page for exempt rulebase.

4. Fill in the information as described in Table 98 on page 666.
5. 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 98: 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:<br>■ To match the from-zone and source address/address sets to the rule, click the left arrow.<br>■ 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 98: 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:<br>■ To match the to-zone and destination addresses/address sets to the rule, click the left arrow.<br>■ 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. | Type a valid predefined attack name and do either one the following:<br>■ To add a predefined attack, type it next to the Add button, and click Add.<br>■ To remove a predefined attack, select it in the Predefined Attacks box, and click Delete. |
| Predefined Attack Groups | Specifies predefined attack groups that are used to match the traffic against known attack objects. | Enter a valid predefined attack group name and do either one the following:<br>■ To add a predefined attack group, type it next to the Add button, and click Add.<br>■ To remove a predefined attack group, select it in the Predefined Attack groups box, and click Delete. |
| 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:<br>■ To match a custom attack to the rule, click the left arrow.<br>■ 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Policy Rules on page 643
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 652.

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

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

## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Policy Rules on page 643
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 652.

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*

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 656
- Using Predefined Policy Templates on page 706
- Inserting a Rule in the Rulebase on page 667

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 672
- Defining Rules for an IPS Rulebase on page 652
- Defining Rules for an Exempt Rulebase on page 656

## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Rule Objects on page 645
   - Configuring Applications or Services for IDP on page 670
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 673.

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

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 672
- Defining Rules for an IPS Rulebase on page 652
- Using Predefined Policy Templates on page 706

## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding Application Sets on page 670
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 673.

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*

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*



## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Policy Rules on page 643
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 51.
5. Configure applications. See "Configuring Applications or Services for IDP" on page 670.

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

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 **Configuration > View and Edit > Edit Configuration**. The Configuration page appears.
   b. Next to Security, click Configure or Edit.

c. Next to Policies, select the check box and click Configure.

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 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.

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 652
- Defining Rules for an Exempt Rulebase on page 656

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 678
- Defining Rules for an IPS Rulebase on page 652
- Defining Rules for an Exempt Rulebase on page 656

## 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Terminal Rules on page 677
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 673.
5. Create security zones. See "Creating Security Zones" on page 51.
6. Define rules. See "Defining Rules for an IPS Rulebase" on page 652.

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

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 **Configuration > View and Edit > Edit Configuration**. The Configuration page appears.
   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 656
- Using Predefined Policy Templates on page 706
- Inserting a Rule in the Rulebase on page 667
- Deactivating and Reactivating Rules in a Rulebase on page 669

## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding Predefined Attack Objects and Groups on page 709

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 682
- Severity on page 682
- Service or Application Binding on page 682
- Protocol or Port Bindings on page 686
- Time Bindings on page 687

- Attack Properties—Signature Attacks on page 688
- Attack Properties—Protocol Anomaly Attacks on page 694
- Attack Properties—Compound or Chain Attacks on page 695
- Related Topics on page 697

### *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 651). 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 99 on page 682 displays supported services and default ports associated with the services.

**Table 99: 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 99: 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 99: 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 email. | 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 99: 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 100 on page 686 lists protocol numbers for different protocols.

**Table 100: 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 |
| COMP | 108 |

**Table 100: Supported Protocols and Protocol Numbers** *(continued)*

| Protocol Name | Protocol Number |
|---|---|
| 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 101 on page 687 displays sample formats for key protocols.

**Table 101: 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><br>■ <Port>TCP/*port* </Port><br>■ <Port>TCP/*minport-maxport* </Port> | Specifying the port is optional for TCP and UDP protocols. For example, you can specify either of the following:<br>■ <Port>UDP</Port><br>■ <Port>UDP/10</Port><br>■ <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 102 on page 691 displays fields and flags that you can set for attacks that use the IP protocol.

**Table 102: IP Protocol Fields and Flags**

| Field | Description |
|---|---|
| **Type of Service** | Specify a value for the service type. Common service types are:<br>■ 0000 Default<br>■ 0001 Minimize Cost<br>■ 0002 Maximize Reliability<br>■ 0003 Maximize Throughput<br>■ 0004 Minimize Delay<br>■ 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 103 on page 692 displays packet header fields and flags that you can set for attacks that use the TCP protocol.

**Table 103: 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 104 on page 693 displays packet header fields and flags that you can set for attacks that use the UDP protocol.

**Table 104: 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 105 on page 693 displays packet header fields and flags that you can set for attacks that use the ICMP protocol.

**Table 105: 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 697
- Defining Rules for an IPS Rulebase on page 652

## 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 640
   - Understanding Custom Attack Objects on page 681
   - Understanding Predefined Attack Objects and Groups on page 709
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 the 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 652.

You can use either J-Web or the CLI configuration editor to create a custom attack object.

This topic contains:



### *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 652
- Configuring Protocol Anomaly-Based Attacks on page 700

## 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 640
   - Understanding Custom Attack Objects on page 681
   - Understanding Predefined Attack Objects and Groups on page 709
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 652.

You can use either J-Web or the CLI configuration editor to create a custom attack object.

This topic contains:



## *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 652
- Configuring Signature-Based Attacks on page 697

## 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 640
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 673.
5. Create security zones. See "Creating Security Zones" on page 51.
6. Define rules. See "Defining Rules for an IPS Rulebase" on page 652.

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*

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 656
- Using Predefined Policy Templates on page 706

Chapter 19

# 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 Web site. You can download this file to protect your network from new threats.

This topic covers:



## 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 640
- Understanding IDP Policy Rulebases on page 641
- Understanding IDP Policy Rules on page 643

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 Web site. 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 652
- Understanding the Signature Database Version on page 717
- Updating the Signature Database Manually on page 712
- Updating the Signature Database Automatically on page 716

## 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 Web

site. To start using a template, you run 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 640
   - Understanding IDP Policy Rulebases on page 641
   - Understanding IDP Policy Rules on page 643
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 106 on page 707 summarizes the predefined IDP policy templates provided by Juniper Networks.

**Table 106: 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 Web site.
- 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 718).

- 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*

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 652
- Defining Rules for an Exempt Rulebase on page 656

## 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 640
- Understanding IDP Policy Rules on page 643
- Understanding IDP Rule Objects on page 645

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 712
- Defining Rules for an IPS Rulebase on page 652
- Defining Rules for an Exempt Rulebase on page 656

## Updating the Signature Database Overview

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks Web site. 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 640
- Understanding the IDP Signature Database on page 705
- Understanding IDP Policy Rules on page 643
- Understanding Predefined Attack Objects and Groups on page 709

To update the signature database, you download a security package from the Juniper Networks Web site. 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 706.)

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.

### *Related Topics*

- Understanding the Signature Database Version on page 717
- Updating the Signature Database Automatically on page 716
- Updating the Signature Database Manually on page 712
- Configuring a Security Package Update—Quick Configuration on page 714

## Updating the Signature Database Manually

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks Web site. 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 711
   - IDP Policies Overview on page 640
   - Understanding the IDP Signature Database on page 705
   - Understanding Predefined Attack Objects and Groups on page 709
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

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 717
- Updating the Signature Database Automatically on page 716
- Updating the Signature Database Manually on page 712
- Configuring a Security Package Update—Quick Configuration on page 714

## Configuring a Security Package Update—Quick Configuration

You can use J-Web Quick Configuration to quickly configure the Security Package Update.

**Figure 152: Quick Configuration Page for Configuring Security Package Update**

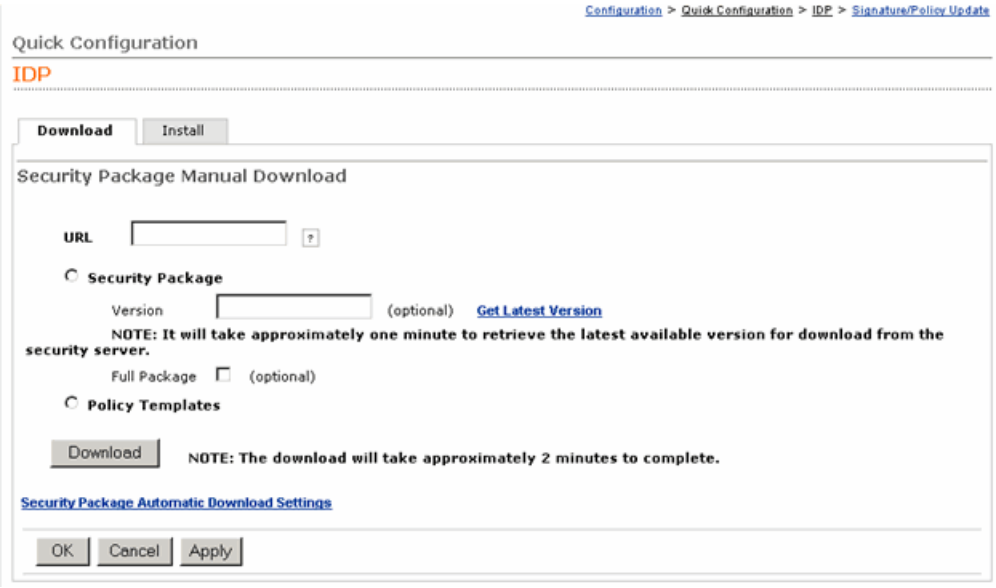


To configure security package update with Quick Configuration:

1. Select Configuration > Quick Configuration > Security Policies > IDP.
2. From the IDP policies page, click Security Package Update.

You can also navigate to this page by selecting Configuration > Quick Configuration > IDP.

3. From the IDP page, click Signature/Policy Update.

   Figure 152 on page 714 shows the Quick Configuration page for Security Package Update.

4. Fill in the information as described in Table 107 on page 715.

5. 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 107: 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. |
| 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. |

**Table 107: Security Package Update Quick Configuration Summary** *(continued)*

| Field | Function | Actions |
|---|---|---|
| Enable Schedule Update | Specifies to enable the device to automatically download the updated signature database from the specified URL. | Click the option button.<br><br>By default, Schedule update is disabled. |
| **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 Web site. 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 711
- Understanding the IDP Signature Database on page 705

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

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

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 705

When updating the signature database, the signature database update client connects to the Juniper Networks Web site 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 718
- Updating the Signature Database Manually on page 712
- Updating the Signature Database Automatically on page 716

## Verifying the Signature Database

This topic contains:

- Verifying the Policy Compilation and Load Status on page 718
- Verifying the Signature Database Version on page 720

### *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 trace option flags and record these operations to a log file in /var/db/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:

```
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)

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.

```
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 20

# 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.

This topic covers:



## 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 640
- Understanding the IDP Signature Database on page 705

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 712.

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 670.

### *Related Topics*

- Understanding Service and Application Bindings on page 722
- Configuring IDP Policies for Application Identification on page 725
- Disabling Application Identification on page 726
- Setting Memory and Session Limits on page 727

## 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 640
- Understanding the IDP Signature Database on page 705

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 108 on page 723 summarizes the behavior of application and service bindings with application identification.

**Table 108: Applications and Services with Application Identification**

| Attack Object Fields | Binding Behavior | Application Identification |
|---|---|---|
| :application (http)<br>:service (smtp) | ■ Binds to the application HTTP.<br>■ 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 109 on page 723 summarizes the binding with the application configuration in the IDP policy.

**Table 109: 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.<br>■ Disabled for service-based attack objects. |
| Specific application | Binds to the application specified in the attack object definition. | Disabled |

**Table 109: Application Configuration in an IDP Policy** *(continued)*

| Application Type in the Policy | Binding Behavior | Application Identification |
|---|---|---|
| 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 725
- Disabling Application Identification on page 726

## 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 721
- IDP Policies Overview on page 640
- Understanding the IDP Signature Database on page 705

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 722
- Configuring IDP Policies for Application Identification on page 725

## 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 721
   - IDP Policies Overview on page 640
   - Understanding the IDP Signature Database on page 705
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 712

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 725
- Related Topics on page 726

### *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 652).
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 726
- Setting Memory and Session Limits on page 727
- Verifying Application Identification on page 729

## 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 721
- IDP Policies Overview on page 640
- Understanding the IDP Signature Database on page 705

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 726
- Related Topics on page 727

### *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 722
- Configuring IDP Policies for Application Identification on page 725
- Setting Memory and Session Limits on page 727
- Verifying Application Identification on page 729

## 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 721
   - IDP Policies Overview on page 640
   - Understanding the IDP Signature Database on page 705
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 712

- 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. Matched application is saved to cache so that the

next session can use it. This prevents 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, when too many connection requests overwhelm and exhaust all the allocated resources on the system.

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 either J-Web or the CLI configuration editor to configure memory and session limits for application identification.

This topic contains:



## 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

## Verifying Application Identification

This topic covers:



### *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.

```
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.

```
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 21

# IDP SSL Inspection



## 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 TCP/UDP 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 110 on page 734 shows the encryption algorithms supported by the SRX-series devices.

**Table 110: 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 383. Table 111 on page 734 shows the supported SSL ciphers.

**Table 111: Supported SSL Ciphers**

| Version | Cipher Suites | Value |
|---|---|---|
| | 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—An RSA key exchange method uses an RSA SecurID external authentication server. SecurID is an authentication method that allows you to enter either static or dynamic passwords as your credentials. A dynamic password is a combination of your PIN and a randomly generated token that is valid for a short period of time, approximately one minute. A static password is preset on the SecurID server.
- 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 440.

**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 SRX 5600 and SRX 5800 series, 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 “\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 {
         session <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 SRX 5600 and SRX 5800 devices.

Chapter 22

# IDP Logging

This topic covers:



## 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 652.

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 643
- IDP Policies Overview on page 640

## 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 739.
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 712

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

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

# Part 4
# Index

■ Index on page 745

# Index

## B



## C

## D



## E



## F

## G



## H



## I

## J



## L



## M

## N



## O



## P

## Q



## R

## S

## T

## U



## V



## W



## X



## Z