JUNOS® Software with Enhanced Services

Design and Implementation Guide

for J-series Services Routers

Release 9.2

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

Part Number: 530-025662-01, Revision 1

Revision History
August 2008—Revision R1

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 Juniper Networks, Inc. and its subsidiaries (collectively "Juniper"), and 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, and updates and releases of such software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller. "Embedded Software" means Software which Juniper has embedded in the Juniper 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 the 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 software on multiple computers requires multiple licenses, regardless of whether such computers 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 the Embedded Software on non-Juniper equipment; (j) use the 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 for the Software are exclusive of taxes, withholdings, duties, or levies (collectively “Taxes”). Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software.

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)).

Table of Contents

About This Guide

This preface provides the following guidelines for using the *JUNOS Software with Enhanced Services Design and Implementation Guide* and related Juniper Networks, Inc., technical documents:

Objectives

This guide contains information you need to design and implement specific deployment scenarios using the JUNOS software with enhanced services on J-series Services Routers, including

- Design instructions and concepts for IPSec VPNs and firewalls
- Detailed instructions on how to implement the designs using the JUNOS software with enhanced services

NOTE: This manual documents Release 9.2 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/.

Supported Routing Platforms

This feature describes features supported on J-series Services Routers running JUNOS software with enhanced services.

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. This 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, wilfully negligent, or hostile manner; and must abide by the instructions provided in the documentation.

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 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, see “JUNOS Software Documentation for J-series Services Routers and SRX-series Services Gateways” on page vi. 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 ■ Installing hardware and establishing basic connectivity ■ Initially setting up the router	J-series Services Routers: ■ *JUNOS Software with Enhanced Services Quick Start Guide* ■ *JUNOS Software with Enhanced Services Hardware Guide* ■ *JUNOS Software with Enhanced Services Release Notes* SRX-series services gateways: ■ *SRX 5600 Services Gateway Getting Started Guide* ■ *SRX 5800 Services Gateway Getting Started Guide*

Table 1: Tasks and Related Documentation (continued)

Task	Related Documentation
Migration from ScreenOS or JUNOS to JUNOS Software with Enhanced Services (if necessary)	
■ Migrating from JUNOS Release 8.2 or higher to JUNOS software with enhanced services ■ Migrating from ScreenOS Release 5.4 or higher to 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	
Configuring device interfaces	■ *JUNOS Software Interfaces and Routing Configuration Guide* ■ *JUNOS Software CLI Reference*
Services Router Deployment Planning and Configuration	
■ Understanding and gathering information required to design network firewalls and IPSec VPNs ■ Implementing a JUNOS software with enhanced services firewall from a sample scenario ■ 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: ■ Stateful firewall policies ■ Zones and their interfaces and address books ■ IPSec VPNs ■ Firewall screens ■ Interfaces modes: Network Address Translation (NAT) mode and Route mode ■ Public Key Cryptography ■ Application Layer Gateways (ALGs) ■ Chassis clusters ■ Intrusion Detection and Prevention (IDP)	■ *JUNOS Software Security Configuration Guide* ■ *JUNOS Software CLI Reference*
Routing Protocols and Services Configuration	
■ Configuring routing protocols, including static routes and the dynamic routing protocols RIP, OSPF, BGP, and IS-IS ■ Configuring class-of-service (CoS) features, including traffic shaping and policing ■ Configuring packet-based stateless firewall filters (access control lists) to control access and limit traffic rates ■ Configuring MPLS to control network traffic pattern	■ *JUNOS Software Interfaces and Routing Configuration Guide* ■ *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)

Table 1: Tasks and Related Documentation (continued)

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

Documentation Conventions

Table 2 defines 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.

Table 3 defines the text and syntax conventions used in this manual.

Table 3: Text and Syntax Conventions (Page 1 of 2)

Convention	Element	Example
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> **configure**
`Fixed-width text like this`	Represents output on the terminal screen.	`user@host> show chassis alarms` `No alarms currently active`
Italic text like this	■ Introduces important new terms. ■ Identifies book names. ■ Identifies RFC and Internet draft titles.	■ A policy *term* is a named structure that defines match conditions and actions. ■ *JUNOS System Basics Configuration Guide* ■ RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*

Table 3: Text and Syntax Conventions (Page 2 of 2)

Convention	Element	Example
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	■ To configure a stub area, include the stub statement at the [edit protocols ospf area *area-id*] hierarchy level. ■ The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric *metric*>;
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast | multicast (*string1* | *string2* | *string3*)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [*community-ids*]

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

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

JUNOS Software Documentation for J-series Services Routers and SRX-series Services Gateways

Table 4 lists the manuals and release notes for J-series Services Routers running JUNOS software with enhanced services and SRX-series services gateways running JUNOS software.

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

Table 4: JUNOS Software Documentation for J-series Services Routers and SRX-series Services Gateways (Page 1 of 2)

Document	Description
All Platforms	
JUNOS Software Interfaces and Routing Configuration Guide	Explains how to configure J-series and SRX-series interfaces for basic IP routing with standard routing protocols, ISDN service, data link switching (DLSw), firewall filters (access control lists), and class-of-service (CoS) traffic classification.
JUNOS Software Security Configuration Guide	Explains how to configure and manage J-series and SRX-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 J-series and SRX-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 J-series and SRX-series devices. This guide also describes the configuration statements and operational mode commands unique to those 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, J-series, and SRX-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, J-series, and SRX-series devices.
J-series Services Routers Only	
JUNOS Software with Enhanced Services Design and Implementation Guide	Provides guidelines and examples for designing and implementing IP Security (IPSec) virtual private networks (VPNs), firewalls, and routing on J-series routers running JUNOS software with enhanced services.
JUNOS Software with Enhanced Services Quick Start Guide	Explains how to quickly set up a J-series 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.
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.

Table 4: JUNOS Software Documentation for J-series Services Routers and SRX-series Services Gateways (Page 2 of 2)

Document	Description
JUNOS Software with Enhanced Services Release Notes	Summarizes new features and known problems for a particular release of JUNOS software with enhanced services on J-series Services Routers, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for JUNOS software with enhanced services.
SRX-series Services Gateways Only	
JUNOS Software for SRX-series Services Gateway Release Notes	Summarizes new features and known problems for a particular release of JUNOS software on SRX-series services gateways, including J-Web interface features and problems. The release notes also contain corrections and updates to the manuals.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. Send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at http://www.juniper.net/techpubs/docbug/docbugreport.html. 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

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 an on-line self-service portal called the Customer Support Center (CSC) 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 Manager: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) tool 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 Manager 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.

Chapter 1

About JUNOS Software with Enhanced Services

JUNOS software with enhanced services integrates and evolves Juniper Network network security and routing features to produce a platform that provides world-class network and security technologies.

JUNOS software with enhanced services implements a modular architecture. The operating system consists of discrete processes that implement specific features and functionality. These processes are compartmentalized with their associated data to provide for functional independence, information privacy, and improved scalability.

JUNOS Software with Enhanced Services Feature Summary

Table 1 identifies the JUNOS software with enhanced services features.

Table 1: JUNOS Software with Enhanced Services Features Summary

Feature Category	JUNOS Software with Enhanced Services Feature
IP address management	Static addresses IPv4 and IPv6
	Dynamic Host Configuration Protocol (DHCP) ■ DHCP server ■ DHCP client ■ DHCP address pools ■ DHCP relay ■ DHCP static mapping
Internet Protocols	IPv4 ■ IP options ■ Broadcast Internet Datagrams IPv6 ■ Routing and forwarding ■ Global address configuration ■ Internet Control Management Protocol (ICMP) Domain Name System (DNS) ■ Proxy ■ DNS with VPN ■ Dynamic DNS (DDNS)

Table 1: JUNOS Software with Enhanced Services Features Summary (continued)

Feature Category	JUNOS Software with Enhanced Services Feature
Routing Protocols and Multicast	■ Open Shortest Path First: OSPFv2 and OSPFv3 ■ Border Gateway Protocol: BGP4 and BGP4 extensions for IPv6 ■ Routing Information Protocol: RIPng and RIPv2 ■ Static routes: IPv4 and IPv6 ■ Neighbor Discovery Protocol and Secure Neighbor Discovery Protocol ■ Compressed Real-Time Transport Protocol (CRTP) ■ Intermediate System-to-System (IS-IS) ■ Routing over IPSec VPN tunnels ■ Multiple virtual routers (VRs) ■ MultiProtocol Label Switching (MPLS) ■ Multicast ■ Internet Management Protocol (ICMP) ■ Protocol Independent Multicast (PIM) ■ Distance Vector Multicast Routing Protocol (DVMRP) ■ Single-source multicast ■ Multicast Source Discovery Protocol (MSDP) ■ Session Announcement Protocol (SAP) and Session Description Protocol (SDP) ■ Remote Procedure Call (RPC ALG) ■ Sun Microsystems RPC Open Network Computing (ONC) ■ Microsoft RPC Distributed Computing Environment (DCE)
Encapsulations	■ Ethernet ■ Media access control (MAC) ■ 802.1p tagging ■ Point-to-Point Protocol over Ethernet (PPPoE) ■ Asynchronous Transfer Mode (ATM) for asymmetric digital subscriber line (ADSL) or symmetric high-speed digital subscriber line (SHDSL) ■ Circuit cross-connect (CCC) ■ Translational cross-connect (TCC) ■ Synchronous Point-to-Point Protocol (PPP) ■ Frame Relay ■ High-Level Data Link Control (HDLC) ■ Serial encapsulation over RS-232, RS-449, X.21, V.35, and EIS-530 connections ■ 802.1Q filtering and forwarding ■ Multilink Frame Relay (not supported on native serial interface) ■ Multilink PPP ■ Generic routing encapsulation (GRE) and IP-over-IP. ■ GRE with IP multicast and GRE keepalive packets. (Fragmentation reassembly is not supported.)
Traffic Management	■ Three-level scheduling with weighted round-robin (WRR) ■ Policing and shaping ■ Class-based queuing with prioritization

Table 1: JUNOS Software with Enhanced Services Features Summary (continued)

Feature Category	JUNOS Software with Enhanced Services Feature
Security	■ Security policies ■ Packet filters ■ Zones ■ User (and administrator) authentication, and firewall authentication: passthrough and Web authentication ■ Screens ■ Network Address Translation (NAT) ■ Application Layer Gateway (ALG) support, including the following gateways: ■ Internet Inter-ORB Protocol (IIOP) ■ Winframe ■ Session Initiation Protocol (SIP) ■ H.323 ■ Media Gateway Control Protocol (MGCP) ■ Skinny Client Control Protocol (SCCP) ■ UNIX remote shell services ■ File Transfer Protocol (FTP) ■ Real-Time Streaming Protocol (RTSP) ■ Trivial File Transfer Protocol (TFTP) ■ SQLNET ■ Domain Name System (DNS) ■ Internet Control Message Protocol (ICMP) ■ TALK ■ NetShow ■ Simple Network Management Protocol-GET (SNMP-GET) ■ Compressed Real-Time Transport Protocol (CRTP) ■ Jumbo Frame Support
IPSec VPNs	■ Policy-based and route-based IPSec VPNs ■ Internet Key Exchange (IKE) key and manual key management ■ IPSec peer types ■ Peers with static IP addresses ■ Peers with dynamic IP addresses (dynamic endpoint site-to-site IPSec VPNs) ■ Remote access by means of user IKE peers and user-group IKE peers (remote access IPSec VPNs) ■ IKE authentication preshared keys and certificates with the following features: ■ Certificate encodings: X509 and PKCS7 ■ RSA signatures ■ Advanced Encryption Standard (AES) 128-bit, 192-bit, and 256-bit encryption ■ 56-bit Data Encryption Standard (DES) and 168-bit 3DES encryption ■ MD5 and Secure Hash Algorithm (SHA-1) authentication (For IPSec IKE Phase 2, SHA-2 is also supported.) ■ Antireplay (packet replay attack prevention) ■ Next-hop tunnel binding

Table 1: JUNOS Software with Enhanced Services Features Summary (continued)

Feature Category	JUNOS Software with Enhanced Services Feature
System Management	■ JUNOScript XML application programming interface (API) ■ J-Web browser interface—for Services Router configuration and management ■ JUNOS command-line interface (CLI)—for Services Router configuration and management through the console, Telnet, SSH, or J-Web CLI terminal ■ Simple Network Management Protocol: SNMPv1, SNMPv2, and SNMPv3 ■ J-Flow flow monitoring and accounting ■ Packet capture (PCAP)
Traffic Analysis	■ Real-time performance monitoring (RPM) ■ J-Web event viewer
Activity Logging and Monitoring	■ Traceroute ■ System log ■ RADIUS external administrator database support

Administration	■ Configuration rollback
	■ Autoinstallation
	■ Button-operated configuration rescue (CONFIG)
	■ Confirmation of configuration changes
	■ Software upgrades
	■ Supports the following features for automating network operations and troubleshooting:
	  ■ Commit scripts
	  ■ Operation scripts
	  ■ Event policies

Class of Service	■ DiffServ ■ Classification ■ Traffic marking ■ Scheduling ■ Policing and shaping ■ Intelligent drop mechanisms ■ Link efficiency mechanism ■ Policy management ■ Queuing for GRE and IP-IP tunnels

Understanding Secure and Router Contexts

As shipped from the factory, a Services Router running the JUNOS software with enhanced services 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, use a predefined template configuration file. If you plan to use the Services Router primarily as a router, change to router context, using 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.

For information describing how to change contexts, see the *JUNOS Software Administration Guide*.

About 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 at 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.

Secure Context Default Configuration

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 is configured to allow for 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, a queue size of 2000, and a timeout value of 20 seconds.)

- The default policy for the untrust zone is to deny all traffic.

Figure 1 shows the default configuration file for secure context.

Figure 1: Default Configuration for Secure Context

```
system {
    autoinstallation {
        delete-upon-commit;
        traceoptions {
            level verbose;
            flag {
                all;
            }
        }
    }
    syslog {
        file messages {
            any any;
        }
    } services {
        ssh;
        web-management {
            http {
                interface [ ge-0/0/0.0 ];

            }
        }
    }
}

interfaces {
        ge-0/0/0 {
            unit 0;
        }
    }

security {
        screen {
            ids-option untrust-screen {
                icmp {
                    ping-death;
                }
                ip {
                    source-route-option;
                    tear-drop;
                }
                tcp {
                    syn-flood {
                        alarm-threshold 1024;
                        attack-threshold 200;
                        source-threshold 1024;
                        destination-threshold 2048;
                        queue-size 2000;
                        timeout 20;

                    }
                    land;
                }
            }
        }

        policies {
            from-zone trust to-zone trust {
```

```
                policy default-permit {
                    match {
                        source-address any;
                        destination-address any;
                        application any;
                    }
                    then {
                        permit;
                    }
                }
            }
            from-zone trust to-zone untrust {
                policy default-permit {
                    match {
                        source-address any;
                        destination-address any;
                        application any;
                    }
                    then {
                        permit;
                    }
                }
            }
            from-zone untrust to-zone trust {
                policy default-deny {
                    match {
                        source-address any;
                        destination-address any;
                        application any;
                    }
                    then {
                        deny;
                    }
                }
            }
        }

        zones {
            security-zone trust {
                tcp-rst;
                interfaces {
                    ge-0/0/0.0 {
                        host-inbound-traffic {
                            system-services {
                                http;
                                https;
                                ssh;
                                telnet;
                                dhcp;
                            }
                        }
                    }
                }
            }
            security-zone untrust {
                screen untrust-screen;
            }
        }
    }
}
```

About Router Context

Router context allows a Services Router to act as a router, in which all management and transit traffic is allowed. In router context, the Services Router forwards all packets unless you configure a security policy that denies specific traffic.

JUNOS software with enhanced services is a hardened operating system. You can use it with more relaxed checks for host-inbound traffic and configure the dataplane 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.

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 router context configuration details, see “About Router Context” on page 9.

For information describing how to change contexts, see the *JUNOS Software Administration Guide*.

Router Context Default Configuration

Router context allows the router to act as a router in which all management and transit traffic is allowed.

All interfaces are bound to the trust zone. In router context, all packets are forwarded unless there is a security policy defined that denies specific traffic. Host inbound traffic from all predefined services is allowed. Router context configuration values are defined in a configuration file named jsr-series-routermode-factory.conf.

Figure 2 shows the default configuration file for router context.

Figure 2: Default Configuration for Router Context

```
system {
    syslog {
        file messages {
            any any;
        }
    }
    services {
        telnet;
        ssh;
        web-management {
            http {
                interface [ ge-0/0/0.0 ];
            }
        }
    }

interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 192.168.1.1/24;
            }
        }
```

```
    }
}

security {
    flow {
        allow-dns-reply;
        tcp-session {
            no-syn-check;
            no-syn-check-in-tunnel;
            no-sequence-check;
        }
    }
    forwarding-options {
        family {
        iso {
                mode flow-based;
            }
        inet6 {
                mode packet-based;
            }
        }
    }
    policies {
        default-policy {
            permit-all;
        }
    }
    zones {
        security-zone trust {
            tcp-rst;
            host-inbound-traffic {
                system-services {
                    any-service;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                all;
            }
        }
    }
    alg {
        dns disable;
        ftp disable;
        h323 disable;
        mgcp disable;
        real disable;
        rsh disable;
        rtsp disable;
        sccp disable;
        sip disable;
        sql disable;
        talk disable;
        tftp disable;
        pptp disable;
        msrpc disable;
        sunrpc disable;
    }
}
```

Chapter 2

Designing Firewalls for Branch Offices

This chapter describes how to design a JUNOS software with enhanced services stateful firewall for a mid-size branch bank office of an enterprise corporation. The branch site intends to deploy the router to be used as both an exterior Internet-facing firewall and an interior corporate-facing firewall. This chapter explores some of the practices to follow in defining a deployment design.

For details on implementing a firewall design, see “Implementing Firewall Deployments for Branch Offices” on page 37. For more information on firewall configuration, see the *JUNOS Software Security Configuration Guide*.

This chapter contains the following sections:

Business Case for Stateful Firewalls

Businesses today are highly networked. They must provide their employees with access to systems and applications internal to the local area network (LAN) and at corporate headquarters. They must selectively provide their users with Internet services in a manner that allows them to secure their assets against intrusions and attacks of any kind originating outside their LANs. They must also protect their networks from attempts of employees and other corporate personnel to access network resources for which they are not authorized and to launch attacks within the LAN and across the Internet, either unwittingly or intentionally.

Firewalls of the recent past, based on packet filtering strategies, satisfy some of these security requirements. However, enterprise corporations are finding that packet-filtering firewalls do not give them the protection that they require against increasingly common and sophisticated intrusions and attacks.

To solve their problem, corporations turn to deployment of stateful firewalls that provide features of packet filtering firewalls but also monitor the state of the communication and connections within and across sessions.

About the Design and Deployment of Stateful Firewalls

Stateful firewalls allow a system to monitor and manage its session connections by adding transport layer (Layer 4) awareness to the capabilities of firewalls. Regardless of the protocol and how it manages its state of communication, a stateful firewall needs to keep track of the communication status between a source and destination host. This information is stored in what is called a state table.

Stateful firewalls add the concept of history to packet filtering firewalls. They can assess the validity of packets within the context of previously transmitted packets within and across sessions.

A stateful firewall assesses the connection state differently based on the protocol it manages. For example, some stateful firewalls enforce security requirements against attacks launched through misuse and abuse of the TCP/IP protocol suite.

About the JUNOS Software with Enhanced Services Stateful Firewall

The design for the deployment explored in this document is based on the JUNOS software with enhanced services stateful firewall and the router's policing, traffic shaping, and class-of-service (CoS) filtering capability. JUNOS software with enhanced services is a session-based operating system that monitors traffic from source to destination bidirectionally. A packet is assessed in relation to other packets of its session and across sessions.

Table 2 lists some highlights of the firewall's features.

Table 2: Summary of the JUNOS Software with Enhanced Services Firewall Features

Firewall Features	Purpose
JUNOS software with enhanced services router policing	■ Rate-limit, traffic shape, and assign CoS attributes to packets at the ingress and egress interfaces. For details, see the *JUNOS Software Interfaces and Routing Configuration Guide*.
JUNOS software with enhanced services screens capability	■ Screen traffic, and permit or deny it based on the type of packet and packet header information. Set thresholds and rate-limits to prevent denial-of-service (DoS) floods before traffic is allowed to pass into the system. Audit and log traffic counts and alarms. ■ Balance protection and access. ■ Thwart attacker attempts to map your network and thereby gaining information about its entry points. Screens can ■ Identify and thwart reconnaissance attempts to sweep the network for IP addresses using ICMP pings. ■ Detect and block attacker's attempts to scan hosts and ports for entry points into the LAN using anomalous TCP packet headers. ■ Protect the network and the firewall against denial-of-services (DoS) attacks, which lock out legitimate users from accessing resources they need, and from other kinds of attacks that damage your network. Using screens, you can ■ Set threshold and rate limits to drop traffic to guard against DoS flood attacks intended to absorb sessions, bring down one or more hosts, and do other kinds of damage. ■ Set limits on the number of sessions from any one IP address or to any one IP address to block attempts at session table exhaustion leading to lockout of legitimate sessions and damage to the firewall and network. ■ Set alarm thresholds to trigger an alarm when a threat occurs that could indicate an attack. ■ Protect against common operating system DoS attacks, such as teardrop, land, ping-of-death, and WinNuke attacks. You apply screens at the zone level. You can use screens to protect one zone differently from another, depending on the requirements of its hosts, applications, and other resources.
JUNOS software with enhanced services policy-based access control	■ Control access to network resources within and across zones through policies, which define who has access to which resources and when. ■ Direct traffic to a tunnel for secure transmission across the Internet using an IPSec VPN. ■ Audit traffic. ■ Set a schedule for when a policy is active.

Overall Firewall Design Guidelines

Although every network has unique characteristics that require unique firewall deployment solutions, the following principles can guide you in the design of any stateful firewall:

- Keep it simple. The simpler the firewall, the more likely it is to be secure and the easier it will be to manage. Allow for complexity to meet special requirements, but consider that unwarranted design complexity lends itself to configuration errors.

- Within the concept of keep-it-simple:

- Create zones that are specific to functional requirements—for example, groups of employees sharing the same job functions and access rights, or applications and database servers that are accessible by various groups of users belonging to other zones.
- Do not mix groups of users and servers in the same zone.
- Design for use of policies that are specific rather than general. Place general policies at the bottom of your policy list.
- Summarize groups of users by using subnets. Apply policies to a group of users based on its subnet rather than to individual users, for ease of use and to avoid configuration errors.

Table 3 provides an overall strategy to follow in designing the deployment of a JUNOS software with enhanced services stateful firewall.

Table 3: Designing a JUNOS Software with Enhanced Services Firewall

Task	Instructions
■ Obtain the corporate security policy for the branch office, or define one.	See "Obtaining or Defining a Security Policy for the Site" on page 15.
■ Define a firewall philosophy.	See "Defining and Documenting the Firewall Philosophy" on page 17.
■ Understand your current network resources and traffic patterns.	See "Developing a Network Profile" on page 20.
■ Determine the topology for the firewall deployment.	See "Determining Placement of the Firewall in the Network Topology" on page 25.
■ Determine the JUNOS software with enhanced services firewall features to be used to execute the firewall policy.	See "Implementing Firewall Deployments for Branch Offices" on page 37.
■ Implement the firewall design	See "Implementing Firewall Deployments for Branch Offices" on page 37.

Obtaining or Defining a Security Policy for the Site

Before a network can be secured for business, a network security policy must be established. A network security policy identifies all the network resources belonging to the site and the required security for each of them. Usually it includes a network map, updated regularly as systems are added or removed from the topology.

An effective network security policy:

- Is current and complete.
- Defines the organization's culture with respect to security and how its policies are applied.
- Is readily available through some means on the LAN to employees and other responsible parties.

Usually corporate policy dictates the network security policy for its branch and regional sites. But smaller enterprises should also establish security policies that their administrators can rely on for direction.

A firewall, and other security measures, such as deployment of VPNs, are designed to execute a portion of the security policy. A well-documented firewall philosophy, which is a more specific part of a security policy, can guide network administrators in managing evolution of the firewall and maintaining it. A firewall security philosophy might also define security threats and the actions to be taken to thwart those threats. "Defining and Documenting the Firewall Philosophy" on page 17 explains what a firewall philosophy is and how to define one.

Table 4 identifies some of the procedures that a corporation might follow in establishing its security policy. There are many commercial products available that you can purchase to help you define a security policy. Because corporations vary in the way they construct their security policies, the procedures shown in Table 4 serve as an example only.

Table 4: Network Security Policy Definition

Task	Instructions
Define your Environment	■ Document the site's network assets to be protected. Identify the services and systems you want to protect. You cannot deploy a robust firewall to be used successfully unless you have determined what you must protect.
Identify resources, systems critical to the network, and other systems requiring strong defense tactics.	■ Create network diagrams and maps that identify the following information: ■ The locations of all hosts in your system and the operating systems that they run. ■ The types and locations of other devices, such as bridges, routers, switches. ■ The types and locations of terminal servers and remote connections. ■ Descriptions and locations of any network servers, including operating system, configuration information, and application software and versions that they run. ■ Location and description of any network management systems used.
Define your current security policy implementation.	Describe any existing security mechanisms used. Identify the following technology and any other mechanisms you use: ■ Anti-virus programs ■ Firewalls, if others ■ Security hardware, such as encryptors for servers ■ VPNs
Define the main threats in plain language and the response to be taken in the event of a security breach or attack.	Define threats to the system. Define the actions administrators take after an attack has been identified and resolved. For example: ■ Will you attempt to identify the attacker? ■ Do you plan to prosecute? ■ Do administrators contact the Internet service provider (ISP) to report the attack?

Defining and Documenting the Firewall Philosophy

Documentation of the firewall philosophy is the part of a site's network security policy that applies strictly to the firewall. It defines the overall goals for the firewall and states explicitly what services and applications are allowed for use on the LAN. Setting a firewall philosophy provides guidelines that you can follow in implementing the firewall deployment. If you identify how services are to be protected and who is allowed to use them, it is much easier to define and configure the firewall itself.

A firewall philosophy is also essential as new hosts and software are added to the network. Firewall philosophy documentation serves as a means of communicating the current firewall deployment and factors that contribute to its deployment to succeeding management personnel and employees.

Even less complex firewalls need a well documented firewall philosophy to guide their deployment design and maintenance. Without a philosophy to guide implementation and administration, the firewall itself might become a security problem. Table 5 identifies some firewall philosophy components you can include in your documentation.

Table 5: Firewall Philosophy Guidelines

Task	Steps
Identify how the firewall will handle services and applications traffic.	Analyze the current applications, services, and protocols, that are allowed, and publish the results of the analysis as a model for future decisions.
Document the criteria to be used to determine that an application is required for business.	■ Answer the following questions: ■ Why is an application allowed? ■ Why is an application denied?
Identify vulnerabilities associated with each of the allowed applications to be used. For each application and service to be allowed, identify the internal host and operating system on which it runs.	■ List the allowed applications and how they are secured. The list must identify any known vulnerabilities associated with each application and the methods used to secure the application. It might also include known attacks related to the applications. Here is a brief example of some of the things you might consider: ■ Is it policy to restrict traffic to secure protocols, such as HTTPS and SSH? ■ Are HTTP, FTP and Telnet traffic allowed? Under what circumstances—for example, only for use in an IPSec VPN? ■ Is DNS supported? Most Internet services rely heavily on it. But is it required?
Create a work flow document that describes by role the interactions your employees and other personnel have with the network system.	■ List the job functions of each role. ■ List the data and network accesses that each function requires. ■ List the applications that are allowed for each role. ■ Document the list of services on the Internet that are permissible for use by your employees. For an example workflow, see Table 6 on page 19.

Table 5: Firewall Philosophy Guidelines (continued)

Task	Steps
Identify the objectives for your firewall deployment.	Define your primary goals. Are they: ■ To protect against threats from outside your organization? ■ To protect against insider attacks? ■ For uses unrelated to security, such as to maintain control over network usage? Define are your goals in regard to integrity, confidentiality, and availably. Define what are your requirements for manageability versus sophistication.
Define what constitutes an attack.	Determine, for example, whether you consider information-gathering (reconnaissance missions) an attack. Do you restrict qualification of attacks to incidents that do damage?
Specify if private addressing is to be used.	Identify the subnetworks to be used. Specify whether you plan to use Network Address Translation (NAT).
Specify how the firewall is to be managed and updated.	Identify management tools, audits, and scheduled downtime for periodic testing. Define how alerts and alarms are to be used.
Identify security vulnerabilities in the network and rectify them.	Record this information in your network philosophy document for historical purposes.
Test the network to ascertain that it has not been breached and that it is not infected with viruses.	Document the process to be followed.

To guide the development of your firewall philosophy, you can establish an overall approach or security stance of least privilege or greatest privilege, depending on your network requirements.

- Least privilege. Block all network connections in both directions. After all interzone and intrazone traffic is blocked, you can unblock it selectively through policy configuration. The policy configuration can then define precisely and incrementally what is allowed.

  Least privilege is the more common approach to deployment of a firewall. It is the default position of the JUNOS software with enhanced services system, when you start the system in security context.

- Greatest privilege. Trust everything inside the network. The policy can then designate specific denial of access, to close down access as appropriate.

  This stance is sometimes taken when the firewall is deployed inline while network activity continues. In this case, the stance allows the firewall to be deployed without disturbing normal business activity that is conducted using the network.

Some sites deploy the firewall inline and set and use logs to capture information to identify common, successful attacks. In this case, parts of the network might succumb to an attack. However, based on the logged information, the network administrator can have a better sense of common attacks on the LAN and, thus understand more definitely the appropriate firewall screens and thresholds to put in place.

Documenting Allowed Applications and Employee Role Workflow

Before you can define polices for your firewall, you need to understand and characterize your network environment, including identification of applications that are currently used on the network. In some cases, network administrators are unaware of some of the applications that employees use, especially in regard to use of the Internet. For example, employers might not know if employees use Instant Messaging, and other similar applications. Both employers and employees might not be aware that these kinds of applications open entry points into the network that provide easy access for attackers.

It is good practice to maintain a list of allowed applications, any known security risks associated with them, and the means used to secure the application. This kind of information can be maintained on a corporate intranet and made available to employees.

It is also important to understand and document the work flow in your organization, based on employee roles and the applications that are allowed and required for each role. Table 6 gives an example of a spreadsheet that is used for this purpose.

Table 6: Employee Roles, Access Rights, and Allowed Applications

Employee Role	Access Rights	Allowed Applications
Tellers	■ Allowed access to the customer checking and savings records at corporate headquarters. ■ Allowed access to the mail server. ■ Not allowed Internet access.	■ SNMP ■ SSH ■ HTTPS ■ custom application (custApp)
Bank Managers	■ Allowed access to both servers at corporate headquarters: the customer checking and savings records and the customer special services records. ■ Allowed access to the mail server. ■ Allowed Internet access.	■ SNMP ■ FTP ■ SSH ■ HTTPS ■ Telnet ■ rlogin ■ custApp

Table 6: Employee Roles, Access Rights, and Allowed Applications (continued)

Employee Role	Access Rights	Allowed Applications
Financial Managers	■ Allowed access to the customer special services records at corporate headquarters. ■ Allowed access to the mail server. ■ Allowed Internet access.	■ SNMP ■ FTP ■ SSH ■ HTTPS ■ Telnet ■ custApp
IT Operations Personnel	■ Allowed access to both servers at corporate headquarters: the customer checking and savings records and the customer special services records. ■ Allowed access to the mail server. ■ Allowed Internet access.	■ SNMP ■ FTP ■ SSH ■ HTTPS ■ Telnet ■ rlogin ■ custApp

Creating these documents can help you define the firewall policy to be configured for your employees. Most of the work will be done, and the firewall policy configuration then becomes a software configuration task.

After you have defined the allowed applications for employees, communicate the information in a way that is visible to all current users and available to new employees.

Developing a Network Profile

Before you can design an effective firewall policy that protects your network resources against attack without severely restricting legitimate users, you need to understand your network traffic patterns. Knowing what is normal for your network and setting a baseline against which to measure what you think is anomalous behavior, are key to understanding how to use screens and set their thresholds to protect against attacks.

To define a baseline for your network under normal operating conditions, you must monitor the network for at least a week. Many commercial and open-source tools are available for this purpose, such as MRTG and NetMGR. You can use SNMP.

NOTE: In most cases, you can use any router that is already deployed to gather the information required to establish a network baseline. After you have configured the JUNOS software with enhanced services firewall and deployed it, you can also use the JUNOS software with enhanced services CLI to collect information about your normal network traffic patterns to tune your network security.

Table 7 identifies some of the tasks involved in creating a detailed profile of your network's normal behavior.

Table 7: Creating a Network Traffic Profile

Task	Instructions
Create a network traffic baseline profile.	See "Setting a Network Traffic Baseline Profile" on page 22.
Create a profile to characterize network host connectivity. Using a JUNOS software with enhanced services firewall screen, you can rate-limit the number of sessions per IP address to avoid session table flood.	See "Defining an IP Source and Destination Sessions Profile" on page 23.
Determine the normal ICMP traffic flow. ■ You can use this information to set boundaries on ICMP traffic to avoid an ICMP address sweep. ■ Many systems use ICMP for error reporting. It is important to understand what normal ICMP traffic flow is so that you do not impede genuine error-reporting information by setting thresholds that are too low.	For details on using the CLI for this purpose, after the device is deployed, see "Collecting Network Data with JUNOS Software with Enhanced Services CLI Statements" on page 24.
Determine the normal TCP packet traffic flow. Many network attacks use malformed or hijacked TCP packets to carry out their nefarious missions.	See "Counting TCP Packets with SYN and RST Flags Set" on page 25.

Setting a Network Traffic Baseline Profile

You can use the packet-filtering features of the JUNOS software with enhanced services router to rate-limit certain types of traffic. You can also use the router's screens for this purpose. However, you cannot effectively determine the thresholds to set for specific types of traffic unless you know the normal traffic flow patterns for your network. Table 8 explains what a traffic profile is and the kind of information that contributes to its definition.

Table 8: Network Baseline Traffic Profile

About the Network Traffic Baseline Profile	
What is it?	Detailed Layer 3 to Layer 7 characterization of network traffic.
How do I create it?	1. Measure and collect session, flow, and packet statistics from real-time traffic. 2. From these statistics, create a model that describes both average aggregate behavior and average individual behavior on the network.
Information the Network Traffic Baseline Profile Provides	
What Layer 3 to Layer 7 aggregate information can I deduce from the traffic baseline I create?	■ The number of users on the network ■ How many applications these users are running ■ What percentage of sessions are of a certain protocol type
What Layer 3 to Layer 7 individual information can I deduce from the traffic baseline I create?	■ The average bandwidth consumed per user ■ The average number of sessions per user ■ The average session rate per user
What information can I obtain by comparing this data to Layer 2 to Layer 3 statistics?	■ The average packet size on your network ■ The normal error rate on your network ■ The normal fragmentation rate on your network.
Measurements Required to Create a Network Traffic Baseline Profile	
What measurements do I need to collect to calculate the average transport layer statistics?	■ Bandwidth You can collect this data from SNMP using tools such as MRTG, or you can monitor it using the command line interface (CLI) of a currently deployed router. After the device is deployed, you can use the JUNOS software with enhanced services software CLI for this purpose. ■ Session count ■ Session rate The preceding three measurements contribute to determining the average aggregate model. These measurements plus the following one constitute the average individual model. ■ User count

Table 8: Network Baseline Traffic Profile (continued)

About the Network Traffic Baseline Profile	
Average Aggregate Model Calculations	
How do I calculate the average aggregate model?	■ Session time = session count / session rate ■ Bandwidth per session = bandwidth per user / sessions per user ■ Data per session = bandwidth per session x session time
Average Individual Model Calculations	
How do I calculate the average individual model?	■ Session rate per user = session rate / user count ■ Bandwidth per user = bandwidth / user count ■ Session per user = session count / user

After you create a traffic model, you can use it to validate the methodology that you used to define the baseline. Then you can program traffic-generating test equipment to fit the traffic model and take the same measurements. If they match the measurements from the real traffic, then the model is correct.

After you deploy the JUNOS software with enhanced services router, you can use its CLI to continue to collect this information. Then you can use the results to fine-tune your firewall. Here are some of the things you can use to obtain this information:

- Set SNMP for collecting bandwidth session, and possibly session rate (by zone or interface).
- Set policy rules to generate traffic logs. Collect with the system logs.

Defining an IP Source and Destination Sessions Profile

You can determine the current connectivity of every host in the network for which you are designing the deployment by correlating session activity per source IP address and destination IP address.

Using the information you collect, you can better define a firewall to protect against session table floods without denying legitimate users session connections.

Table 9: Collecting Session information

Establishing a Profile for Active Sessions	
What do you need to know to establish a profile for active sessions?	■ How many active source IP addresses your network has ■ How many active destination IP addresses your network users communicate with ■ How many active IP sources communicate with your network. ■ How many active IP destinations your network has

Collecting Network Data with JUNOS Software with Enhanced Services CLI Statements

You can use theJUNOS software with enhanced services CLI to count certain types of packets transiting your network and use the information to fine-tune your security configurations.

Collecting TCP Control Packet Data

After you have deployed the router and you want to tune the configuration to protect against DoS attacks, you can use the JUNOS software with enhanced services counters statements to count TCP control packets. You can compare these numbers with the total amount of packets traversing your network to gain a sense of the proportion of these packets in relation to the throughput total number of packets.

TCP control traffic is generally less than 5% of all packets and 1% of bytes traversing the network.

By monitoring interface counters from the command line, you can calculate the percentage of TCP SYN packets that traverse your network.

The following sample calculations compare the number of SYN packets in a normal baseline TCP control packet profile to a measurement that shows anomalous behavior:

- Example 1: TCP control packets baseline—When TCP packets = 814280 and all traffic = 43059013:

  (814280/43059013) x 100 = 1.8% of all packets are SYN packets

- Example 2: Anomalous TCP control packet behavior—When TCP packets = 14241013 TCP and all traffic = 78210483:

  14241013 / 78210483 x 100 = 18.2% of all packets are SYN packets

  In this case, the system is most likely under a DoS attack.

Counting ICMP Ping Packets

After you have deployed the JUNOS software with enhanced services router, you can use the following CLI statements to count ICMP ping packets. You can then tune screen thresholds based on this information. After you know the normal throughput for ICMP packets, you can determine the ICMP rate-limits to apply to a zone to avoid attackers' use of massive amounts of ICMP ping packets to sweep the network for IP addresses.

user@host# **set firewall filter input-transit term ping from protocol icmp**

user@host# **set firewall filter input-transit term ping from icmp-type-echo-request**

user@host# **set firewall filter input-transit term ping from icmp-type echo-reply**

user@host# **set firewall filter input-transit term ping then count count-ping-transit**

user@host# **set firewall filter input-transit term icmp from protocol icmp**

user@host# **set firewall filter input-transit term icmp then count count-icmp-transit**

Counting SYN-NO-ACK Packets

After you have deployed the router, you can use the following JUNOS software with enhanced services CLI commands to count the SYN-NO-ACK packets and tune threshold values based on the results.

user@host# **set firewall filter input-transit term syn-no-ack from protocol tcp**

user@host# **set firewall filter input-transit term syn-no-ack from tcp-initial**

user@host# **set firewall filter input-transit term tcp-close then count count-tcp-open-transit**

user@host# **set firewall filter input-transit term tcp-close from protocol tcp**

Counting TCP Packets with SYN and RST Flags Set

After you have deployed the router, you can use the following JUNOS software with enhanced services CLI statements to count the TCP packets with SYN and RST flags set.

user@host# **set firewall filter input-transit term tcp-close from tcp-flags fin**

user@host# **set firewall filter input-transit term tcp-close from tcp-flags rst**

user@host# **set firewall filter input-transit term tcp-close then count count-tcp-close-transit**

user@host# **set firewall filter input-transit term default-permit then accept**

Determining Placement of the Firewall in the Network Topology

A fundamental decision involved in designing a firewall deployment is where to place it. As a rule, its primary use dictates largely where the firewall is deployed in the network environment. Many firewalls are deployed at the edge, or border, between the private LAN and a public network, such as the Internet. However, there are other ways in which a firewall can be deployed.

An enterprise network generally comprises two areas: the core (or internal network) and the edge. The network can also be extended to include an area called the Demilitarized Zone (DMZ) (also known as a perimeter, or bastion network). Firewalls are designed and deployed differently at these areas of a network because each area has its own requirements in regard to how resources are to be protected. Table 10 shows some of the uses of a firewall in these areas of the network.

Table 10: Network Areas and Types of Firewalls

Firewall Deployment Network Area
Edge: Internet-facing firewall ■ Protects the border of the network against unauthorized access from the Internet. Defends its hosts against all forms of attack from outside the LAN. ■ Ensures that authorized users are able to perform required tasks by thwarting denial-of-service (DoS) and other forms of lock-out attacks launched from outside the LAN. ■ Guards the entry points to the LAN by checking each packet to determine if it is allowed through.
Core: corporate-facing firewall ■ Protects corporate resources from internal opportunistic, accidental, or malicious attacks, such as data theft or DoS floods instigated through a virus. ■ Provides outgoing traffic-handling policies. Ensures that employees have access only to the Internet services they require. ■ Protects against employee use of the network to launch outside attacks.
Firewall in the DMZ: ■ Provides additional security by creating a less secure area in front of the private network to provide a first-line-of-defense behind which the internal LAN hosts can safely exist. ■ Usually contains publicly accessible servers and bastion hosts. If these servers are attacked, hosts within the LAN are not compromised.

Defining Use of Alerts and Audit Trails

Real-time alerts send system log error messages to central managements consoles when suspicious activity is detected. Enhanced audit trail features use the system log to track all transactions and to record time-stamp, source host, destination host, and ports.

After you establish a baseline profile for your network, you can set alarms to detect anomalous activity such as SYN attacks and bandwidth usage beyond the threshold. Traffic exceeding specified thresholds triggers an alarm to call the network administrator's attention to the deviation from the baseline. Network administrators can then evaluate the situation to determine if they need to take action.

For this purpose, you can use the count and alarm parameters of a policy statement.

JUNOS Software with Enhanced Services Firewall Configuration Considerations

This section explores firewall design considerations. For firewall implementation details, including the configuration of policies that dictate who has access to hosts and other resources protected by zones, see “Implementing Firewall Deployments for Branch Offices” on page 37.

Designing a firewall is an exercise in risk management:

- You do not want to lock out users and guests from resources they require—files on servers, e-mail, access to your online banking system—and you do not want to leave systems that provide this information wide open to attack.
- You do want to minimize as much as possible legitimate packet loss while defending the network by dropping attack traffic.

Risk management entails taking into account the following considerations:

- How high is the threat associated with a type of attack?
- How common is the attack for your network?
- What is the cost to protect against the attack, in terms of system performance?

  For example, what is the CPU utilization and memory consumption associated with the method, such as a screen, used to protect against the attack, given normal traffic patterns?

Security mechanisms, such as screens, are applied at the zone level to ensure that traffic initiating from hosts within a zone is not used to launch an attack. For example, screens applied to the Internet zone filter out attack traffic coming from outside the LAN. Screens applied to the Financial zone, on the LAN, assure that any attacks launched from hosts within the zone were dropped at the source, before they left the zone, to protect all other zones and the firewall itself. Screens applied to a zone do not protect the hosts within that zone against attacks. Rather, they defend against attacks launched from within that zone. When attack traffic is dropped at the source zone, it cannot be used against any other zone.

You can create zones for various segments of a network and distinguish the kind of security applied to one zone from another though policies and screens.

In determining whether to use a particular screen, consider that every screen utilizes CPU cycles and takes up memory. The number of screens you use and the amount of traffic they process can adversely affect system performance, depending on the rest of your configuration. To understand performance issues, you must look at your deployment as a whole, factoring in all parts of the configuration.

Consider the following questions:

- Should tighter constraints be put in place at the risk of loss of legitimate packets?
- Is it better to let some attack traffic through to ensure that all legitimate traffic is delivered, but not enough to flood the system into denial of service?

You might be willing to tolerate a specific kind of attack, such as a reconnaissance attack. You might discover that for your environment, certain kinds of attacks are uncommon and certain screens have little relevance.

You can rely on the network map and identification of allowed services and applications for employee portions of your security plan to better decide what screens to use. For example, if ICMP applications are not allowed by the policy for a specific zone and its hosts, you do not need to use the ICMP screens for that zone.

As mentioned previously, before you can design an effective firewall policy that protects your network resources against attack without severely restricting legitimate users, you need to understand your network patterns. Knowing what is normal for your network and setting a baseline against which to measure what you think is anomalous behavior, is key to understanding how to use screens and set their thresholds. “Developing a Network Profile” on page 20 explains how to set up a network profile.

Chapter 3

Designing IPSec VPNs for Branch Offices

This chapter explains concepts required to understand and design JUNOS software with enhanced services IP Security virtual private networks (IPSec VPNs). For details on implementing IPSec VPNs, see “Implementing Route-based IPSec VPNs for Branch Offices” on page 103, “Implementing Policy-based IPSec VPNs Branch Offices” on page 137, and “Implementing Remote Access IPSec VPN for Branch Offices” on page 167.

This chapter contains the following sections:

Business Case for IPSec VPNs

Corporate enterprises are finding it increasingly important that they provide their employees at branch and regional offices with secure access to resources at the central site. To this end, they pursue cost-effective ways to connect their outlying offices to corporate sites and to give employees who travel or work at home and business partners remote access to corporate resources. To solve this problem, corporations increasingly turn to virtual private networks (VPNs) to replace expensive wide area networks (WANs). This decision is premised on the fact that VPNs provide the same assurances of security, reliability, and scalability.

Preliminary Assumptions and IPSec VPN Design Considerations

Implementation of the deployments described in this book depends on the assumptions addressed in Table 11.

Table 11: Preliminary Assumptions for Deployment Scenario

Preliminary Assumptions
The basic network security policy has already been established. (See “Designing Firewalls for Branch Offices” on page 11.)
The network has been tested for basic connectivity and security, and problems, if any, have been resolved.
Network integration issues have been addressed. For example, plans have been made for VPN communication with network services, such as CA authorities and server directories.
A corporate network administrator, system integrator, or service engineer has installed the device and tested it for basic connectivity.
One or more root, network, or security administrators who will implement the deployment and manage the system afterward have been authenticated using one of the following authentication servers: ■ Local server/password ■ RADIUS ■ TACACS + This guide does not address authentication. For authentication concepts and configuration, see the *JUNOS Software Security Configuration Guide*.
Identities of users at the branches and remote access users have been set up in one of the following authentication databases: ■ Local server ■ LDAP server ■ SecureID server ■ RADIUS This guide does not address authentication. For authentication concepts and configuration, see the *JUNOS Software Security Configuration Guide*.
All routing end-point reachability has been established. For details on configuring routing, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Overall IPSec VPN Design Considerations

Table 12 addresses factors to consider when designing your IPSec VPNs.

Table 12: IPSec VPN Deployment Considerations

Deployment Considerations	Decision Process
What types of VPNs do you require?	■ Consider that for policy-based IPSec VPNs, new Security Associations (SAs) are created each time a policy allows traffic to use that VPN. (SAs consume resources.) ■ Consider that for route-based IPSec VPNs, the route directs traffic to the VPN, based on the destination address, and only one pair of SAs is used for all traffic.
What keying method do you plan to use for each site?	■ Internet Key Exchange (IKE) ■ manual
Do you plan to use class of service (CoS) and traffic engineering?	See the *JUNOS Software Interfaces and Routing Configuration Guide* for details.

IPSec VPN Planning Considerations

Table 13 identifies some of the information to take into account in planning your IPSec VPN deployment.

Table 13: Assessing Your VPN Requirements

VPN Design Considerations	Decision Process
How many VPNs do you require?	Determine how many users need to communicate with headquarters across the VPN. ■ Take the worst-case scenario of all users active concurrently. ■ Consider the traffic throughput on a normal daily and hourly basis. Some companies deploy primary and secondary IPSec VPNs and use the backup for load balancing.
What IPSec VPN features do you plan to configure?	See “Implementing Policy-based IPSec VPNs Branch Offices” on page 137 for a description of the supported IPSec VPN features.
What kinds of applications and services do you plan to use? What sizes of packets will traverse the VPN? Will the traffic be mixed? Do you want the VPN to simply pass packets or participate in the network?	This information should be available in your network traffic philosophy document, if you have created one. For details on creating a network traffic philosophy, “Developing a Network Profile” on page 20.

Table 13: Assessing Your VPN Requirements (continued)

VPN Design Considerations	Decision Process
Will the JUNOS software with enhanced services router also be configured for other security context features?	If you are also designing the system to include a firewall, take into account the cost in CPU and memory usage of the features used for the firewall.
What routing protocols do you plan to configure, if any?	Consider the overhead cost of routing protocols in assessing performance and capacity requirement for your configuration.

Capacity Planning and Performance Considerations

Planning for performance takes into account use of system resources for all features that are part of your deployment solution as well as your expected traffic types and throughput on a regular basis.

There are a number of things that you can do to improve scalability and performance. Where possible, use subnets to allow for summarization of configuration information. Address summarization conserves router resources and makes routing tables smaller in size. It also improves manageability.

Table 14 identifies some of the factors that can affect your IPSec VPN performance.

Table 14: Capacity and Performance Planning

Planning and Performance Considerations	
Assess the capacity of the device as it pertains to the features that you will use:	
Which JUNOS software with enhanced services Services J-Series Services Router will you deploy at the branch office?	■ What is the maximum number of tunnel interfaces supported on the router? ■ For policy-based IPSec VPNs, what is the maximum number of policies support on the router, and how many policies do you plan to configure for the VPN? ■ For route-based IPSec VPNS, what is the maximum number of routes configurable in the router's routing table? Consider these factors in relation to your complete solution. For details on specific Services Routers, refer to the *JUNOS Software with Enhanced Services Hardware Guide* or the product data sheet.
Assess the effect of the aggregation system on the IPSec VPN deployment:	
What system is used as the aggregation router at the central site? ■ What is the configuration and performance of the head-end router at the central site? ■ How many IPSec VPNs from all branch offices will be terminated at the aggregation router? ■ What applications will run on that router concurrent with termination of the IPSec VPNs?	Consider that the head-end aggregation device can be a main determining factor in how many VPNs you can deploy collectively from outlying offices. Consider the CPU usage of the head-end device, taking into account the desired number of VPNs and traffic profile to determine if performance of that system would be negatively impacted, therefore negatively impacting VPN performance.

IPSec VPN Concepts

This section explains at a high level some of the concepts underlying IPSec VPNs. If you are familiar with IPSec VPN concepts, you can skip this section.

About IPSec

Internet Security Protocol (IPSec) is a suite of related protocols for cryptographically securing communications at the IP packet layer.

IPSec virtual private networks (VPNs) are logical secure connections between gateways on two local area networks (LANs). They also support users working at home or traveling who require remote access connections to a corporate site. In this case, they are considered end-to-site VPNs. The security characteristics negotiated and agreed upon by the IPSec VPN endpoints—either peer gateways or client software and a gateway—protect data transmitted between them as it transits the Internet, which is otherwise unsecure. Traffic that flows between these endpoints passes through shared resources such as routers, switches, and other network equipment that make up the public WAN.

Policy-based IPSec VPNs

For policy-based IPSec VPNs, a policy specifies as its action the IPSec VPN tunnel to be used for transit traffic that meets the policy's match criteria. An IPSec VPN is configured independent of a policy statement. The policy statement refers to the VPN by name to specify the traffic that is allowed access to the tunnel. For policy-based IPSec VPNs, a new tunnel is generated for each flow of traffic that matches the policy. For example, if a policy contains a group source address and a group destination address, whenever one of the users belonging to the address set attempts to communicate with any one of the hosts specified as the destination address, a new tunnel is negotiated and established. Because each tunnel requires its own negotiation process and separate pair of security associations (SAs), use of policy-based IPSec VPNs can be more expensive in terms of resources than route-based IPSec VPNs.

Route-based IPSec VPNs

With a route-based approach to IPSec VPNs, you can configure dozens of policies to regulate traffic flowing through a single VPN tunnel between two sites, and there is just one set of IKE and IPSec SAs at work.

Unlike policy-based IPSec VPNs, for route-based IPSec VPNs, a policy refers to a destination address not an IPSec VPN tunnel. When the JUNOS software with enhanced services looks up a route to find the interface to use to send traffic to the packet's destination address, it finds a route via a secure tunnel interface (st-x). The tunnel interface is bound to a specific IPSec VPN tunnel, and the traffic is routed to the tunnel if the policy action is permit. The policy dictates if a packet from a source address containing a payload of a certain application type (or service type) is to be delivered to the destination secure tunnel interface and on through the IPSec VPN.

Site-to-Site IPSec VPNs

For both policy-based IPSec VPNs and route-based IPSec VPNs, JUNOS software with enhanced services supports the most common physical IPSec VPN topology called a site-to-site VPN. A site-to-site IPSec VPN connects two peers across a secure connection referred to as a tunnel.

Hub-and-Spoke IPSec VPNs

For route-based IPSec VPNs, JUNOS software with enhanced services supports a topology called hub-and-spoke. In this case, you can create two IPSec VPN tunnels that terminate at a gateway and set up a pair of routes so that the router directs traffic exiting from one tunnel to the other tunnel. If both tunnels are contained within a single zone, you do not need to create a policy to permit the traffic to pass from one tunnel to the other. You only need to define the routes.

Dynamic Endpoint Remote Access IPSec VPNs

For dynamic endpoint remote access IPSec VPNS, the interface used for the IPSec VPN on one peer gateway has a dynamically assigned IP address, and the outgoing interface used for the other peer has a static IP address. For a gateway to reach its peer, the peer must have a static IP address. For remote access IPSec VPNs, only users behind the gateway with the dynamically assigned IP address can initiate the establishment of IPSec VPN tunnels because only their gateway's peer has a static IP address. To support remote access IPSec VPNs, the configuration for the gateway with the static IP address must specify a dynamic address for its peer. In all other respects, remote access IPSec VPNs are the same as basic site-to-site IPSec VPNs.

Remote Access Connections

Remote access connections provide IPSec VPN support for users working at home or traveling. Using IPSec client software, a user can connect to a gateway at a corporate site and gain access to resources without the requirement of a peer gateway on the user side. This topology is sometimes referred to as an end-to-site tunnel.

IPSec Security Associations

A security association (SA) is a data structure that maintains information that constitutes an agreement between the participants in an IPSec VPN regarding the methods and parameters to use in securing communication between themselves and transit data across the Internet.

Each IPSec tunnel consists of a pair of SAs—an Internet Key Exchange (IKE) SA and an IPSec SA—at each end of the tunnel. These SAs are generated from the two phases of cryptography negotiation between the VPN peers: Phase 1 and Phase 2.

An SA is uniquely identified by a Security Parameter Index (SPI) value. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, keys, the protocol mode, the key management method, and the lifetime for the SA.

For policy-based IPSec VPNs, for each tunnel that is generated, a new set of SAs is created. For example, as part of the IPSec VPN policy statement, you can specify groups of source addresses, destination addresses, and applications, rather than only one of each. If you create an IPSec VPN policy statement that includes a group of users (an address set, for example) as its source address, a separate tunnel is generated whenever one of the users belonging to the address set attempts to communicate with any one of the hosts specified as the destination address. Because an IPSec VPN policy statement can include many source addresses, destination addresses, and applications, you can configure a single policy-based IPSec VPN that can result in generation of many tunnels.

When you consider the effect of your configuration as a whole on capacity and performance, take into account the number of tunnels that your policy configurations can generate, considering that each pair of SAs consumes memory.

IKE and IPSec

An IPSec VPN is created as a result of two phases of negotiation of security parameters between the participants. (The participants can be two peer gateways or, for a remote access connection, a gateway and a user's client software.)

You can configure parameters required for the two phases of negotiation using a method called Internet Key Exchange (IKE), or you can configure them manually. Use of IKE is the preferred method. When IKE is used, IPSec key data is automatically derived and managed based on negotiated information configured at the peer gateways.

As was mentioned previously, each of the phases of negotiation culminates in generation of an SA at each peer gateway. The SA generated (at either end) as a result of the first phase, referred to as IKE Phase 1, defines and establishes a secure channel that allows the peers to establish a trusted relationship to be used for their Phase 2 negotiations. For example, during IKE Phase 1, the peers negotiate the kind of encryption to be used to secure the messages that they will exchange in Phase 2.

The IPSec SA generated during Phase 2 contains the parameters that specify the cryptography to be used to secure transit traffic exchanged between the two peers the user data that transits the IPSec VPN. During Phase 2, the two peers— exchange messages to negotiate these parameters. The outcome of this negotiation is an IPSec tunnel. The Phase 2 IPSec SA effectively defines the VPN tunnel.

The Phase 1 IKE SAs manage and initiate the IPSec Phase 2 process with activities that include generation of keying material for the IPSec process, initiation of the IPSec negotiation process, and management of the Phase 2 IPSec SAs. In effect, the life of a Phase 2 IPSec SA is tied to the Phase 1 IKE SA that generated it.

Chapter 4

Implementing Firewall Deployments for Branch Offices

This chapter explains how to implement a JUNOS software with enhanced services stateful firewall for a mid-size branch office of an enterprise corporation. The firewall implementation is based on and follows the design principles explained in “Designing Firewalls for Branch Offices” on page 11.

This scenario describes the firewall deployment for a site called the Albuquerque branch of the New Bank of the Southwest. It is intended to convey a general idea of how to deploy the router in an environment representative of a bank branch and does not match any one, real branch topology. The implementation shows the configuration for a representative group of employees belonging to various departments, not for all employees of the branch.

The example deployment used in this chapter also includes two route-based IPSec VPNs. For more information, see “Implementing Firewall Deployments for Branch Offices” on page 37.

The chapter contains the following sections:

- Configuring the Firewall to Protect Against Penetration Attacks: Firewall and Network DoS Attacks on page 91
- Preventing Operating System DoS Attacks on page 100
- Monitoring for Attacks on page 102
- Testing the Firewall on page 102

In this scenario, the users have already been authenticated.

About the Branch Office Firewall Deployment

The Albuquerque branch of the New Bank of the Southwest is a mid-size branch with two hundred employees. Because of its size, the network architect is deploying the firewall in a JUNOS software with enhanced services J4350 Services Router.

Figure 3 shows the topology for the Albuquerque branch office, including the firewall zones defined for it by the network architect. See “Configuring Zone Services and Protocols” on page 52 for a description of the branch zones.

In addition to regular checking and savings account banking services, the branch offers its customers mortgage and business loan services, insurance services, credit card services, and investment services.

Figure 3: Albuquerque Branch Deployment Topology

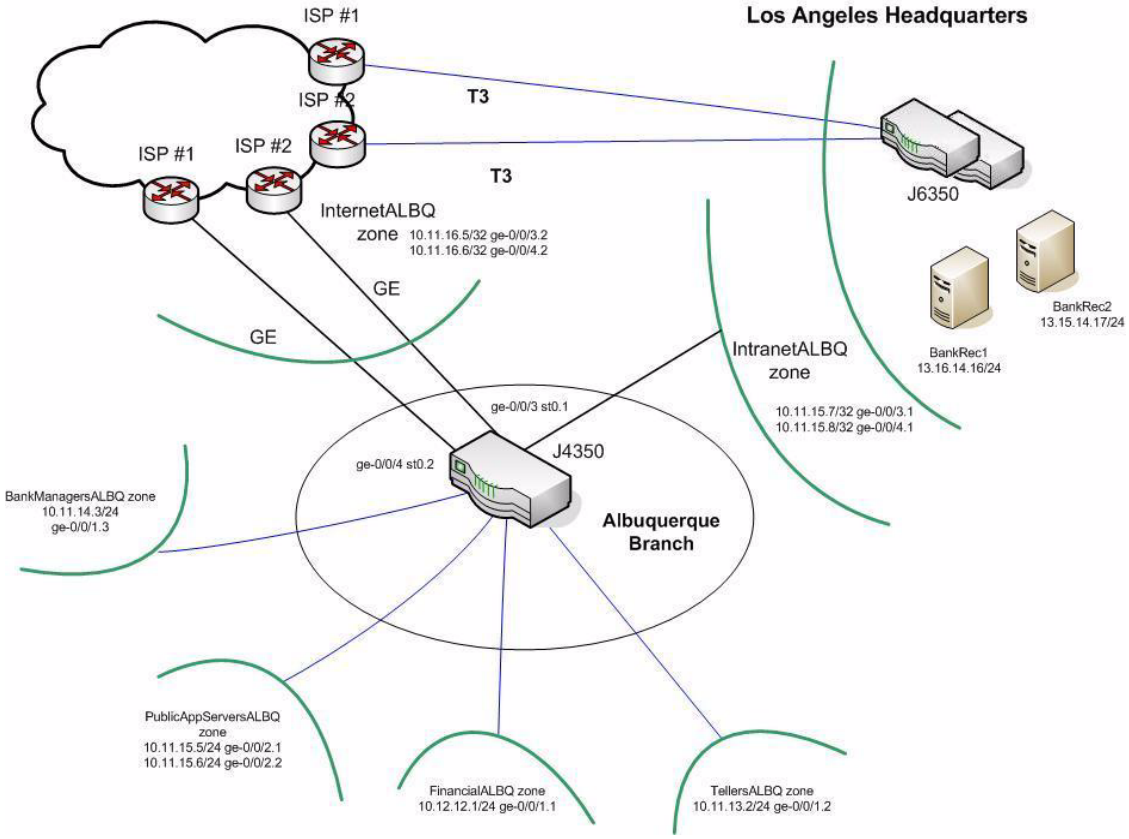

Tasks for Implementing a Firewall Deployment

Table 15 shows the overall tasks needed to implement a firewall deployment.

Table 15: Steps for Implementing a Firewall

Task	Instructions
Configure interfaces and zones.	See: ■ "Configuring Interfaces and Assigning Them to Zones" on page 49. ■ "Configuring TCP Reset" on page 52. ■ "Configuring Allowed Inbound Services" on page 52.
Configure routes.	See: ■ "Configuring Virtual Routing Instances" on page 55. ■ Table 20, "Virtual Routing Instances Definition and Interface Configuration Statement," on page 57.
Configure policies.	See: ■ "Creating Policies" on page 64. ■ "Configuring Address Books" on page 66. ■ "Configuring Applications and Application Sets" on page 69. ■ "Configuring Policies" on page 71. ■ "Configuring Schedulers" on page 76.
Configure screens.	See: ■ "Configuring Screens for the Branch's Zones" on page 79. ■ "Preventing Reconnaissance Attacks" on page 82. ■ "Configuring the Firewall to Protect Against Penetration Attacks: Firewall and Network DoS Attacks" on page 91. ■ "Preventing Operating System DoS Attacks" on page 100.

Albuquerque Branch Security Zones

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. You bind interfaces to zones to associate segments of a network with them. You also configure screens and address books for zones:

- A zone's address book identifies all hosts that belong to the zone so that those hosts are reachable and able to be referred to in policies.

- A zone's screens filter traffic before it leaves the zone. Identifying and stopping attacks at the source ensures that the traffic does not reach hosts in other zones.

Zones are specified in policies. Policies control user access to resources in their own or a different zone. A policy specifies the source zone and the destination zone.

Figure 3 on page 39 shows the zones configured for the example scenario explored in this chapter.

You create zones for segments of your network. You can base zones logically on department, job function, or services. You can devise your own method for segmenting groups of users and hosts into zones. For example, you might create a zone containing all hosts used by tellers and another containing servers. To create a zone, you add interfaces for segments of the network to it. Hosts belonging to the zone must be on one of the network segments whose interface is bound to the zone.

To give employees access to the resources they require located at the branch and at headquarters, and Internet access, if appropriate, the network architect for the Albuquerque branch has defined the following zones:

- A security zone called FinancialALBQ. This zone includes bank officers who manage special services, including mortgage and personal loan services, insurance services, and credit card services. Members of this zone have access to both servers in the PublicAppServersALBQ zone, and they have access to one of the bank records servers, BankRec2, at corporate headquarters in the IntranetALBQ zone. They are also given Internet access via the InternetALBQ zone.

- A security zone called TellersALBQ. This zone includes all tellers who work at the branch, including part-time tellers and trainees. Members of this zone have access to the BankRec1 database server at corporate headquarters in the IntranetALBQ zone. The BankRec1 server contains customer checking and savings information. They also have access to the two servers in the PublicAppServersALBQ zone, which deliver Web, mail, and DNS services. They are not granted Internet access.

- A security zone called BankManagersALBQ. This zone includes bank officers who manage the day-to-day activities and personnel at the branch office. Members of this zone have access to both bank records servers in the IntranetALBQ zone at corporate headquarters and both servers in the PublicAppServersALBQ zone. They also have Internet access.

- A security zone called PublicAppServersALBQ. This zone contains two public servers that deliver e-mail, other Web services, and DNS support.

- A security zone called IntranetALBQ. This zone includes the IPSec VPN tunnels that connect the branch office to corporate headquarters.

  There are two bank record servers at headquarters—BankRec1, whose IP address is 13.16.14.16/32 and BankRec2 whose IP address is 13.16.14.17/32. Users at the branch must access these servers.

  - BankRec1 contains a database that delivers customer checking and savings records.
  - BankRec2 contains a database that delivers records for customers who participate in special services, such as loans, credit card services, and financial investment.
- A security zone called InternetALBQ, which allows for access to the Internet.

Firewall Configuration Deployment Commands

This section shows the complete set of statements used to configure the firewall for the Albuqueque branch office described in "About the Branch Office Firewall Deployment" on page 39 and "Albuquerque Branch Security Zones" on page 40.

The complete set of statements is shown before the sections that describe how to use them to configure parts of the firewall.

NOTE: This deployment does not include an authentication component. See the *JUNOS Software Security Configuration Guide* for information about user and firewall authentication.

- In the examples, complete set commands are used. You can also enter these statements at the [edit security] and security hierarchy level, in which case, you do not need to specify the word *security* as part of a statement.
- For each group of related statements, the left margin identifies the part of the firewall configuration that the statements pertain to.
- To go to the section providing details for a part of the configuration, click the link in the left margin.

Zones, their interfaces, and VLAN tagging. (See "Configuring Interfaces and Assigning Them to Zones" on page 49.)

```
user@host# set interfaces ge-0/0/1 unit 1 family inet address 10.12.12.1/24
user@host# set security zones security-zone FinancialALBQ interfaces ge-0/0/1.1

user@host# set interfaces ge-0/0/1 unit 2 family inet address 10.11.13.2/24
user@host# set security zones security-zone TellersALBQ interfaces ge-0/0/1.2

user@host# set interfaces ge-0/0/1 unit 3 family inet address 10.11.14.3/24
user@host# set security zones security-zone BankManagersALBQ interfaces
ge-0/0/1.3

user@host# set interfaces ge-0/0/2 unit 1 family inet address 10.11.15.5/32
user@host# set security zones security-zone PublicAppServersALBQ interfaces
ge-0/0/2.1
user@host# set interfaces ge-0/0/2 unit 2 family inet address 10.11.15.6/32
user@host# set security zones security-zone PublicAppServersALBQ interfaces
ge-0/0/2.2

user@host# set interfaces ge-0/0/3 vlan-tagging
user@host# set interfaces ge-0/0/3 unit 1 family inet address 10.11.15.7/32
user@host# set interfaces ge-0/0/3 unit 1 vlan-id 299
user@host# set security zones security-zone IntranetALBQ interfaces ge-0/0/3.1
```

```
user@host# set interfaces ge-0/0/4 vlan-tagging
user@host# set interfaces ge-0/0/4 unit 1 family inet address 10.11.15.8/32
user@host# set interfaces ge-0/0/4 unit 1 vlan-id 300
user@host# set security zones security-zone IntranetALBQ interfaces ge-0/0/4.1

user@host# set interfaces ge-0/0/3 unit 2 family inet address 11.11.16.5/32
user@host set interfaces ge-0/0/3 unit 2 vlan-id 398
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2

user@host# set interfaces ge-0/0/4 unit 2 family inet address 11.11.16.6/32
user@host# set interfaces ge-0/0/4 unit 2 vlan-id 399
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/4.2
```

TCP reset. (See "Configuring TCP Reset" on page 52.)

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

user@host# set security zones security-zone TellersALBQ tcp-rst

user@host# set security zones security-zone PublicAppServersALBQ tcp-rst

user@host# set security zones security-zone BankManagersALBQ tcp-rst

user@host# set security zones security-zone IntranetALBQ tcp-rst

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

Host inbound services. (See "Configuring Allowed Inbound Services" on page 52.)

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

user@host# set security zones security-zone BankManagersALBQ interfaces
ge-0/0/1.3 host-inbound-traffic system-services dhcp

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

user@host# set security zones security-zone PublicAppServersALBQ
host-inbound-traffic system-services traceroute

user@host# set security zones security-zone PublicAppServersALBQ
host-inbound-traffic system-services snmp

user@host# set security zones security-zone PublicAppServersALBQ
host-inbound-traffic system-services snmp-trap

user@host# set security zones security-zone TellersALBQ host-inbound-traffic
system-services traceroute

user@host# set security zones security-zone TellersALBQ host-inbound-traffic
system-services snmp
user@host# set security zones security-zone TellersALBQ host-inbound-traffic
system-services snmp-trap

user@host# set security zones security-zone InternetALBQ host-inbound-traffic
system-services telnet except

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

Virtual routing instances. (See "Configuring Virtual Routing Instances" on page 55.)	**user@host# set routing-instances IntranetALBQ-vr instance-type virtual-router interface ge-0/0/3.1** **user@host# set routing-instances IntranetALBQ-vr instance-type virtual-router interface ge-0/0/4.1** **user@host# set routing-instances InternetALBQ-vr instance-type virtual-router interface ge-0/0/3.2** **user@host# set routing-instances InternetALBQ-vr instance-type virtual-router interface ge-0/0/4.2** **user@host# set routing-instances FinancialALBQ-vr instance-type virtual-router interface ge-0/0/1.1** **user@host# set routing-instances TellersALBQ-vr instance-type virtual-router interface ge-0/0/1.2** **user@host# set routing-instances BankManagersALBQ-vr instance-type virtual-router interface ge-0/0/1.3** **user@host# set routing-instances PublicAppServersALBQ-vr instance-type virtual-router interface ge-0/0/2.1** **user@host# set routing-instances PublicAppServersALBQ-vr instance-type virtual-router interface ge-0/0/2.2**
Static routes. (See "Configuring Static Routes" on page 57.)	**user@host# set routing-instances IntranetALBQ-vr routing-options static route 13.16.14.16 next-hop st0.1 preference 5** **user@host# set routing-instances IntranetALBQ-vr routing-options static route 13.16.14.16 next-hop st0.2 preference 7** **user@host# set routing-instances IntranetALBQ-vr routing-options static route 13.16.14.17 next-hop st0.1 preference 5** **user@host# set routing-instances IntranetALBQ-vr routing-options static route 13.16.14.17/24 next-hop st0.2 preference 7** **user@host# set routing-instances InternetALBQ-vr routing-options static route 9.10.14.5 next-hop 10.11.16.5/32 preference 5** **user@host# set routing-instances InternetALBQ-vr routing-options static route 11.12.15.7 next-hop 10.11.16.6/32 preference 7**
Routing tables. (See "Installing Routes into More Than One Routing Table" on page 61.)	**user@host# set routing-options rib-groups intranet-rg import-rib IntranetALBQ-vr.inet.0** **user@host# set routing-options rib-groups intranet-rg import-rib TellersALBQ-vr.inet.0** **user@host# set routing-options rib-groups intranet-rg import-rib FinancialALBQ-vr.inet.0** **user@host# set routing-options rib-groups intranet-rg import-rib BankManagersALBQ-vr.inet.0**

```
user@host# set routing-options rib-groups intranet-rg import-rib
PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
TellersALBQ-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
FinancialALBQ-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
BankManagersALBQ-vr.inet.0

user@host# set routing-options rib-groups tellers-rg import-rib TellersALBQ-vr.inet.0

user@host# set routing-options rib-groups tellers-rg import-rib
PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups tellers-rg import-rib IntranetALBQ-vr.inet.0

user@host# set routing-options rib-groups bankManagers-rg import-rib
BankManagersALBQ-vr.inet.0

user@host# set routing-options rib-groups bankManagers-rg import-rib
PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups bankManagers-rg import-rib
IntranetALBQ-vr.inet.0

user@host# set routing-options rib-groups bankManagers-rg import-rib
InternetALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib
FinancialALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib
PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib
IntranetALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib
InternetALBQ-vr.inet.0
```

Policies. (See "Creating Policies" on page 64.)

```
user@host# set security policies from-zone BankManagersALBQ to-zone
PublicAppServersALBQ policy BankMantoWebApps match source-address bankMan
destination-address PubApps2Web application ssh

user@host# set security policies from-zone BankManagersALBQ to-zone
PublicAppServersALBQ policy BankMantoWebApps then permit
```

Address books. (See "Configuring Address Books" on page 66.)

```
user@host# set security zones security-zone IntranetALBQ address-book address
BankRec1 13.16.14.16/32
user@host# set security zones security-zone IntranetALBQ address-book address
BankRec2 13.15.14.17/32

user@host# set security zones security-zone IntranetALBQ address-book address-set
HeadSerAll address BankRec1

user@host# set security zones security-zone IntranetALBQ address-book address-set
HeadSerAll address BankRec2

user@host# set security zones security-zone TellersALBQ address-book address
tellers 10.11.13.0/24

user@host# set security zones security-zone FinancialALBQ address-book address
finManagers 10.12.12.0/24

user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan 10.11.14.0/24

user@host# set security zones security-zone PublicAppServersALBQ address-book
address PubApps1Mail 11.10.15.5/32

user@host# set security zones security-zone PublicAppServersALBQ address-book
address PubApps2Web 10.11.15.6/32

user@host# set security zones security-zone PublicAppServersALBQ address-book
address-set PublicAppBoth address PubApps1

user@host# set security zones security-zone PublicAppServersALBQ address-book
address-set PublicAppBoth address PubApps2
```

Applications and application sets. (See "Configuring Applications and Application Sets" on page 69.)

```
user@host# set applications application-set BankAppSet application ssh
user@host# set applications application-set BankAppSet application telnet
user@host# set applications application-set BankAppSet application custApp

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
```

Policies. (See "Configuring Policies" on page 71.)

```
user@host# set security policies default-policy deny-all

user@host# set security policies from-zone TellersALBQ to-zone IntranetALBQ policy
TellCustRec match source-address tellers destination-address BankRec1 application
custRecApp

user@host# set security policies from-zone TellersALBQ to-zone
PublicAppServersALBQ policy TellMailOut match source-address tellers
destination-address PubApps1Mail application [POP3 smtp]

user@host# set security policies from-zone TellersALBQ to-zone
PublicAppServersALBQ policy TellMailOut then permit

user@host# set security policies from-zone FinancialALBQ to-zone IntranetALBQ
policy FinRecPolOut match source-address finManagers destination-address
BankRec2 application-set BankAppSet

user@host# set security policies from-zone FinancialALBQ to-zone Intranet policy
FinRecPolOut then permit

user@host# set security policies from-zone FinancialALBQ to-zone
PublicAppServersALBQ policy FinMailWebOut match source-address finManagers
destination-address PublicAppBoth application [smtp POP3]

user@host# set security policies from-zone FinancialALBQ to-zone
PublicAppServersALBQ policy FinMailWebOut then permit

user@host# set security policies from-zone BankManagersALBQ to-zone
IntranetALBQ policy BankManHQServOut match source-address bankMan
destination-address PublicAppBoth application-set BankAppSet

user@host# set security policies from-zone BankManagersALBQ to-zone
IntranetALBQ policy BankManHQServOut then permit

user@host# set security policies from-zone BankManagersALBQ to-zone
PublicAppServersALBQ policy BankManMailWebOut match source-address bankMan
destination-address all application-set WebMailApps

user@host# set security policies from-zone BankManagersALBQ to-zone
PublicAppServersALBQ policy BankManMailWebOut then permit

user@host# set security policies from-zone PublicAppServersALBQ to-zone
BankManagersALBQ policy ManRecServIn match source-address PublicAppBoth
destination-address bankMan application-set CustAppMessage

user@host# set security policies from-zone PublicAppServersALBQ to-zone
BankManagersALBQ policy ManRecServIn then permit

user@host# set security policies from-zone financialALBQ to-zone InternetALBQ
policy InternetAccess match source-address finManagers
destination-address all application-set all

user@host# set security policies from-zone financialALBQ to-zone InternetALBQ
policy InternetAccess match then permit
```

user@host# set security policies from-zone BankManagersALBQ to-zone InternetALBQ policy InternetAccess match source-address bankMan destination-address all application-set all

user@host# set security policies from-zone BankManagersALBQ to-zone InternetALBQ
policy InternetAccess match then permit

Schedulers. (See "Configuring Schedulers" on page 76.)

user@host# set schedulers scheduler RegularHours start-date 2007-10-01.07:00 stop-date 2008-10-01.18:00

user@host# set schedulers scheduler RegularHours saturday exclude

user@host# set schedulers scheduler RegularHours sunday exclude

user@host set security policies from-zone TellersALBQ to-zone IntranetALBQ policy TellCustRec scheduler-name RegularHours
user@host# set schedulers scheduler SaturdayHrs saturday start-time 12:00 stop-time 17:00

user@host# set schedulers scheduler SundayHrs sunday start-time 12:00 stop-time 17:00

user@host# set security policies from-zone BankManagersALBQ to-zone IntranetALBQ policy BankManRecServOut scheduler-name SaturdayHrs

user@host# set security policies from-zone BankManagersALBQ to-zone IntranetALBQ policy BankManRecServOut scheduler-name SundayHrs

Screens: network reconnaissance. (See "Preventing Reconnaissance Attacks" on page 82.)

user@host# set security screen ReconnScrn icmp ip-sweep threshold 1000

user@host# set security screen ReconnScrn tcp fin-no-ack port-scan threshold 1000

user@host# set security screen ReconnScrn tcp syn-fin tcp-no-flag syn-frag

user@host# set security screen ReconnScrn ip spoofing loose-source-route-option strict-source-route-option

user@host# set security screen ReconnScrn ip record-route-option timestamp-option bad-option security-option stream-option block-frag unknown-protocol

user@host# set security screen ReconnScrn icmp fragment large

Screens: firewall and network DoS attack. (See "Configuring the Firewall to Protect Against Penetration Attacks: Firewall and Network DoS Attacks" on page 91.)

user@host# set security screen FireDoSScrn limit-session source-ip-based 100 destination-ip-based 100

user@host# set security screen FireDoSScrn tcp syn-ack-ack-proxy threshold 1000

user@host# set security screen NetDoSScrn tcp syn-flood attack-threshold 1000 alarm-threshold 512 source-threshold 4000 destination-threshold 4000 timeout 30 queue-size 1024

Screens: ICMP. (See "Preventing ICMP Flood Attacks" on page 98.)

user@host# set security screen ICMPfloodScrn icmp flood threshold 1500

Screens: UDP. (See "Preventing UDP Flood Attacks" on page 99.)	*user@host# set security screen UDPfloodScrn udp flood threshold 1500*
Screens: land attack. (See "Preventing Land Attacks" on page 100.)	**user@host# set security screen LandAttackScrn tcp land**
Screens: OS attacks. (See "Preventing Operating System DoS Attacks" on page 100.)	**user@host# set security screen OSDoSAttacksScrn icmp ping-death** **user@host# set security screen OSDoSAttacksScrn ip tear-drop** **user@host# set security screen OSDoSAttacksScrn tcp winnuke**

Configuring Interfaces and Assigning Them to Zones

For the branch zones, the network architect has determined that the security administrator should configure the interfaces and subinterfaces (logical interfaces, or logical units) identified in this section. Subinterfaces borrow the bandwidth they require from the physical interface from which they are derived. This section shows how to define the router's interfaces, assign IP addresses to them, and associate an interface with a zone.

About the JUNOS Software with Enhanced Services J-series Services Routers Interfaces

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 router. In addition, each router 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.

J-series 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 1of 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 J-series interfaces and interface names, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Assigning Security Zones to Interfaces

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

- To configure the interface and its IP address for the FinancialALBQ zone, enter the following statement in configuration mode:

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

- To configure the FinancialALBQ zone and assign the interface to it, enter the following statement in configuration mode:

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

- To configure the interface and its IP address for the TellersALBQ zone, enter the following statement in configuration mode:

  user@host# set interfaces ge-0/0/1 unit 2 family inet address 10.11.13.2/24

- To configure the TellersALBQ zone and assign the interface to it, enter the following statement in configuration mode:

  user@host# set security zones security-zone TellersALBQ interfaces ge-0/0/1.2

- To configure the interface and its IP address for the BankManagersALBQ zone, enter the following statement in configuration mode:

  user@host# set interfaces ge-0/0/1 unit 3 family inet address 10.11.14.3/24

- To configure the BankManagersALBQ zone and assign the Interface to it, enter the following statement in configuration mode:

  user@host# set security zones security-zone BankManagersALBQ interfaces ge-0/0/1.3

- To configure the interfaces and their IP addresses for the PublicAppServersALBQ zone, enter the following statements in configuration mode:

  user@host# set interfaces ge-0/0/2 unit 1 family inet address 10.11.15.5/32

  user@host# set interfaces ge-0/0/2 unit 2 family inet address 10.11.15.6/32

- To configure the PublicAppServersALBQ zone and assign the interfaces to it, enter the following statements in configuration mode:

  user@host# set security zones security-zone PublicAppServersALBQ interfaces ge-0/0/2.1

  user@host# set security zones security-zone PublicAppServersALBQ interfaces ge-0/0/2.2

- To enable VLAN tagging on the two interfaces shared for IPSec VPNs and Internet connections, enter the following statements in configuration mode:

  user@host# set interfaces ge-0/0/3 vlan-tagging

  user@host# set interfaces ge-0/0/4 vlan-tagging

- To configure the interfaces and their IP addresses for the IntranetALBQ zone, which allows for two IPSec VPN connections to headquarters, enter the following statements in configuration mode:

  user@host# set interfaces ge-0/0/3 unit 1 family inet address 10.11.15.7/32
  user@host# set interfaces ge-0/0/3 unit 1 vlan-id 299

  user@host# set interfaces ge-0/0/4 unit 1 family inet address 10.11.15.8/32
  user@host# set interfaces ge-0/0/4 unit 1 vlan-id 300

- To add the interfaces to the IntranetALBQ zone, enter the following statements in configuration mode:

  user@host# set security zones security-zone IntranetALBQ interfaces ge-0/0/3.1
  user@host# set security zones security-zone IntranetALBQ interfaces ge-0/0/4.1

- To configure the interfaces and their IP address for the InternetALBQ zone, enter the following statements in configuration mode:

  user@host# set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32
  user@host set interfaces ge-0/0/3 unit 2 vlan-id 398
  user@host# set interfaces ge-0/0/4 unit 2 family inet address 11.11.16.6/32
  user@host# set interfaces ge-0/0/4 unit 2 vlan-id 399

- To configure the zones for the InternetALBQ zone and assign the interfaces to them, enter the following statements in configuration mode:

  user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2
  user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/4.2

Table 16 gives the interface configuration syntax.

Table 16: Interfaces Configuration Statement

Interface Configuration	Syntax
Configures an interface and its IP address to be used for a zone. When you configure a physical interface, you create a (virtual) logical interface for it to identify its connection.	`set interfaces <port-type> unit <logical-interface number> family inet address <IP-address>` ■ `port-type`: Type of port and its location, that is, the number of the port on the physical interface module (PIM). ■ `logical-interface number`: A value from 0 through 16384. If no logical interface number is specified, the default number of unit 0 is used, but you must explicitly configure it. A logical interface is equivalent to a subinterface. ■ `IP-address`: Address to assign to the interface.

Configuring Zone Services and Protocols

This section explains how to configure some zone-level parameters. You can configure parameters to control the kind of traffic that is allowed inbound for the interfaces of a zone, or a particular interface. By controlling the allowed inbound traffic and limiting it to specific applications and protocols, you can better protect the device from attacks launched against it from a system that is directly connected to one of its interfaces.

Configuring TCP Reset

The network architect has determined that for all zones, the TCP reset is to be set. 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.

To set the parameter for the zones, enter the following statements in configuration mode:

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

user@host# set security zones security-zone TellersALBQ tcp-rst

user@host# set security zones security-zone PublicAppServersALBQ tcp-rst

user@host# set security zones security-zone BankManagersALBQ tcp-rst

user@host# set security zones security-zone IntranetALBQ tcp-rst

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

Table 17 gives an overview of the TCP Reset parameter used for the deployment scenario and the statement syntax.

Table 17: TCP Reset Configuration Statement

TCP Reset	Syntax
Enables the TCPO Reset 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.	`set security zones security-zone <zone-name> tcp-rst` ■ `zone-name`: Name of the zone for which you are enabling the feature.

Configuring Allowed Inbound Services

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

- You can configure parameters at the zone level to 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.

This feature allows you to protect the device from attacks launched against it 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 attempt to log into the device because you would not want them connecting to your router.

NOTE: You can also configure an interface to allow for dynamic routing protocols. These examples do not include specification of protocols because the Albuquerque site uses only static routes and not dynamic routing protocols.

To allow inbound traffic on an interface from a directly connected device, no policy is required; Only inbound service configuration is required.

However, for traffic sent from a directly connected device to cross from one interface to another, a policy is required even if the interfaces are in the same zone. Also, both interfaces must have the inbound service allowed. For example, suppose a user whose system was directly connected to interface 1.3.1.4 in zone A wanted to telnet into interface 2.1.2.4 in zone A. 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 18 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 18: Supported Inbound System Services

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

System services for the BankManagersALBQ zone

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

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

For this zone, DHCP is configured at the interface level. You must configure DHCP at the interface level, not the zone level. For incoming DHCP requests to be recognized, the interface must be known to the server.

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

System services for the FinancialALBQ zone

To configure the FinancialALBQ zone to allow all application services as inbound services, enter the following statements in configuration mode:

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

System services for the PublicAppServersALBQ zone

To configure the PublicAppServersALBQ zone to allow traceroute, SNMP, and SNMP trap traffic as inbound services, enter the following statements in configuration mode:

```
user@host# set security zones security-zone PublicAppServersALBQ
host-inbound-traffic system-services traceroute

user@host# set security zones security-zone PublicAppServersALBQ
host-inbound-traffic system-services snmp

user@host# set security zones security-zone PublicAppServersALBQ
host-inbound-traffic system-services snmp-trap
```

System services for the TellersALBQ zone

To configure the TellersALBQ zone to allow traceroute, SNMP, and SNMP trap traffic as inbound services, enter the following statements in configuration mode:

```
user@host# set security zones security-zone TellersALBQ host-inbound-traffic
system-services traceroute

user@host# set security zones security-zone TellersALBQ host-inbound-traffic
system-services snmp

user@host# set security zones security-zone TellersALBQ host-inbound-traffic
system-services snmp-trap
```

System services for the InternetALBQ zone

To configure the InternetALBQ zone to exclude inbound Telnet traffic, enter the following statement in configuration mode:

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

System services for the IntranetALBQ zone

To configure the IntranetALBQ zone to allow all application services as inbound services, enter the following statements in configuration mode:

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

Table 19 gives an overview of the host inbound traffic parameter and the statement syntax.

Table 19: Inbound Services Configuration Statement

Inbound Services Configuration	Syntax
Configures the interfaces of a zone to allow (or deny) traffic of a specific type from interfaces that are directly connected to the router.	`set security zones security-zone <zone-name> host-inbound-traffic system-services <service-name> [except]` ■ `zone-name`: Zone for which the inbound service is allowed. ■ `service-name`: Name of a predefined or custom service. Custom services must be explicitly configured. ■ `except`: Optional, Excludes traffic from the specified service.

Configuring Routing Instances and Static Routes

This section describes how to configure virtual-router routing instances including their interfaces and static routes on the Albuquerque branch's J4350 Services Router. Routes are required for traffic to traverse the two IPSec VPNs configured for this example and for traffic to be sent from the hosts in the TellersALBQ, FinancialALBQ, and BankManagersALBQ zones at the local branch to the servers in the PublicAppServersALBQ zone. Addresses for directly connected interfaces bound to zones are added to the default routing table for a routing instance automatically.

This section also describes how to create routing table groups known as RIB (routing information base) groups to give a routing instance access to the routing table of another routing instance.

This scenario uses static routes. It does not address dynamic routing protocols. For information describing how to configure dynamic routing protocols, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Configuring Virtual Routing Instances

To configure a routing instance on a JUNOS software with enhanced services J-series Services Router, you create a virtual router (VR) routing instance. You can configure multiple virtual routing instances, each of which creates

- A set of routing tables
- A set of interfaces that belong to these routing tables
- A set of routing options

You must give each routing instance a unique name. If you configure a routing instance with the name IntranetALBQ-vr, its corresponding tables are named accordingly. For example, its unicast IPv4 routing table is called Intranet-vr.inet.0. In this case, all static IPv4 routes that you configure explicitly for the Intranet-vr instance are installed into the IntranetALBQ-vr.inet.0 table.

NOTE: By default, routes are installed into the default routing instance table inet.0 unless you specify a particular routing instance when you configure the route.

When you boot the system, routes for all the directly connected interfaces are added to their respective routing-instance routing tables with a status of Connected.

For the Albuquerque site, you must configure routing instances for each of the zones. Then you can create routing groups for the routing instance of each zone to include the routing tables of other appropriate routing instances. For example, for the TellersALBQ-vr routing instance, you want to add the routing tables of only the IntranetALBQ-vr routing instance and the PublicAppServersALBQ-vr routing instance because their zones are the only ones to which hosts in the TellersALBQ zone are allowed access. If you were to create a single routing table group, the routes for all routing instances (in this case, zones) would be visible to all other routing instances (zones).

To create a virtual router, you specify the following information for it.

- its name
- its type
- the interfaces to be bound to it

You can also specify its static routes. “Table 20 gives an overview of the routing instances definition and syntax.” on page 57 describes how to configure them.

Virtual routing instance for IntranetALBQ traffic

To create the virtual router for routing encrypted traffic to corporate headquarters using one of the route-based IPSec VPNs in the IntranetALBQ zone, enter the following statements in configuration mode:

```
user@host# set routing-instances IntranetALBQ-vr instance-type virtual-router
interface ge-0/0/3.1

user@host# set routing-instances IntranetALBQ-vr instance-type virtual-router
interface ge-0/0/4.1
```

Virtual routing instance for InternetALBQ traffic

To create a virtual router to be used for Internet traffic, enter the following statements in configuration mode:

```
user@host# set routing-instances InternetALBQ-vr instance-type virtual-router
interface ge-0/0/3.2

user@host# set routing-instances InternetALBQ-vr instance-type virtual-router
interface ge-0/0/4.2
```

Virtual routing instance for FinancialALBQ traffic

To create a virtual router to be used for traffic from the FinancialALBQ zone, enter the following statements in configuration mode:

```
user@host# set routing-instances FinancialALBQ-vr instance-type virtual-router
interface ge-0/0/1.1
```

Virtual routing instance for TellersALBQ traffic

To create a virtual router to be used for traffic from the TellersALBQ zone, enter the following statements in configuration mode:

```
user@host# set routing-instances TellersALBQ-vr instance-type virtual-router
interface ge-0/0/1.2
```

Virtual routing instance for BankManagersALBQ traffic

To create a virtual router to be used for traffic from the BankManagersALBQ zone, enter the following statements in configuration mode:

```
user@host# set routing-instances BankManagersALBQ-vr instance-type
virtual-router interface ge-0/0/1.3
```

Virtual routing instance for PublicAppServersALBQ traffic

To create a virtual router to be used for traffic from the PublicAppServersALBQ zone, enter the following statements in configuration mode:

```
user@host# set routing-instances PublicAppServersALBQ-vr instance-type
virtual-router interface ge-0/0/2.1

user@host# set routing-instances PublicAppServersALBQ-vr instance-type
virtual-router interface ge-0/0/2.2
```

Alternatively, you can enter the statements at the Edit routing-instances hierarchy and omit the term routing-instances.

Table 20 gives an overview of the routing instances definition and syntax.

Table 20: Virtual Routing Instances Definition and Interface Configuration Statement

Virtual Routing Instances	Syntax
Defines a routing instance by name. Establishes its type as a "virtual router" (VR), and adds the VR interfaces to its routing tables.	`set routing-instances <routing-instance-name> instance-type (forwarding | l2vpn | no-forwarding | virtual-router | vpls | vrf) interface <interface-name>` ■ `routing-instance-name`: Name of the virtual router. The name is used to refer to the routing instance in configuring its routes and attributes, including ribs (routing tables). ■ `(instance-type)`: One of the types shown in the syntax given in this table. For JUNOS software with enhanced services software, virtual-router routing instances are supported. ■ `interface-name`: The logical, private interface between the provider edge and the customer edge.

Configuring Static Routes

For a network that has few connections to other networks or for networks where inter-network connections are relatively unchanging, it is usually more efficient to define static routes rather than to use dynamic routing protocols. This is the case for the Albuquerque branch deployment.

Here are some characteristics of static routes:

- A static route is a manually configured mapping of an IP network address to a next-hop destination that you define.
  - Static routes are never changed unless you explicitly update them.
  - Static routing avoids the bandwidth cost and route import latency of dynamic routing.
- You explicitly configure static routes for a routing instance by entering them into the virtual router's tables as permanent additions.
  - Static routes are automatically imported into the specific routing table when the system first comes online.
  - A static route is inserted into the virtual router's forwarding table when the next hop associated with it is reachable.
- A routing table can contain more than one static route to the same destination. It is useful to have alternate routes if that the primary route becomes unavailable. Also, you can add the same routes to more than one routing table, because each table is considered unique.

The router determines the zone for a packet based on its destination address. Here is the part of the JUNOS software with enhanced services packet flow process that pertains to the route lookup:

1. When a host sends packets to another host that resides on a different network, each packet header contains the address of the destination host. When the router retrieves a packet from its queue, it polices and filters the packet first, then it checks to determine if the packet belongs to an existing session.

2. If the router does not find a matching session, it applies the firewall screens associated with the source zone to the packet.

3. The router compares the packet's destination address to all addresses contained in its routing table. The system selects the most specific route in the forwarding table to that destination address, and from the selected route entry, determines the next hop to forward the packet to. The route-lookup also identifies the destination zone for the packet, based on the next-hop address.

4. The packet is checked against security policies that you have defined to determine how it is to be treated.

Table 21 gives an overview of the static route basic configuration and its statement syntax.

Table 21: Static Routes Configuration Statement

VR Routing Options: Static Routes Configuration	Syntax

Adds a static route to a virtual router, including next hop and preference.
 	
To create a static route in the routing table, you must, at a minimum, define the route as static and associate a next-hop address with it.
 	
Note: This table gives the rudimentary syntax for creating a static route. There are other attributes that you can specify, as described in the *JUNOS Software Interfaces and Routing Configuration Guide*.
	`set routing-instances <instance-name> routing-options static route <destination-address> next-hop <address-prefix> preference <value>`
	
	■ `instance-name`: Routing instance name, specified when the virtual router was created.
	
	■ `destination-address`: Destination of the static route.
	
	■ `destination-prefix/netmask`: The destination prefix is the network portion of the IP address, and the prefix length is the netmask.
	
	■ `(next-hop)`
	  ■ `address-prefix:` How to reach the destination.
	  ■ `interface-name`—Reach the next-hop router by using the specified interface address.
	
	■ `(preference) value`: A lower number indicates a more preferred route. Range: 1 through 255. Default: 5 (for static routes).

You can specify the address of another routing table as the next hop (next-table), or you can specify that packets for the destination are to be discarded (discard) or rejected (reject). For details, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

You can also add a qualified next hop to a static route.

- You can configure next-hop addresses for a particular route and have them treated differently by defining one or more as a qualified next hop.

- You can set an overall preference for a static route, and specify a different preference for each next hop.

Table 22 gives an overview of the qualified next hop feature and its statement syntax.

Table 22: Static Route Qualified Next Hop Configuration Statement

Static Routes: Qualified Next Hop	Syntax
Specifies the qualified-next hop for a static route, including the preference for it.	`set routing-instances <instance-name> static route <destination-address> qualified-next-hop <address> preference <value>` ■ `instance-name`: Routing instance name, specified when the virtual router was created. ■ `destination-address`: Destination of the static route. ■ destination-prefix/netmask. The destination prefix is the network portion of the IP address, and the prefix length is the destination netmask. ■ `(qualified-next-hop) address:` How to reach the destination. ■ `address-prefix: or` ■ `interface-name`—Reach the next-hop router by using the specified interface address. ■ `(preference) value`: A lower number indicates a more preferred route. Range: 1 through 255. Default: 5 (for static routes).

Configuring Static Routes to Los Angeles Corporate Headquarters

The following statements add static routes to the Albuquerque's IntranetALBQ-vr routing tables to direct all traffic destined to corporate headquarters through one of the IPSec VPNs. Because the primary VPN has the lower preference, it is used first. If the route for the primary VPN is not available, then the backup VPN is used. Each of the VPNs uses a different ISP to provide for failover if one ISP is unavailable.

There are two servers at headquarters: BankRec1 with an IP address of 13.16.14.16, and BankRec2, with an IP address of 13.15.14.17. Traffic is sent to these servers from users at the branch via an IPSec VPN in the IntranetALBQ zone.

Two static routes to BankRec1 server at headquarters

To configure a primary static route to the BankRec1 server (13.16.14.16) at corporate headquarters with a next hop through the primary IPSec VPN tunnel interface (st0.1) and with a preference of 5, enter the following statement in configuration mode:

```
user@host# set routing-instances IntranetALBQ-vr routing-options static route
13.16.14.16 next-hop st0.1 preference 5
```

To configure a secondary static route to the BankRec1 server (13.16.14.16) at corporate headquarters with a next hop through the secondary IPSec VPN tunnel interface (st0.2) and with a preference of 7, enter the following statement in configuration mode:

```
user@host# set routing-instances IntranetALBQ-vr routing-options static route
13.16.14.16 next-hop st0.2 preference 7
```

Two static routes to BankRec2 at headquarters

To configure a primary static route to the BankRec2 server (13.16.14.17) at corporate headquarters with a next hop through the primary IPSec VPN tunnel interface (st0.1) with a preference of 5, enter the following statement in edit routing-instances mode:

```
user@host# set routing-instances IntranetALBQ-vr routing-options static route
13.16.14.17 next-hop st0.1 preference 5
```

To configure a secondary static route to the BankRec2 server (13.15.14.17) at corporate headquarters with a next hop through the secondary IPSec VPN tunnel interface (st0.2) and with a preference of 7, enter the following statement in configuration mode:

```
user@host# set routing-instances IntranetALBQ-vr routing-options static route
13.15.14.17/24 next-hop st0.2 preference 7
```

Configuring Static Routes to the Internet

The following statements add static routes to the Albuquerque's JUNOS software with enhanced services InternetALBQ-vr routing tables for traffic destined to the Internet. A route is included to each of the Internet service providers (ISPs). The IP address for ISP1 is 9.10.14.5, and the IP address for ISP2 is 11.12.15.7.

Because the routes go to different ISPs connected via separate WAN lines, they act independently of each other. The IP addresses of the WAN interfaces used to connect to the ISP's routers are used as the next-hop addresses. (Alternatively, you could specify the interface names as the next hop.)

Creating a static route to connect to ISP1

To configure a static route with preference of 5 to be used to connect to the primary service provider (ISP1) for Internet access, enter the following statement in configuration mode:

```
user@host# set routing-instances InternetALBQ-vr routing-options static route
9.10.14.5 next-hop 10.11.16.5/32 preference 5
```

Creating a static route to connect to ISP2

To configure a static route with preference of 7 to be used to connect to the secondary service provider (ISP2) for Internet access, enter the following statement in configuration mode:

```
user@host# set routing-instances InternetALBQ-vr routing-options static route
11.12.15.7 next-hop 10.11.16.6/32 preference 7
```

Controlling Importation of Static Routes

You can control the importation of different kinds of routes into the routing and forwarding tables in a number of different ways, including whether the route should be retained in the forwarding table after the routing process shuts down (retain), whether the system should reject traffic destined for the route (passive), and whether the route should not be advertised (no-re advertise).

For this purpose, you assign a route one of the following attributes:

- retain

  Keeps the route in the forwarding table after the routing process shuts down or the next- hop router goes down. (By default, routes are not retained in the routing table. When the routing process starts up again, any routes configured as static routes must be added by the system to the forwarding table again. To avoid this latency, routes can be flagged as retain, so that they are kept in the forwarding table even after the routing process shuts down.

- passive

  Rejects traffic destined for the route. Generally, only active routes are included in the routing and forwarding tables. If a static route's next-hop address is unreachable, the route is marked ineligible, and it is not included in the routing or forwarding tables. To force a route to be included in the routing tables regardless of next-hop reachability, you can flag the route as passive. If a route is flagged as passive and its next-hop address is unreachable, the route is included in the routing table and all traffic destined for it is rejected.

Installing Routes into More Than One Routing Table

By default, IPv4 interface routes (also called direct routes) are imported into routing table inet.0. When an interface is bound to a routing instance, its direct route is added into the instance inet.0 table.

You can add interface routes to other routing tables on an interface basis. To do so, you create a (RIB) routing table RIB group.

When you specify a RIB group, the primary routing table—the first one you specify—determines the address family of the routing table group (for example, IPv4, IPv6, and so on). If the primary routing table is deleted, the secondary ones are also deleted.

For the Albuquerque site, each zone has its own routing instance. This distinction makes it easy to create a routing group for each routing instance. You add to the routing group only the tables from other routing instances whose interfaces should be visible to it. If hosts within a zone should not have access to hosts within another zone, the second zone's interfaces should not be exposed to the first zone.

To create a routing group to be used to add other routing tables to the IntranetALBQ-vr routing instance, enter the following statements in configuration mode:

```
user@host# set routing-options rib-groups intranet-rg import-rib
IntranetALBQ-vr.inet.0

user@host# set routing-options rib-groups intranet-rg import-rib
TellersALBQ-vr.inet.0

user@host# set routing-options rib-groups intranet-rg import-rib
FinancialALBQ-vr.inet.0
```

```
user@host# set routing-options rib-groups intranet-rg import-rib
BankManagersALBQ-vr.inet.0
user@host# set routing-options rib-groups intranet-rg import-rib
PublicAppServersALBQ-vr.inet.0
```

To create a routing group to be used to add other routing tables to the InternetALBQ-vr routing instance, enter the following statements in configuration mode:

```
user@host# set routing-options rib-groups internet-rg import-rib
InternetALBQ-vr.inet.0

user@host# set routing-options rib-groups internet-rg import-rib
FinancialALBQ-vr.inet.0

user@host# set routing-options rib-groups internet-rg import-rib
BankManagersALBQ-vr.inet.0
```

To create a routing group to be used to add other routing tables to the PublicAppServersALBQ-vr routing instance, enter the following statements in configuration mode:

```
user@host# set routing-options rib-groups publicAppSrv-rg import-rib
PublicAppServers-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
TellersALBQ-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
FinancialALBQ-vr.inet.0

user@host# set routing-options rib-groups publicAppSrv-rg import-rib
BankManagersALBQ-vr.inet.0
```

To create a routing group to be used to add other routing tables to the TellersALBQ-vr routing instance, enter the following statements in configuration mode:

```
user@host# set routing-options rib-groups tellers-rg import-rib
TellersALBQ-vr.inet.0 set routing-options rib-groups tellers-rg import-rib
PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups tellers-rg import-rib
IntranetALBQ-vr.inet.0
```

To create a routing group to be used to add other routing tables to the BankManagersALBQ-vr routing instance, enter the following statements in configuration mode:

```
user@host# set routing-options rib-groups bankManagers-rg import-rib
BankManagersALBQ-vr.inet.0

user@host# set routing-options rib-groups bankManagers-rg import-rib
PublicAppServersALBQ-vr.inet.0
```

user@host# set routing-options rib-groups bankManagers-rg import-rib IntranetALBQ-vr.inet.0

user@host# set routing-options rib-groups bankManagers-rg import-rib InternetALBQ-vr.inet.0

To create a routing group to add other routing tables to the FinancialALBQ-vr routing instance, enter the following statements in configuration mode:

user@host# set routing-options rib-groups financial-rg import-rib FinancialALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib PublicAppServersALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib IntranetALBQ-vr.inet.0

user@host# set routing-options rib-groups financial-rg import-rib InternetALBQ-vr.inet.0

Alternatively, you can enter these statements at the edit routing-options hierarchy level and omit the term routing-options.

Table 23 gives the RIB group's statement syntax.

Table 23: RIB Group Configuration Statement

RIB group	Syntax
Specifies a routing table group and adds routing tables to it.	`set routing-options rib-groups <group-name> import-rib [<routing-table-name>...]` The first routing table you list in the import-rib statement must be the one you configured in the rib-group statement. ■ `group-name`: Name of the routing table group. The name must begin with a letter and it can include letters, numbers, and hyphens. ■ `routing-table-name`: Names of the routing tables into which the software should import routing information. The first routing table name you enter is the primary routing table. Any additional names you specify identify secondary routing tables, Specify inet.0 as the primary routing table for IPv4 routing table groups.

Creating Policies

For transit traffic to pass between zones (intrazone) and across interfaces of the same zone (intrazone), there must be a security policy that allows it.

This section explains how to configure the components required to create a policy. using the Albuquerque branch office scenario example.

About Policies

You configure policies to control access to hosts and other resources within a zone. Policies enforce rules that act on traffic based on the information contained in the policy. In addition to permitting or denying traffic or any of the other actions, a policy can specify accounting and auditing actions and a schedule stipulating when the policy is active. A policy contains five match criteria. If a packet matches the components of a policy, the action that the policy specifies is applied to the packet. You can use the predefined keyword ANY as a wildcard for any of the five match criteria to broaden the scope of a policy.

In the from-zone parameter, a policy specifies the incoming zone. In the to-zone parameter, a policy specifies the outgoing zone. Traffic enters one security zone, the from-zone, and goes out on another zone, the to-zone. In combination, these zones create a context.

Here are the actions that you can specify in a policy:

- Permit: Allows traffic to flow through the router.
- Deny: Drops the traffic.
- Reject: Drops the traffic and returns an error.
- Firewall authentication: authenticates the client before forwarding the traffic.
  - Passthrough: Access profile or client match
  - Web authentication: Client match
- Network Address Translation (NAT): Performs address and port translation on the permitted traffic.

Apart from the default policies to deny all or permit all, you can create wide-spanning policies. At the broadest level, you can create a policy to allow traffic of all kinds from any source in one zone to any destination in all other zones without time-based restrictions. At the narrowest level, you can create a policy that allows only one kind of traffic between a specific host in one zone and another specific host in another zone during a scheduled period.

Policies are unidirectional, which means that if you want traffic to be allowed both to and from a zone, you must configure two polices for it. For example, if you want to allow HTTP traffic to be sent from zone A to zone B, you must configure a policy for it. After the session is established, zone B can exchange traffic with zone A. However, if you want to allow zone B to initiate a session to send HTTP traffic to zone A, you must configure a separate policy for it.

To configure a policy, you use two statements:

- The first statement specifies the match criteria for the policy.
- The second statement specifies the action for the policy.

The following example shows a policy configuration that allows SSH traffic from the BankManagersALBQ zone at the branch to the PublicAppServersALBQ zone, also at the branch.

- The BankManagersALBQ zone's address book contains the bankMan address.
- The PublicAppServersALBQ zone's address book contains the PubApps2Web address.

  user@host# set security policies from-zone BankManagersALBQ to-zone PublicAppServersALBQ policy BankMantoWebApps match source-address bankMan destination-address PubApps2Web application ssh

  user@host# set security policies from-zone BankManagersALBQ to-zone PublicAppServersALBQ policy BankMantoWebApps then permit

Table 24 gives an overview of the policy configuration statement syntax.

Table 24: Policy Configuration Statement

Policy Configuration	Syntax
Configure a policy.	`set security policies from-zone <zone-name> to-zone <zone-name> policy <policy-name> match source-address (<address-name> | <address-set-name>) destination-address (<address-name> | <address-set-name>) application (<application-name> | <application-set-name>)` `set security policies from-zone <zone-name> to-zone <zone-name> policy <policy-name> then (permit | firewall-authentication | tunnel ipsec-vpn <vpn-name> pair-policy <pair-policy-name> | source-nat (pool <pool-name>) | pool-set <pool-set-name> | interface) | destination-nat <name> | deny | reject) schedule <scheduler-name> log <session-init session-close> count alarm per-second-threshold <value> per-minute-threshold <value>`  ■ `(from-zone) zone-name`: Source zone. Name of the zone from which traffic is sent. ■ `(to-zone) zone-name`: Destination zone. Name of the zone to which traffic is to be sent. ■ `policy-name`: Unique name used to refer to the policy. ■ `(source-address) address-name`: Name of an address (or address set) as entered in the source zone's address book. ■ `(destination-address) address-name`: Name of an address (or address set) as entered in the destination zone's address book. ■ `(application) application-name`: Name of a preconfigured or custom application (or application set). ■ `(then)`: Any one of the actions listed in the policy syntax following the word "then". ■ `scheduler-name`: Optionally, the name of a scheduler whose schedule determines when the policy is active and when it can be used. ■ `(log) session-init session-close`: Log the traffic that matches the policy ■ `(count) values`: Count setters.

Steps for Creating a Policy

The following sections explain the tasks you perform to create a policy. Table 25 shows the steps to follow.

Table 25: Creating a Policy

Task	Instructions
1. Configure a zone's hosts addresses and subnet addresses in its address book.	See "Configuring Address Books" on page 66.
2. Configure applications.	See "Configuring Applications and Application Sets" on page 69.
3. Create security policies.	See "Configuring Policies" on page 71
4, Create schedulers, if you plan to use them for your policies.	See "Configuring Schedulers" on page 76.

Configuring Address Books

Before you can set up policies, you must define addresses for each zone's address book. A zone's address book must contain entries for the addressable systems (and, thus, users) belonging to the zone.

An address book for a security zone contains the IP addresses or domain names of hosts and subnets whose traffic is either allowed, blocked, encrypted, or user-authenticated. Additionally, you can define address groups and refer to the members of the group collectively.

Be sure to account for all hosts within a zone with static addresses, not just its subnets, if individual hosts have special access requirements. For example, the PublicAppServersALBQ zone has two servers. Addresses for each of these servers must be added to the zone's address book because each server must be reachable independently of the other.

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 destination zone and destination 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 source address in policies.

If traffic matches the source and destination information together with the application, or service, of a policy, the action specified by the policy is applied to the packet.

Table 26 gives an overview of the address book feature and its syntax,

Table 26: Address Book Configuration Statement

Address Book Addresses Configuration	Syntax
Adds a host or subnet address to the zone, and specifies a name to use to refer to it in a policy.	`set security zones security-zone <zone-name> address-book address (<IPV4-address> | <hostname>)` ■ `zone-name`: The name of the zone for which you are creating the address. ■ `(address)`: IPv4 addresses with the number of prefix bits, or a DNS name. You can only use domain names if the system is configured to use DNS services.

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 refer to an address set in a policy instead of an individual address book entry.

Table 27 gives an overview of the address set feature and its statement syntax.

Table 27: Address Set Configuration Statement

Address Book Address Set Configuration	Syntax
Defines an address set containing more than one address. ■ To define an address set, first define individual addresses, then add them to a set. ■ Individual names belonging to a set are each used by security policies as criteria. For the policy to apply to the packet, any one address in the set can be used as matching criteria when all other parts of the policy match.	`set security zones security-zone <zone-name> address-book address-set <address-book-set-name> address <address-name>` ■ `zone-name`: The name of the zone for which you are creating the address. ■ `address-book-set-name`: Name to be used to refer to the set of addresses collectively in a policy. ■ `address-name`: Predefined name of an address to be added to the set.

NOTE: Consider that for each address set, the router 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 do not consider this expansion when you create address books, you could 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.

Configuring Address Books for Zones

This section describes the addresses and address sets configured for the address books of the Albuquerque branch zones.

Address book for the IntranetALBQ zone

The IntranetALBQ zone contains two servers. These servers belong to the same subnet, but they are separately addressable. To configure address book entries for the Intranet zone, enter the following statements in configuration mode.

This example adds individual addresses for both of the 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.

```
user@host# set security zones security-zone IntranetALBQ address-book address
BankRec1 13.16.14.16/32

user@host# set security zones security-zone IntranetALBQ address-book address
BankRec2 13.15.14.17/32

user@host# set security zones security-zone IntranetALBQ address-book
address-set HeadSerAll address BankRec1

user@host# set security zones security-zone IntranetALBQ address-book
address-set HeadSerAll address BankRec2
```

Address book for the TellersALBQ zone

The tellers of the Albuquerque branch belong to the same subnet and they share the same access rights. To define an address book entry for the tellers subnet in the TellersALBQ zone, enter the following statement in configuration mode:

```
user@host# set security zones security-zone TellersALBQ address-book address
tellers 10.11.13.0/24
```

Address book for the FinancialALBQ Zone

The branch's financial managers share the same subnet and are granted the same access rights. To define an address book entry for the managers in the FinancialALBQ zone, enter the following statement in configuration mode:

```
user@host# set security zones security-zone FinancialALBQ address-book
address finManagers 10.12.12.0/24
```

Address book for the BankManagersALBQ Zone

The bank managers belonging to the BankManagersALBQ zone share the same subnet and have the same access rights. To define an address book entry for the bank managers in the BankManagersALBQ zone, enter the following statement in configuration mode:

```
user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan 10.11.14.0/24
```

Address book for the PublicAppServersALBQ Zone

To set up the address book for the servers in the PublicAppServersALBQ zone, enter the following statements in configuration mode. An address is added for each of the servers so that they can be addressed individually. The address-set parameter is used to create a single address that includes both of the servers.

```
user@host# set security zones security-zone PublicAppServersALBQ
address-book address PubApps1Mail 11.10.15.5/32
```

```
user@host# set security zones security-zone PublicAppServersALBQ
address-book address PubApps2Web 10.11.15.6/32

user@host# set security zones security-zone PublicAppServersALBQ
address-book address-set PublicAppBoth address PubApps1

user@host# set security zones security-zone PublicAppServersALBQ
address-book
address-set PublicAppBoth address PubApps2
```

Configuring Applications and Application Sets

When you create a policy, you specify an application, or service. The application (or application set) serves 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.

Here is how application sets work:

- You add predefined or custom applications separately to an application set, which you name. (After you create a custom service, you can refer to it in a policy.)
- Instead of a discrete application, you specify the name of the application set to be used as part of a policy.
- If a packet matches all other match criteria of the policy and any one of the applications included in the set, the policy applies to the packet.

Creating an Application Set

Instead of creating individual policies to grant the financial managers the right to use these applications, you could create an application set and refer to the set as an item from within the policy. Here is how application sets work:

- You add predefined or custom applications separately to an application set, which you name.
- Instead of a discrete application, you specify the application set as part of a policy.
- If a packet matches all other match criteria of the policy and any one of the applications included in the set, the policy applies to the packet.

Suppose that as part of your firewall philosophy, you created a table similar to the example shown in Table 6 on page 19. In this case, for each group of employees, you could create an application set containing the applications allowed them.

For example, all financial managers on the Financial zone's subnet 10.12.12.0/24 are allowed to use SSH, Telnet, and CustApp, a custom application, to access customer records for special financial services on the BankRec2 server of the IntranetALBQ zone.

The following example shows how to create an application set for the three applications that allow managers in the FinancialALBQ zone and the BankManagersALBQ zone to perform the following tasks:

- Log into the servers in the IntranetALBQ zone.
- Use the database.
- Transfer files.

The application set is given the generic name BankAppSet.

```
user@host# set applications application-set BankAppSet application ssh

user@host# set applications application-set BankAppSet application telnet

user@host# set applications application-set BankAppSet 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 PublicAppServersALBQ 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
```

Table 28 gives an overview and syntax for the application set statement.

Table 28: Application Set Configuration Statement

Application Set Configuration	**Syntax**
Creates an application set, which combines more than one application in a group. A policy can refer to an application set rather than a single service as its application match criteria. In that case, if the traffic type matches any one application of the set and the other policy criteria match, then the policy's action is applied to the packet.	`set applications application-set <application-set-name> application <application>` ■ `application-set-name`: The name of the set of applications, which is to be referred to in a policy ■ `application`: The individual application to add to the set.

Configuring Policies

A security policy applies security or logging and accounting rules to transit traffic.

A packet is matched against configured 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 that the policy specifies. If the packet matches all tuples of a policy, that policy's action is applied to the packet. 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 router has looked up its route. The packet's destination address determines its destination zone.

Figure 4 illustrates how the router uses policies to determine the action to apply to transit traffic:

- A policy stipulates that traffic from the FinancialALBQ zone is allowed access to the BankRec2 server in the IntranetALBQ zone.
- Traffic from the TellersALBQ zone is denied access to the BankRec2 server in the IntranetALBQ zone. Either there is a specific policy that denies this traffic access to the server, or the default policy that denies all traffic is the only other policy that applies to the packet.

Figure 4: Firewall Policies Determining the Fate of a Packet

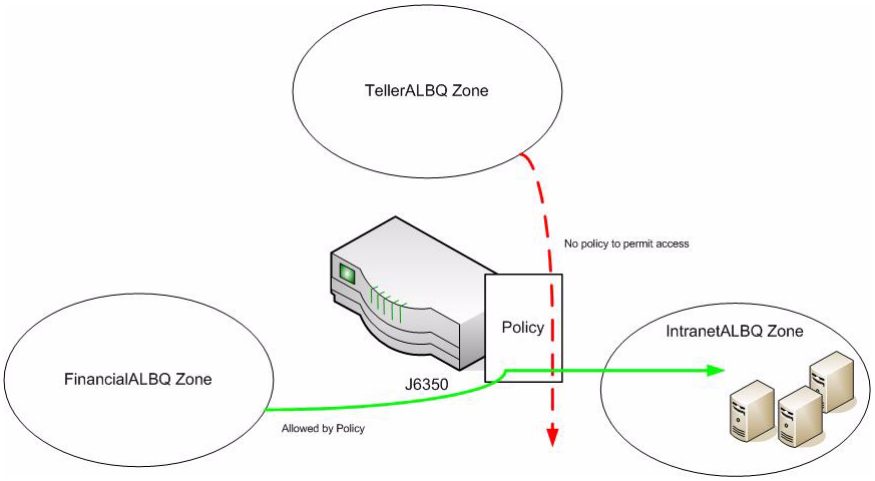

You can use any of the keywords shown in Table 29 to specify a policy's action.

Table 29: Policy Actions

Policy Action	Description
Permit	Allows the packet to pass through the firewall.
Reject	Blocks the packet from traversing the firewall. The firewall drops the packet and sends a TCP reset (RST) segment to the source host for TCP traffic and an ICMP destination unreachable, port unreachable message (type 3, code 3) for UDP traffic. In this case, for types of traffic other than TCP and UDP, the firewall drops the packet without notifying the source host.
Deny	Blocks and drops the packet from traversing the firewall, but sends no notification back to the source.
Log	Logs packets that enter the firewall. This action produces a full log of firewall traversal if the packet traverses the firewall.
Count	Counts the number of packets and bytes that enter the firewall for a given policy. For counts, you can specify that alarms be generated whenever the traffic exceeds specified thresholds.
Authenticate	Enables user authentication to allow traffic to pass securely through the router.
Tunnel	Encapsulates outgoing IP packets and decapsulates incoming IP packets.

About Creating Policies with the CLI

To create policies using the CLI, you enter a single statement that specifies the matching criteria and another statement that specifies the action to be taken if a packet matches all criteria.

Creating Policies for the Albuquerque Branch

The default policy for the Albuquerque branch is to deny all traffic unless explicitly permitted through another policy statement.

The following sections give example policy rules for the Albuquerque branch employees' hosts. These are only some of the policies required for this site.

Default policy of deny all

The following statement implements the default security stance for the Albuquerque branch, which is to deny all users access to any of the resources outside their own zones. Intrazone traffic is also blocked by default. This statement is shown for completeness. You do not need to configure a default policy of deny-all unless you have previously changed the default to permit-all.

```
user@host# set security policies default-policy deny-all
```

The default policy should be at the bottom of your policy list because it is the most general.

Table 30 shows the statement syntax for configuring the default policy.

Table 30: Default Policy Configuration Statement

Default Policy Configuration	Syntax
Defines a default security stance, for the router. By default, when the device is in secure context, the position is to deny all traffic between zones and within a zone. However, in that context, policies are configured to allow traffic from the trust zone to the untrust zone. (The untrust zone encompasses all interfaces other than ge-0/0/0 when the system is initialized in secure context.) Also, all transit traffic within the trust zone is allowed.	`set security policies default-policy permit-all | deny-all` `permit-all`: Allow all traffic within and between zones. `deny-all`: Deny all traffic within and between zones.

Policies for members of the TellersALBQ zone

Adhering to the design guidelines of the network architect, the security administrator configures the following policies for the TellersALBQ zone.

IntranetALBQ: The following policy grants tellers access to customer records on the BankRec1 server in the IntranetALBQ zone using the custom database access application, custApp. The BankRec1 server is located at the Los Angeles corporate headquarters, and the tellers systems are located at the branch office. Traffic sent from any of the tellers to the BankRec1 server in the IntranetALBQ zone travels through one of the IPSec VPN tunnels, depending on the route lookup.

The following route specifies that all traffic destined for BankRec1 (13.16.14.16) is to be sent to st0.1 as the next hop. The primary IPSec tunnel is bound to st0.1, so it is used to route traffic to the BankRec1 server in the IntranetALBQ zone, if a policy permits it.

```
user@host# set IntranetALBQ-vr routing-options static route 13.16.14.16/32
next-hop st0.1 preference 7
```

To configure the policy to allow traffic of the custom application type (custRecApp) from the tellers to be sent to the BankRec1 server in the IntranetALBQ zone, enter the following statement in configuration mode.

```
user@host# set security policies from-zone TellersALBQ to-zone IntranetALBQ
policy TellCustRec match source-address tellers destination-address BankRec1
application custRecApp

user@host# set security policies from-zone TellersALBQ to-zone IntranetALBQ
policy TellCustRec then permit
```

PublicAppServersALBQ: The tellers require access to the mail server in the PublicAppServersALBQ zone. Because the server and the tellers are located at the branch office, the traffic does not need to pass through an IPSec VPN. The PublicAppServersALBQ zone is configured to allow POP3 and SMTP traffic.

To enable tellers to send e-mail using the mail server, enter the following statements in configuration mode:

```
user@host# set security policies from-zone TellersALBQ to-zone
PublicAppServersALBQ policy TellMailOut match source-address tellers
destination-address PubApps1Mail application [POP3 smtp]

user@host# set security policies from-zone TellersALBQ to-zone
PublicAppServersALBQ policy TellMailOut then permit
```

Policies for members of the FinancialALBQ zone

IntranetALBQ BankRec2 server: To create a policy to allow financial officers to gain access to customer investment and special services records on the BankRec2 server at headquarters, enter the following statements in configuration mode.

Recall that during route lookup, a static route to the destination address is discovered. The route specifies that traffic for this destination is to be sent to the st0.1 tunnel interface. The st0.1 interface is bound to an IPSec VPN, and so the traffic is forwarded to the destination specified in the policy (BankRec2) via the primary IPSec VPN.

```
user@host# set IntranetALBQ-vr routing-options static route 13.15.14.17/32
next-hop st0.1 preference 7
```

The following policy allows that traffic transmission to occur. In this case, an application set called BankAppSet, which refers to the applications SSH, Telnet, and custApp, is used.

```
user@host# set security policies from-zone FinancialALBQ to-zone IntranetALBQ
policy FinRecPolOut match source-address finManagers destination-address
BankRec2 application-set BankAppSet

user@host# set security policies from-zone FinancialALBQ to-zone Intranet
policy FinRecPolOut then permit
```

PubAppServersALBQ: To create a security policy to allow members of the FinancialALBQ zone's subnet to send e-mail using the PubApps1MailALBQ server and to use Web applications using the PubApps2Web server, both of which are in the PublicAppServersALBQ zone, enter the following statements in configuration mode:

```
user@host# set security policies from-zone FinancialALBQ to-zone
PublicAppServersALBQ policy FinMailWebOut match source-address
finManagers
destination-address PublicAppBoth application [smtp POP3]

user@host# set security policies from-zone FinancialALBQ to-zone
PublicAppServersALBQ policy FinMailWebOut then permit
```

Policies for members of the BankManagersALBQ zone

IntranetALBQ: To create a policy for members of the BankManagersALBQ zone to access database records on either of the customer records servers at headquarters, enter the following statements in configuration mode. In this case, the PublicAppBoth destination address set, which includes both servers, is used.

```
user@host# set security policies from-zone BankManagersALBQ to-zone
IntranetALBQ policy BankManHQServOut match source-address bankMan
destination-address PublicAppBoth application-set BankAppSet

user@host# set security policies from-zone BankManagersALBQ to-zone
IntranetALBQ policy BankManHQServOut then permit
```

PublicAppServersALBQ: To create a policy to allow members of the BankManagersALBQ zone to send email using the PubApps1Mail server and to use Web applications using the PubApps2Web server, both of which are in the PublicAppServersALBQ zone, enter the following statement in configuration mode. In this case, an address set, PublicAppBoth, is used to refer to the two servers collectively

The following policy allows members of the BankManagersALBQ zone to send email and use Web applications:

```
user@host# set security policies from-zone BankManagersALBQ to-zone
PublicAppServersALBQ policy BankManMailWebOut match source-address
bankMan destination-address all application-set WebMailApps

user@host# set security policies from-zone BankManagersALBQ to-zone
PublicAppServersALBQ policy BankManMailWebOut then permit
```

Policies for the servers in the PublicAppServersALBQ zone

To create a policy to allow the two servers in the PublicAppServersALBQ zone to automatically send information to managers belonging to the BankManagersALBQ zone, enter the following statements in configuration mode:

```
user@host# set security policies from-zone PublicAppServersALBQ to-zone
BankManagersALBQ policy ManRecServIn match source-address PublicAppBoth
destination-address bankMan application-set CustAppMessage

user@host# set security policies from-zone PublicAppServersALBQ to-zone
BankManagersALBQ policy ManRecServIn then permit
```

Policies to allow Internet access

To create a policy to allow the financial managers access to the Internet, enter the following commands in configuration mode.

```
user@host# set security policies from-zone FinancialALBQ to-zone InternetALBQ
policy InternetAccess match source-address finManagers
destination-address all application-set all

user@host# set security policies from-zone FinancialALBQ to-zone InternetALBQ
policy InternetAccess match then permit
```

To create a policy to allow the bank managers access to the Internet, enter the following commands in configuration mode.

```
user@host# set security policies from-zone BankManagersALBQ to-zone
InternetALBQ policy InternetAccess match source-address bankMan
destination-address all application-set all

user@host# set security policies from-zone BankManagersALBQ to-zone
InternetALBQ policy InternetAccess match then permit
```

Configuring Schedulers

Schedulers are a powerful feature that allows you to set times to specify when a policy is active. You can define schedulers and refer to them from any policies. If a policy contains a reference to a scheduler, the schedule determines when the policy is active and 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.

For example, if a policy permits access, you can use a scheduler to deactivate the policy for a specified time to lock down a resource from access by the entity or group the policy refers to. Alternatively, if a policy specifies deny as its action, you can create a scheduler that removes that constraint for a period of time. However, in this case, for the user or group to gain access to the resource, there must be another policy to permit it.

This section shows how to configure several schedulers. Using example policies configured for the deployment, it shows how to refer to a scheduler from the policy.

Regular hours scheduler

The following three statements set a schedule allowing a policy that refers to it to be used for packet match checks from 7 a.m. to 6 p.m. all days of the week from October 1, 2007 to October 1, 2008 except Saturday and Sunday. Otherwise, the policy is inactive.

```
user@host# set schedulers scheduler RegularHours start-date
2007-10-01.07:00
stop-date 2008-10-01.18:00

user@host# set schedulers scheduler RegularHours saturday exclude

user@host# set schedulers scheduler RegularHours sunday exclude
```

The following statement applies the schedule to the policy for bank tellers, which gives them access to customer checking and savings records during regular work hours, as specified by the scheduler.

```
user@host# set security policies from-zone TellersALBQ to-zone IntranetALBQ
policy TellCustRec scheduler-name RegularHours
```

Schedule for weekend hours

The following statements set schedulers allowing policies that refer to them to be used to check for packet matches from noon to 5 p.m. on Saturdays and Sundays. Otherwise the policy is inactive.

```
user@host# set schedulers scheduler SaturdayHrs saturday start-time 12:00
stop-time 17:00
```

```
user@host# set schedulers scheduler SundayHrs sunday start-time 12:00
stop-time 17:00
```

The following statements apply these schedules to the policy for bank managers to give them access to both of the servers whose databases contain customer records during the specified weekend hours:

```
user@host# set security policies from-zone BankManagersALBQ to-zone
IntranetALBQ policy BankManRecServOut scheduler-name SaturdayHrs
```

```
user@host# set security policies from-zone BankManagersALBQ to-zone
IntranetALBQ policy BankManRecServOut scheduler-name SundayHrs
```

Policy Ordering Reminders

Take care to arrange your policies in the order that is appropriate for your network.

- Policy statements are processed in a top-down fashion. If a statement matches the packet under evaluation, the policy action is executed and no other policy lines are searched.
- When you add new policies, they are added to the end of the policy list, which is most likely not what you intended. Reorder the list to accommodate new entries.
- Place the most specific entries at the top of the list and the more general entries at the bottom.
- Place IPSec VPN tunnel entries close to the top of the list.
- Your policy list must have at least one permit statement. If it consists of all deny statements, no traffic is allowed.
- If no matches are found, the traffic is denied by default.

Identifying Potential Attacks and Configuring Firewall Defense Mechanisms

Attacks fall into two main categories:

- Information gathering, or reconnaissance, attacks are meant to identify points of entry into the network.

  Though seemingly harmless themselves, these attacks are preludes to more serious and damaging ones.

- Penetration attacks are meant to damage to the network or its resources, for example, by locking out legitimate users or stealing information.

Many attacks are perpetrated through the use of various aspects of the TCP, IP, and UDP protocols. In some cases, an attacker could use the same tactic for either reconnaissance or penetration attacks. For example, through the onslaught of TCP SYN packets an intruder might either attempt to map the network's IP addresses or lock up its resources so that they are no longer available to legitimate users.

Screens can protect against attacks. Here are some things to consider about screens:

- How Screens Work—You apply screens to a zone, and they filter traffic before it leaves the zone to ensure that any traffic ensuing from an attack within the zone is dropped at the source, or otherwise treated appropriately. Screens filter traffic before it leaves a zone so that attack traffic cannot pass through to other zones.

- About Use of Screens—Screens consume memory and CPU cycles. The number of screens you use and the amount of traffic they process when used in combination with other features can affect system performance, possibly introducing some latency, Therefore, you should consider carefully which screens to use for each zone. For example, if your network has no policy allowing ICMP traffic to be sent from a zone, then applying any of the ICMP related screens to that zone would be effectively meaningless but costly in regard to CPU utilization.

Here are some of factors to take into account in determining whether to use a specific screen:

- How high is the threat associated with a type of attack?

- How common is the attack for your network?

- What is the cost in terms of system performance for use of a screen?

Other questions to consider that pertain to rate limiting traffic are:

- Should tighter constraints be put in place to protect against certain kinds of attacks at the risk of loss of legitimate packets?

- Is it better to let some attack traffic through to ensure that all legitimate traffic is delivered, but not enough to flood the system into denial of service?

You may be willing to tolerate a specific kind of attack, such as a reconnaissance attack; You may discover that for your environment, certain kinds of attacks are uncommon; And you may realize that certain screens have little relevance for your network environment.

The following sections describe the screens you can configure to protect zones against these attacks:

- Preventing Reconnaissance Attacks on page 82
- Configuring the Firewall to Protect Against Penetration Attacks: Firewall and Network DoS Attacks on page 91
- Preventing Operating System DoS Attacks on page 100

Configuring Screens for the Branch's Zones

You apply screens to zones, and they filter outgoing traffic. The firewall inspects a packet, checking it against the zone's screens, including their rate-limiting constraints. In most cases, the firewall drops traffic in violation of a screen's parameters and tests. Depending on the screen, the firewall may only generate and log an alarm message.

This section identifies the screens to be assigned to the zones of the Albuquerque branch.

Assessing Risk and Defense Requirements

For the Albuquerque office of the New Bank of the Southwest, the main thrust of the deployment design is toward greater security at the risk of legitimate packet loss. Here is the basic strategy the network architect has specified:

- Protecting the network and the firewall from outside attacks is of great importance. For this reason, all screens are to be applied to the InternetALBQ zone.
- Protecting the network, customer data and other records, and corporate resources against breaches of any kind launched from within the LAN, whether intentional or accidental, is of equal importance for a bank. Many attacks which are costly in terms of real dollars have been perpetrated from within organizations by disgruntled or malicious employees, contractors, or associates.

  Furthermore, in addition to use of their own systems to launch attacks, users can exploit a LAN server, using it as a zombie agent. For this reason, zones within the LAN are to be protected with many of the firewall's screens, even those zones with only server hosts. However, rate limits are set to allow for business to proceed reasonably without compromising the system.

About Creating Screens

Use of some screens makes sense only if the system's configuration lends itself to the kind of attack that the screen protects against. For example, you can configure a screen to protect against IP address sweep reconnaissance attacks. However, this screen is only useful for a security zone for which there is a policy permitting ICMP traffic from that zone.

Flow-based Filters

There are two flow-based screens, or filters, that are used to protect the firewall against DoS attacks. They are effective at the gross system level across all zones, not at the discrete zone level. The Albuquerque site uses these screens with the following settings:

```
user@host# set security flow aging early-ageout 4 high-watermark 80
low-watermark 60

user@host# set security flow syn-flood-protection-mode syn-cookie
```

For details, see the following sections:

- Using Aggressive Aging to Protect Against Session Table Flood on page 92
- About SYN Cookie on page 96

Screen Sets for the Albuquerque Branch Office

For the Albuquerque site, the network designer has determined that the following screen sets, designed along functional lines, are to be created. For most of the zones, the same parameter values are to be used for the screens, so it is easier to create one set of a particular kind of screens and use it for all zones for which it is appropriate. However, if a zone requires unique parameter values for a particular screen, you should create a separate screen set for it.

After this section shows the common screen sets, it shows how these sets are used for two zones: the InternetALBQ zone and the TellersALBQ zone. The security administrator would apply them to other zones of the branch, based on the requirements for a zone and the kind of traffic it transmits.

Because there is no policy allowing ICMP-based applications to be transmitted from the TellersALBQ zone, a special set is defined for it that does not include screens that handle ICMP packets.

Here are the common screen sets defined for the Albuquerque branch office:

Network reconnaissance and identity evasion	**user@host# set security screen ReconnScrn icmp ip-sweep threshold 1000**
	user@host# set security screen ReconnScrn tcp fin-no-ack port-scan threshold 1000
	user@host# set security screen ReconnScrn tcp syn-fin tcp-no-flag syn-frag
	user@host# set security screen ReconnScrn ip spoofing loose-source-route-option strict-source-route-option
	user@host# set security screen ReconnScrn ip record-route-option timestamp-option bad-option security-option stream-option block-frag unknown-protocol
	user@host# set security screen ReconnScrn icmp fragment large
Firewall DoS attacks	**user@host# set security screen FireDoSScrn limit-session source-ip-based 100 destination-ip-based 100**
	user@host# set security screen FireDoSScrn tcp syn-ack-ack-proxy threshold 1000
Network DoS attacks	**user@host# set security screen NetDoSScrn tcp syn-flood attack-threshold 1000 alarm-threshold 512 source-threshold 4000 destination-threshold 4000 timeout 30 queue-size 1024**
ICMP and UDP flood attacks	**user@host# set security screen ICMPfloodScrn icmp flood threshold 1500**
	user@host# set security screen UDPfloodScrn udp flood threshold 1500
Land attack	**user@host# set security screen LandAttackScrn tcp land**
Operating system DoS attacks	**user@host# set security screen OSDoSAttacksScrn icmp ping-death**
	user@host# set security screen OSDoSAttacksScrn ip tear-drop
	user@host# set security screen OSDoSAttacksScrn tcp winnuke

IntranetALBQ zone screens

To apply screen sets to the InternetALBQ zone, you enter the following statements in configuration mode:

user@host# set security zones security-zone InternetALBQ screen ReconnScrn

user@host# set security zones security-zone InternetALBQ screen FireDoSScrn

user@host# set security zones security-zone InternetALBQ screen NetDoSScrn

user@host# set security zones security-zone InternetALBQ screen ICMPfloodScrn

user@host# set security zones security-zone InternetALBQ screen UDPfloodScrn

```
user@host# set security zones security-zone InternetALBQ screen
LandAttackScrn

user@host# set security zones security-zone InternetALBQ screen
OSDoSAttacksScrn
```

TellersALBQ zone screens

Because there is not a policy allowing ICMP traffic from the Tellers zone, the ICMPfloodScrn is not applied to the zone. To apply screen sets to the TellersALBQ zone, you enter the following statements in configuration mode:

```
user@host# set security zones security-zone InternetALBQ screen ReconnScrn

user@host# set security zones security-zone InternetALBQ screen FireDoSScrn

user@host# set security zones security-zone InternetALBQ screen NetDoSScrn

user@host# set security zones security-zone InternetALBQ screen UDPfloodScrn

user@host# set security zones security-zone InternetALBQ screen
LandAttackScrn

user@host# set security zones security-zone InternetALBQ screen
OSDoSAttacksScrn
```

Preventing Reconnaissance Attacks

In almost all cases, an intruder precedes a malicious attack with a network scanning and reconnaissance attack. These attacks are meant to gain information about the logical topology of the network: IP addresses of hosts composing the network's IP address space, active hosts and active ports on active hosts, and the operating system that a host runs.

Information-gathering attacks have the following purposes and traits:

- They enable the attacker to learn the network's characteristics and vulnerabilities.

  They occur when a perpetrator probes a network or system by sending unusual, foreign, or altered packets into its flow.

  They are meant to uncover weaknesses in the system and expose entry points, making the network's security easier to breach.

- They do not compromise systems. They are commonly preludes to more serious and damaging attacks, although they themselves are seemingly harmless.

You can apply screens to a zone to ensure that traffic leaving the zone is not used for the following purposes:

- Network mappings and host and port scans

  See "Preventing IP Address Sweeps and Host and Port Scans" on page 83.

- Operating system and network identification

  See Figure 7, “TCP Header With the SYN and FIN Flags Set,” on page 86.

- Identity evasion, used in conjunction with various other kinds of attacks

  See “Accountability Evasion” on page 90.

Information-gathering attacks are easy and common because the attacker can use any one of the many readily accessible tools designed for legitimate purposes, such as NMap (meant to map the network) and SuperScan (meant to scan it for TCP ports) to nefarious ends.

For some screens, you can specify threshold values to rate-limit the allowed number and kind of packets. The stateful firewall inspects a packet to determine if it is anomalous or, in the context of other packets in the session and packets of other sessions, if it is part of an attack.

The following sections address reconnaissance attacks and the screens used to identify handle traffic of that type.

Preventing IP Address Sweeps and Host and Port Scans

An attacker can use several approaches to discover information about the network. Figure 5 shows how the router protects against an ICMP IP address sweep attack when the IP address sweep (icmp ip-sweep) screen is used.

Figure 5: Protection Against IP Address Sweep

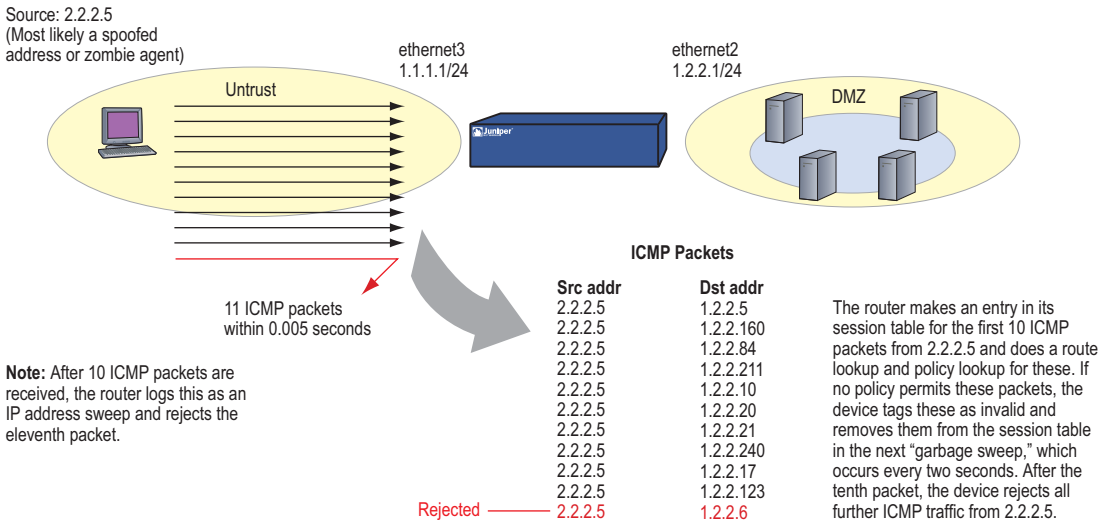

Table 31 describes the kind of address sweep attack and the firewall screen used to defend against it

This table shows individual statements. However, you can enter tcp screen options on a single statement.

Table 31: IP Address Sweep and Host and Port Scan Screen Configuration Statements

Attack	Screen Syntax
IP Address Sweep (ICMP Network Scan), or Ping Sweep	Screen for IP Address Sweep `set security screen <screen-name> icmp ip-sweep threshold <interval>` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. ■ `interval`: Interval in microseconds. (Default: 5000) The valid threshold range is 1000 through 10000. On a router with the default settings, if a remote host sends ICMP traffic to 10 addresses in 0.005 seconds (5000 microseconds), the firewall flags this as an address sweep attack and rejects all further ICMP echo requests from that host for the remainder of the threshold period. (The firewall detects and drops the tenth packet that meets the address sweep attack criteria.) Example: `set security screen MyScreen icmp ip-sweep threshold 1000`
TCP Host and Port Scans	
Host/FIN scan (TCP FIN) An attacker attempts to discover an active host or an active port on an active host. Attackers sometimes take this approach rather than a TCP SYN scan, which is more widely blocked by firewalls. They use this tactic to evade detection.	Screen for host scan (TCP FIN) `set security screen <screen-name> tcp fin-no-ack` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device blocks a TCP FIN scan by dropping the packets. Example: `set security screen MyScreen tcp fin-no-ack`
Port scan (TCP SYN) Using a single source IP address, an attacker sends IP packets containing SYNchronize segments to a number of different ports at the same destination IP address within a defined interval with the hope that one will respond. The goal is to identify a service to target.	Port Scan `set security screen <screen-name> tcp port-scan threshold <interval>` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. ■ `(threshold) interval`: Interval in microseconds. The default threshold is 5000. The valid threshold range is 1000 through 10000. Controls the number of TCP SYN packets allowed per second from one source IP address (attacker), regardless of the destination IP address. On a router with the default settings, if a remote host scans 10 ports in 0.005 seconds (5000 microseconds), the firewall flags this as a port scan attack and rejects all further packets from the remote source for the remainder of the specified time-out.The firewall detects and drops the tenth packet that meets the port scan attack criterion. Example: `set security screen MyScreen tcp port-scan threshold 1000`

Table 31: IP Address Sweep and Host and Port Scan Screen Configuration Statements (continued)

Attack	Screen Syntax
TCP SYN Using a single source IP address, an attacker sends IP packets containing SYNchronize segments to a number of different ports at the same destination IP address within a defined interval with the hope that one will respond. The goal is to identify a service to target.	SYN check `set security flow tcp-syn-check` When this filter screen is enabled, if the device receives a non-SYN TCP segment that does not belong to an existing session, it drops the packet and sends the source host a TCP RST unless the initial non-SYN TCP packet also has a RST flag set. In the latter case, it simply drops the packet.

JUNOS Software with Enhanced Services CLI Statements for Counting ICMP Ping Packets

After you have configured the firewall and have the router up and running, you can use the following CLI statements to count ICMP ping packets. You can use this information to tune screen thresholds. The CLI includes the following statements:

```
user@host# set firewall filter input-transit term ping from protocol icmp

user@host# set firewall filter input-transit term ping from icmp-type-echo-request

user@host# set firewall filter input-transit term ping from icmp-type echo-reply

user@host# set firewall filter input-transit term ping then count
count-ping-transit
```

Preventing Operating System and Network Reconnaissance Probes

An attacker can use many methods to discover which operating system is running on a host. These include use of TCP control packets and badly formed ICMP packets. JUNOS software with enhanced services provides screens to protect against abuse using these kinds of packets.

TCP Three-Way Handshake

Many attacks are perpetrated using standard TCP control packets exchanged during the three-way handshake. Figure 7 shows a state transition table for the TCP three-way handshake. Notations in the figure indicate where some of the attacks occur during the exchanges involved in setting up a session. Note the indication of packets with SYN and FIN flags set and SYN packets with data during the SYNchronize send and receive process. These packets are abnormal for the protocol and are dropped, as are packets sent during the FINish process with the FIN but no ACK flag set.

Figure 6: Using Parts of the TCP Three-Way Handshake for Reconnaissance Attacks

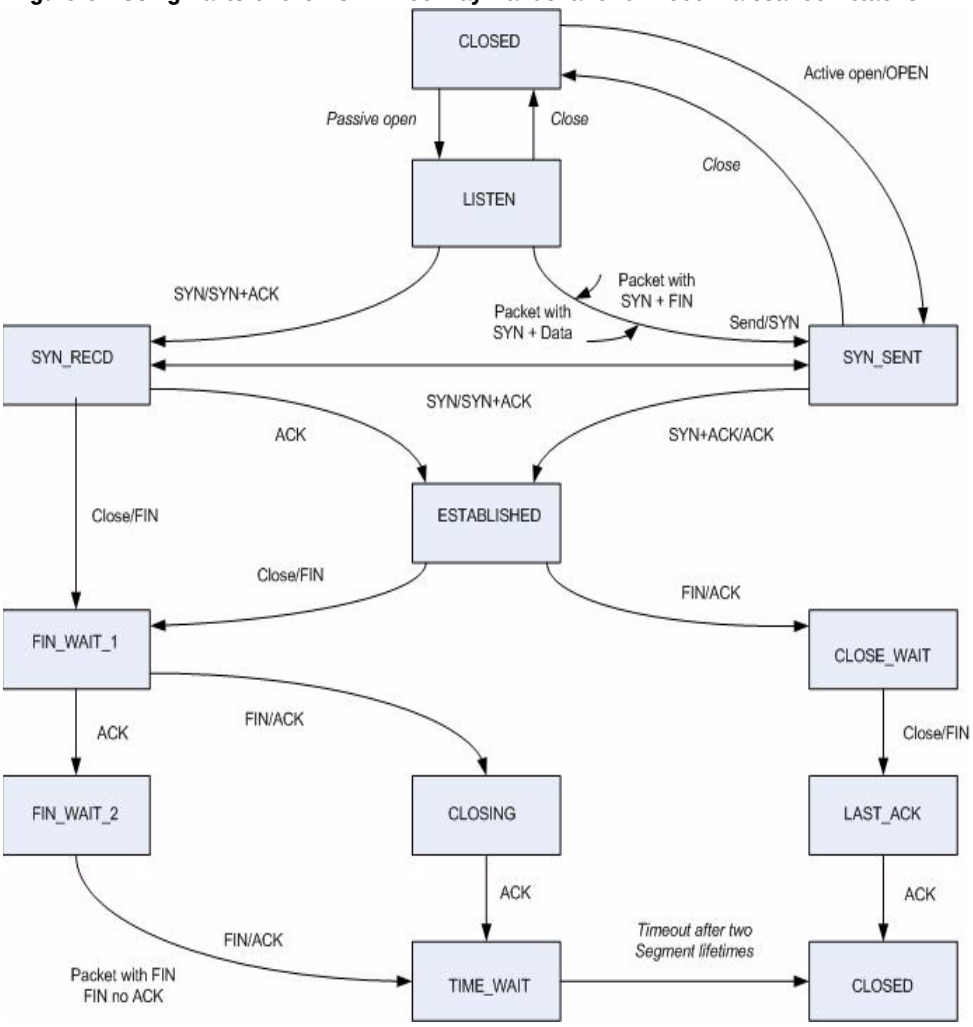

Figure 7 shows the TCP header of a packet with both SYN and FIN flags set that is used as part of an attack to discover the host's operating system. You can use the tcp syn-fin screen to defend against this kind of attack.

Figure 7: TCP Header With the SYN and FIN Flags Set

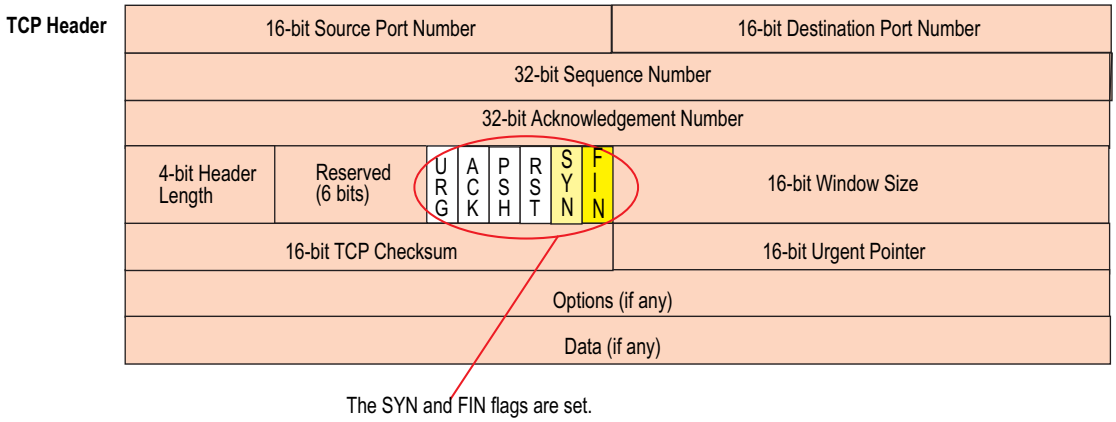

Figure 8 shows the TCP header of a packet with just the FIN flag set that is used as part of an attack to discover the host's operating system. You can use the tcp syn-fin screen to defend against this kind of attack.

Figure 8: TCP Header with Only the FIN Flag Set

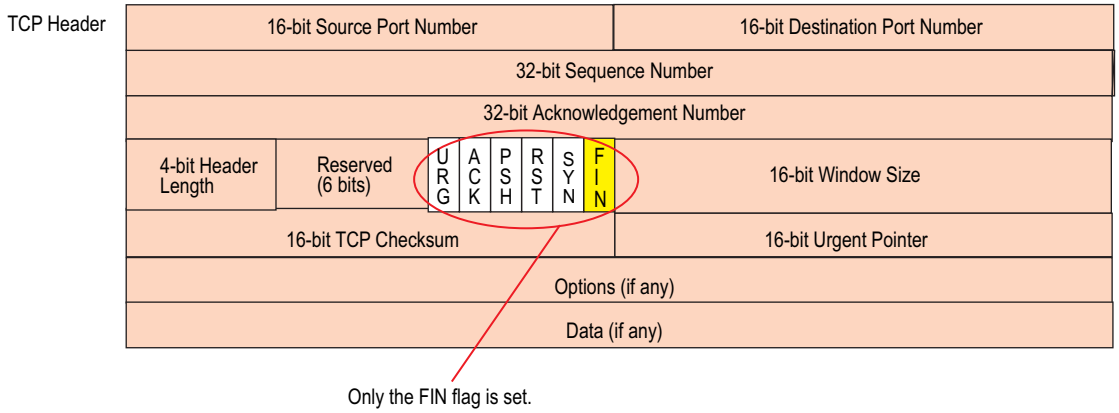

Figure 9 shows the TCP header of a packet with no flags set that is used as part of an attack to discover the host's operating system. You can use the tcp-no-flag screen to defend against this kind of attack.

Figure 9: TCP Header with No Flags Set

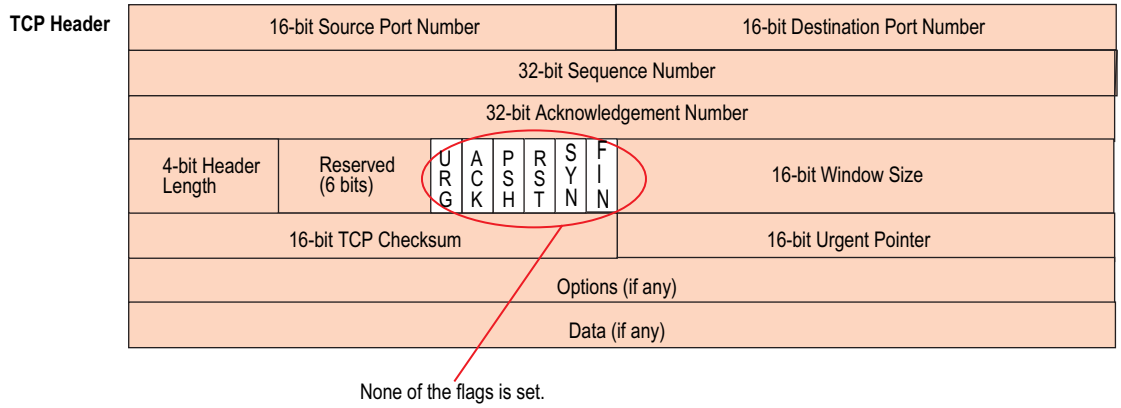

Table 32 summarizes the attacks perpetrated to discover the operating system that a host is running. It also shows the firewall screens used to defend against them. To see at which point an attack takes place during the three-way handshake, see Figure 6, "Using Parts of the TCP Three-Way Handshake for Reconnaissance Attacks," on page 86.

Table 32: Operating System Discovery and Reconnaissance Attack Screen Configuration Statements

Attack	Screen Syntax
Network Reconnaissance Using Anomalous Control Packets	
NOTE: This table shows individual statements. However, you can enter multiple tcp screen options in a single statement.	
For example: `set security screen <screen-name> tcp syn-fin fin-no-ack tcp-no-flag syn-frag`	

Table 32: Operating System Discovery and Reconnaissance Attack Screen Configuration Statements (continued)

SYN(chronize) and FIN(ish) flags set SYN begins a session. FIN ends it. Normally both flags are not set in the same TCP header. An attacker can set them and send the packet to try to get a unique response from the operating system, thereby identifying it	Screen for SYN and FIN Flag `set security screen <screen-name> tcp syn-fin` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device detects the packet and drops it. Example: `set security screen MyScreen tcp syn-fin`
FIN flag without ACK(nowledgement) flag set A TCP packet with the FIN flag set normally also has the ACK flag set. An attacker can send a packet without ACK to see how the system responds, thereby identifying it.	Screen for FIN no ACK `set security screen <screen-name> tcp fin-no-ack` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device detects the packet and drops it. Example: `set security screen MyScreen tcp fin-no-ack`
TCP header with no control flags Normally a TCP packet header has at least one control flag set. An attacker can send a packet with no flags set to see how the system responds, thereby identifying it.	Screen for TCP header with no flags set `set security screen <screen-name> tcp tcp-no-flag` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device detects the packet and drops it. Example: `set security screen MyScreen tcp tcp-no-flag`

SYN fragments
 	
The SYN packet does not usually contain data because its purpose is to initiate a session connection and invoke a SYN/ACK segment in response. Because the packet is small, there is no legitimate reason for it to be fragmented.
	Screen for SYN packets containing data
	
	`set security screen <screen-name> tcp syn-frag`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	
	The device detects the fragment, drops it, and records the event for the ingress interface.
	
	`Example: set security screen MyScreen tcp syn-frag`

Network reconnaissance using IP options.	

NOTE: This table shows individual statements. However, you can enter all of the IP screen options on a single statement. For example:

```
set security screen <screen-name> ip record-route-option timestamp-option security-option stream-option
block-frag unknown-protocol filter-source loose-source-route-option strict-source-route-option
```

The alarm-without-drop statement allows you to monitor but not block attacks to obtain information about the attacker’s methods, strategies, and objectives. See the *JUNOS Software Security Configuration Guide* for details.

Table 32: Operating System Discovery and Reconnaissance Attack Screen Configuration Statements (continued)

Anomalous use of IP options
 	
Attackers can set these rarely used IP options flags to malicious ends.
	Screen for reconnaissance using IP options
	
	`set security screen <screen-name> ip record-route-option`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device drops the packet if alarm-without-drop is not set.
	
	`set security screen <screen-name> ip timestamp-option`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device drops the packet if alarm-without-drop is not set.
	
	`set security screen <screen-name> ip security-option`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device records the event and does not drop the packet.
	
	`set security screen <screen-name> ip stream-option`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device records the event and does not drop the packet.
	
	Example: `set security screen MyScreen ip record-route-option timestamp-option security-option stream-option`

Reconnaissance attacks using bad IP options An attacker can misuse any of a group of eight rarely used IP packet header options flags to nefarious ends.	Screen for bad IP options `set security screen <screen-name> ip bad-option` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device detects and drops packets when any IP option in the IP packet header is incorrectly formatted, if alarm-without-drop is not set. Example: `set security screen MyScreen ip bad-option`
Attacks using unknown IP protocols Some protocols types in the IP header are reserved and unassigned. Attackers can use the fields, but there is no way to know in advance if a particular packet of this type is benign or malicious.	Screen for unknown IP protocols `set security screen <screen-name> ip unknown-protocol` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device drops the packet when the protocol field contains a protocol ID number of 137 or greater. Example: `set security screen MyScreen ip unknown-protocol`

Table 32: Operating System Discovery and Reconnaissance Attack Screen Configuration Statements (continued)

Attacks using IP fragments As packets traverse different networks, it is sometimes necessary to break a packet into smaller fragments based on the MTU size of a network. IP fragments might indicate an attacker's attempts to exploit vulnerabilities in the packet reassembly code of specific IP stack implementations. When the system receives these altered packets the results can range from processing the packets incorrectly to crashing the entire system.	Screen for altered IP fragments `set security screen <screen-name> ip block-frag` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. The device blocks all IP packet fragments incorrectly formatted. Example: `set security screen MyScreen ip block-frag`
ICMP fragments and large ICMP packets	
NOTE: This table shows individual statements. However, you can enter both of the ICMP screen options on a single statement. For example: `set security screen <screen-name> icmp fragment large`	
Attacks using fragmented ICMP packets Because ICMP packets contain very short messages (for error reporting and network probes mostly), there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, there is a good chance that an attack is underway.	Screen for ICMP fragments `set security screen <screen-name> icmp fragment` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. The device blocks and drops ICMP packets with the More Fragments flag set or with an offset value indicated in the offset field. Example: `set security screen MyScreen icmp fragment`
Attacks using large ICMP packets Because ICMP packets are very short messages, there is no legitimate reason for large ICMP packets to be sent. If the packet is large, there is a good chance that an attack is underway.	Screen for large size ICMP packets `set security screen <screen-name> security icmp large` The device blocks packets with a length greater than 1024 bytes. ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. Example: `set security screen MyScreen icmp large`

Accountability Evasion

Accountability is the ability to connect an event to the party who initiated it. Because for TCP/IP networks the origin of an event is expressed as a source IP address, the evasion of accountability can be carried out through use of a technique called IP spoofing.

Once attackers find their way into your network, they can wreak havoc on it in many different ways. For attacks whose mission is to bring down your network, crowd out users from connecting from one system to another, or limit user access to resources, no information need be returned to the attacker. In these cases, the attacker can hide behind a spoofed IP address and leave without having been identified by all but the most sophisticated security tools and the help of your service provider. IP spoofing is used when the attacker does not require returned information to carry out the attack.

IP spoofing is a process in which the attacker inserts a bogus source address in a packet header to make the packet appear to come from a trusted source. Intruders use IP spoofing as part of many of the attacks that they launch, in particular DoS attacks. Table 33 shows some of the kinds of accountability evasion techniques that attackers use.

Table 33: Evasion Techniques Screens Configuration Statements

TYPE OF ATTACK	SCREEN

IP spoofing
Attempt to insert a bogus source address in the packet header to make the packet appear to come from a trusted source.
	Screen for IP Spoofing
	`set security screen <screen-name> ip spoofing`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device uses the route table to determine the veracity of an IP address, and it discards packets containing spoofed IP addresses and drops them.
	`set security screen MyScreen ip spoofing`

IP source route options
Attackers can use IP source route options to hide their true addresses and access restricted areas of a network by specifying a different path,
	Screen for IP source route options
	`set security screen <screen-name> ip loose-source-route-option`
	■ `screen-name`: `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device detects and records packets with a loose source route.
	`set security screen <screen-name> ip strict-source-route-option`
	■ `screen-name`: Name of the screen to be applied to a zone.
	The device detects and records packets with a strict source route.

Configuring the Firewall to Protect Against Penetration Attacks: Firewall and Network DoS Attacks

There are different kinds of penetration attacks. They include

- Confidentially attacks, in which the intruder attempts to take private information.
- Integrity attacks, in which the intruder attempts to alter the state of the network or its hosts in some way.
- Availability attacks, in which the intruder attempts to overwhelm the system in some manner to lock out legitimate users from accessing resources they require.

About Denial-of-Service (DoS) Attacks

A denial-of-service (DoS) attack is an attempt by an intruder either to prevent legitimate users from accessing resources or to bring the network to a standstill in some way. JUNOS software with enhanced services provides screens to protect against DoS attacks launched against the network and against the firewall.

The most common type of DoS attack occurs when the attacker floods the network with information that consumes limited resources. A DoS attack can consume the system's available session table data structures, its processing power and memory and, its network bandwidth. A successful DoS attack stops the transmission of legitimate traffic.

Protecting the Firewall Against Session Table Floods

If the intruder discovers the presence of a firewall, he or she might launch a DoS attack against the firewall itself. A successful DoS attack against the firewall amounts to a successful DoS attack against the protected network because it prohibits legitimate traffic from passing through the firewall.

About Screens for Session Table Floods

The firewall provides several screens, described in Table 34, that you can use to protect against session table floods. However, before you can define screens to thwart session table floods, you need to understand the normal number of session connections for your environment.

Determining what constitutes an acceptable number of session connection requests requires a period of observation. You must first

- Establish a baseline for normal traffic flows.

  For information on establishing a network baseline, see "Developing a Network Profile" on page 20.

- Consider the maximum number of concurrent sessions that would exhaust the session table of the router that you are using.

  For details on the supported number of sessions, see the data sheet for your router.

Using Aggressive Aging to Protect Against Session Table Flood

By default, an initial TCP session three-way handshake takes 20 seconds to time out. After a TCP session has been established, the timeout value changes to 30 minutes. The timeout for HTTP sessions is 5 minutes. The timeout for UDP sessions is one minute.

Aggressive aging accelerates the timeout process for sessions (or, shortens the default session timeout) when the number of sessions in the session table exceeds the specified high-watermark percentile. When the number drops below the specified low-watermark percentile, the timeout process returns to normal.

During the period when the aggressive aging process is in effect, the router decreases the timeout for all sessions by the aging-out rate that you specify. The device ages out the oldest sessions first. It does this until the number of sessions in the table is under the low-watermark threshold.

Based on the following statement, the router takes the following actions:

user@host# set security flow aging early-ageout 4 high-watermark 80 low-watermark 60

- Decreases the timeout value for all sessions by a rate of 4.
- Begins aging out sessions when the number of sessions in the session table exceeds 80 percent.
- Stops aging out sessions aggressively when the number of sessions in the session table drops below 60 percent.

Figure 10 illustrates how aggressive aging works.

Figure 10: Aggressive Aging

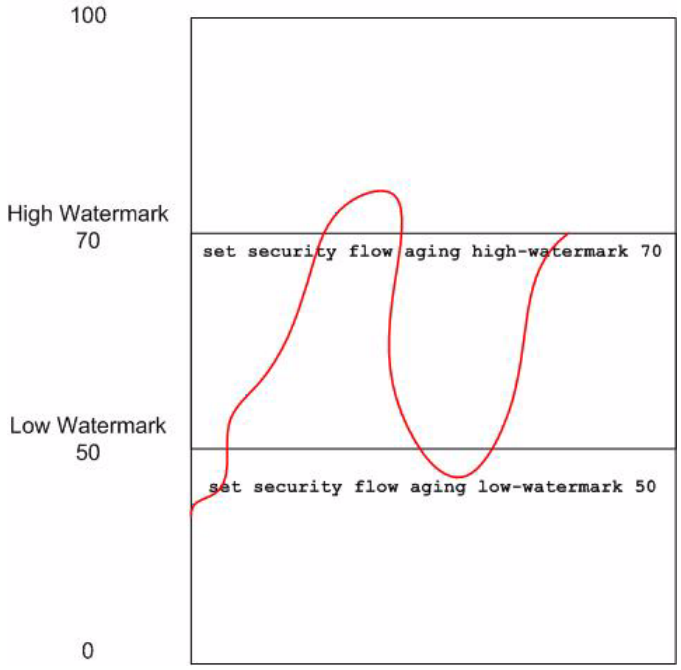

Session Table Flood from Incomplete SYN-ACK-ACK

When an authentication user initiates a Telnet or FTP connection, the following events take place:

- The user's software sends a SYN segment to the Telnet or FTP server.
- JUNOS software with enhanced services software intercepts the SYN segment, creates an entry in its session table, and proxies a SYN-ACK segment to the user.
- The user replies with an ACK segment.

At this point, the initial three-way handshake is complete and the device sends a login prompt to the user. If the user does not log in, but instead continues initiating SYN-ACK-ACK sessions, the firewall table can fill up to the point where the device begins rejecting legitimate connection requests. This is one example of how a network DoS attack can occur.

A network DoS attack 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.

Depending on the attacker's purpose and the extent and success of previous intelligence gathering efforts, the attacker might single out a specific host. Alternatively, the attacker might aim at random hosts across the targeted network. Either approach has the potential of upsetting service.

Figure 11 shows DoS attacks launched using session connections and how the limit-session screen can be used to protect against the attack.

Figure 11: Screens Used to Defend Against DoS Attacks Based on Session Connections

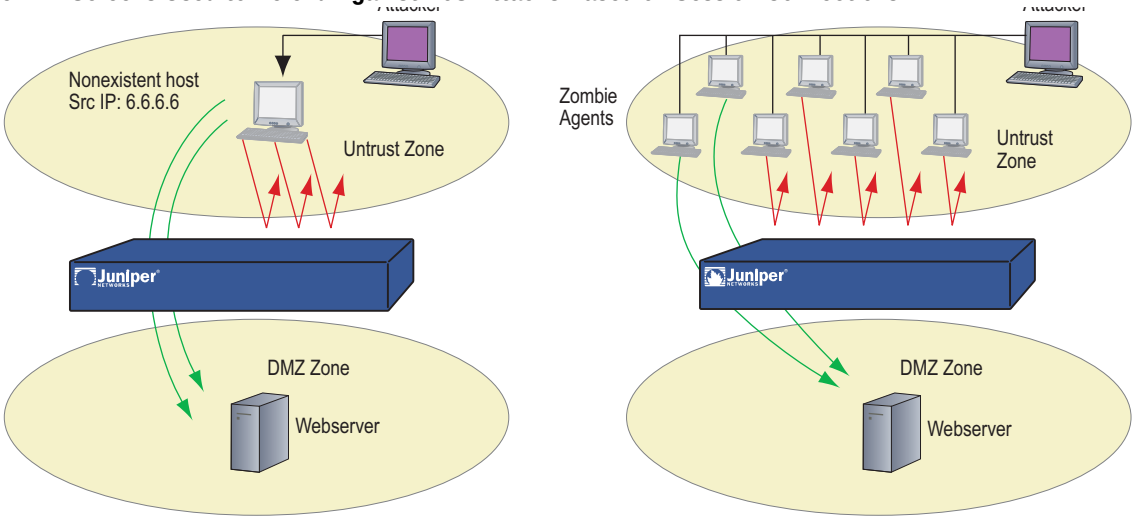

Table 34 identifies the router's screens that protect against firewall session table floods

NOTE: This table shows individual statements. However, you can enter both rate-limiting statements to stem session table flood on a single statement. For example:

set security screen <screen-name> limit-session source-ip-based <maximum-concurrent-sessions> destination-ip-based <maximum-concurrent-sessions>

Table 34: Firewall Session Table Flood Denial-of-Service Attack Screen Configuration Statements

Attack	Screen Syntax
Flooding the firewall's session table	
Attacks using session connections (coming from the same source address) An attacker attempts to overwhelm the firewall's session table with invalid session connection requests to many hosts (IP addresses) from a single source IP address.	Rate-limiting session connections based on source addresses `set security screen <screen-name> limit-session source-ip-based <maximum-concurrent-sessions>` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. ■ `maximum-concurrent-sessions`: Sets a limit on how many IP session connection requests are permitted from the same source IP address. For example, setting a source-based session limit can stem an attack, such as the Nimda virus, which infests a server and then begins generating massive amounts of traffic from it. The router denies session connection requests after the limit is reached. It drops packets for the remainder of the second until the source of the session requests begins to close session connections. Example: `set security screen MyScreen limit-session source-ip-based 100`
Attack using session connections (to the same destination address) From many zombie hosts to a single destination host, an attacker attempts to overwhelm the firewall's session table with invalid session connection requests to tie up its resources.	Rate-limiting session connections based on destination IP addresses `set security screen <screen-name> limit-session destination-ip-based <maximum-concurrent-sessions>` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. ■ `maximum-concurrent-sessions`: Sets limits on how many IP session connection requests are permitted to a single destination IP address. An attacker can launch a distributed denial-of-service (DDoS) attack from hundreds of zombie hosts. Setting a destination-based session limit defends against this attack. The device denies session-connection requests after the limit is reached. It drops packets for the remainder of the second until the source of the session requests begins to close session connections. Example: `set security screen MyScreen limit-session destination-ip-based 100`

Attack using SYN-ACK-ACK to cause session table flood	SYN-ACK-ACK proxy screen
	`set security screen <screen-name> tcp syn-ack-ack-proxy <threshold>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80.
	■ `threshold`: When the number of connections from the same IP address reaches the specified SYN-ACK-ACK proxy threshold, The device rejects further requests from that IP address.
	`set security screen MyScreen tcp syn-ack-ack-proxy 1000`

About SYN Flood Attacks

A SYN flood occurs when a host becomes so overwhelmed by SYN segments initiating incomputable connection requests that it can no longer process legitimate requests. Here are the events that constitute a SYN flood attack:

- A SYN flood attack inundates a site with SYN segments containing forged or spoofed IP source addresses that are unreachable.
- The victim host responds to these addresses with SYN/ACK segments and then waits for ACK segments from the source.
- Because the SYN/ACK segments are sent to nonexistent or unreachable IP addresses, they never elicit responses and the connections eventually time out.

Preventing SYN Flood Attacks: SYN Proxying

You use the SYN flood screen to protect against SYN flood attacks. The screen protects against SYN floods by proxying the SYN connection requests. Here is how it works:

- You can set a limit on the number of SYN segments that are permitted to pass through a zone per second.
- You can set an 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 the thresholds, the firewall begins to proxy incoming SYN segments, replying with SYN/ACK statements 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.
  - When the proxied connection queue is completely filled up, the firewall rejects new incoming SYN segments. This action shields hosts on the protected network from the bombardment of incomplete three-way handshakes.

About SYN Cookie

SYN cookie is a stateless SYN proxy mechanism that you can use to protect against a SYN flood attack. Because SYN cookie is stateless, it does not set up a session or do policy and route lookups when a SYN segment is received, and it maintains no connection request queues. This results in reduced CPU and memory usage.

Here is how it works:

- SYN cookie is activated when the SYN flood attack threshold is exceeded.
- The SYN cookie screen detects and protect against spoofed SYN flood attacks and thus minimizes negative impact to hosts that are secured by the firewall.

  If an attacker is using a legitimate source IP address rather than a spoofed one, the screen will not defend against the attack.

Table 35 gives details on how to configure the firewall for SYN proxy and SYN cookie.

Table 35: Network DoS Attack Screen Configuration Statements

Attack	Screen Syntax
Screen protects against SYN flood attacks using SYN proxying or SYN cookie for session connection attempts NOTE: When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection and source-based SYN flood tracking into effect.	

SYN flood attacks
An attacker floods a host with TCP connection attempts that cannot be completed.
Eventually the memory buffer of the target host fills up. The host becomes overwhelmed by SYN segments initiating uncompleteable connection requests that it can no longer process legitimate requests.
	SYN flood screen
	`set security screen <screen-name> tcp syn-flood`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80.
	Enables the SYN flood defense feature. By default, SYN proxy is used.
	`set security screen MyScreen tcp syn-flood`

	SYN flood attack threshold screen
	`set security screen <screen-name> tcp syn-flood attack-threshold <number>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80.
	■ `(attack-threshold) number`: Number of SYN segments to the same destination address and port number per second required to activate SYN proxying.
	Understand your normal network traffic patterns before you set the threshold. For example. if the firewall normally gets 2000 SYN segments per second, you might want to set the threshold at 3000 packets per second.
	`Example: set security screen MyScreen tcp syn-flood attack-threshold 1000`

	SYN flood alarm threshold screen `set security screen <screen-name> tcp syn-flood alarm-threshold <number>` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. ■ `(alarm-threshold) number`: Number of proxied, half-complete TCP connection requests per second to the same destination IP address and port number after which the system enters an alarm in the event log. For example, if you set the SYN attack threshold at 200 SYN segments per second, then a total of 301 SYN segments to the same destination IP address and port number per second is required to trigger an alarm in the event log. Example: `set security screen MyScreen tcp syn-flood alarm-threshold 512`
	Screen for SYN flood based on source of attack `set security screen <screen-name> tcp syn-flood source-threshold <number>` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80. ■ `(source-threshold) number`: Number of SYN segments received per second from a single source IP address before the firewall begins dropping connection requests from that source. Example: `set security screen MyScreen tcp syn-flood source-threshold 4000`

Table 35: Network DoS Attack Screen Configuration Statements (continued)

Attack	Screen Syntax

	Screen for SYN flood based on destination of attack
	`Example: set security screen <screen-name> tcp syn-flood destination-threshold <number>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80.
	■ `(destination-threshold) number`: Number of SYN packets received per second for a single destination IP address before the firewall begins dropping connection requests to that destination.
	`Example: set security screen MyScreen tcp syn-flood destination-threshold 4000`

	Session timeout for SYN flood screen
	`set security screen <screen-name> tcp syn-flood timeout <seconds>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80.
	■ `(timeout) seconds`: Maximum amount of time before a half-completed connection is dropped from the queue.
	`Example: set security screen MyScreen tcp syn-flood timeout 30`

	Queue size for SYN flood screen
	`set security screen <screen-name> tcp syn-flood queue-size <number>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see "About Creating Screens" on page 80.
	■ `(queue-size) number`: Number of proxied connection requests held in the proxied connection queue before the firewall begins rejecting new connection requests.
	`Example: set security screen MyScreen tcp syn-flood queue-size 1024`

	SYN cookie protection screen
	`set security flow syn-proxy syn-cookie`
	Sets the mode from traditional SYN proxy to SYN cookie. Builds on SYN proxy.
	`Example: set security flow syn-flood-protection-mode syn-cookie`

Preventing ICMP Flood Attacks

An ICMP flood occurs when a router is overloaded with so many ICMP echo requests that it expends all of its resources responding to the requests and can no longer process valid network traffic. You can use the ICMP flood screen summarized in Table 36 to set a threshold which when exceeded causes the firewall to ignore further ICMP echo requests for the remainder of that second, plus the next second. The default threshold value is 1000 packets per second.

Table 36: ICMP Flood Screen Configuration Statement

Attack	Screen Syntax

ICMP flood attack
Attacker sends IP packets containing UDP datagrams to slow down the firewall until it can no longer handle valid connections.
	ICMP flood protection screen
	`set security screen <screen-name> icmp flood threshold <number>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	■ `(threshold) number`: When the specified value is exceeded, ICMP flood attack prevention is initiated. The default value is 1000 packets per second.
	If the threshold is exceeded, The device ignores further ICMP echo requests for the remainder of that second and for the next second.
	`Example: set security screen MyScreen icmp flood threshold 1500`

Preventing UDP Flood Attacks

UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the router to the point that it can no longer handle valid connections. To protect against UDP flooding, you can set a threshold for the UDP flood protection screen. If the number of UDP datagrams from one or more sources to a single destination exceeds this threshold, JUNOS software with enhanced services ignores further UDP datagrams to that destination for the remainder of that second plus the next second.

With the UDP flood protection feature enabled, you can set a threshold that once exceeded invokes the UDP flood attack protection feature. The default threshold value is 1000 packets per second.

Table 37: UDP Flood Screen Configuration Statement

Attack	Screen Syntax

UDP flood attack
Attacker sends IP packets containing UDP datagrams to slow down the firewall to the point where it can no longer handle valid connections.
	UDP flood protection screen.
	`set security screen <screen-name> udp flood threshold <number>`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	■ `(threshold) number`: When the specified value is exceeded, UDP flood attack prevention is initiated. The default value is 1000 packets per second.
	If the number of UDP datagrams from one or more sources to a single host exceeds this threshold, the firewall drops further UDP datagrams to that destination for the remainder of the second and the next second.
	`Example: set security screen MyScreen udp flood threshold 1500`

Preventing 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 target as both the destination and source IP address. The receiving system responds by sending the SYN-ACK packets to itself, creating an empty connection that lasts until the idle timeout value is reached. Flooding the system with such empty connections can overwhelm the system, causing a DoS attack.

Table 38: Land Attack Screen Configuration Statement

Attack	Screen Syntax

Land Attack
Protects against an attack that combines a SYN attack with IP spoofing to overwhelm the target host and tie up session connections.
	`Land attack protection screen`
	`set security <screen-name> tcp land`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	`Example: set security MyScreen tcp land`

Preventing Operating System DoS Attacks

The following penetration attacks attempt to crash the operating system on the router. Firewall screens can protect against these attacks.

Preventing Ping of Death Attacks

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, that is 8 bytes long. Therefore the maximum allowable size of the data area of an ICMP echo request is 65,507 bytes. However, many ping implementations allow the user to specify a packet size larger than 65,507 bytes. Table 39 summarizes how the firewall responds when the ping of death protection screen is enabled.

Table 39: Ping of Death Attack Screen Configuration Statement

Attack	Screen Syntax
Ping of Death Attacker sends grossly oversized ICMP packets to trigger a range of adverse system reactions such as denial of service (DoS), ultimately crashing, freezing, and rebooting the system.	Ping of Death attack protection screen `set security screen <screen-name> icmp ping-death` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. With this option enabled, the firewall detects and rejects such oversized and irregular packet sizes even when the attacker hides the total packet size by purposefully fragmenting it. Example: `set security screen MyScreen icmp ping-death`

Preventing Teardrop Attacks

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 of the data contained in a fragmented packet relative to the data of the original packet. When the sum of the offset and the size of one fragmented packet differ from that of the next fragmented packet, the packets overlap. In this case, the system attempting to reassemble the packet can crash, especially if it is running an older operating system that has this vulnerability.

Table 40: Teardrop Attack Screen Configuration Statement

Attack	Screen Syntax

Teardrop attack
Attacker attempts to crash the system by creating fragmented packets that overlap.
	Screen to protect against teardrop attack
	`set security screen <screen-name> ip tear-drop`
	■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80.
	The device detects the discrepancy in the fragmented packet and drops it.
	`set security screen MyScreen ip tear-drop`

Preventing WinNuke Attacks

WinNuke is a DoS attack that targets any computer on the Internet running Microsoft 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. This introduces a NetBIOS fragment overlap, which causes many machines running Windows to crash. After the attacked machine reboots, the following message appears, indicating that an attack has occurred:

```
An exception of OE has occurred at 0028:[address] in VxD MSTCP (01) AE.
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.
```

Table 41 summarizes the WinNuke attack and the firewall’s response when you configure defense against it.

Table 41: WinNuke Attack Configuration Screen

Attack	Screen Syntax
WinNuke Attack Attacker attempts a DoS attack by sending to a host with an established connection a TCP segment with the URG flag set. This segment is usually to NetBIOS port 139. This introduces NetBIOS fragment overlap.	Screen to protect against a WinNuke attack `set security screen <screen-name> tcp winnuke` ■ `screen-name`: Name of the overall screen to be applied to a zone. For details, see “About Creating Screens” on page 80. The device scans any incoming Microsoft NetBIOS session service (port 139) packets. If it detects 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 logs the attempted WinNuke attack event. Example: `set security screen MyScreen tcp winnuke`

Monitoring for Attacks

Normally you want to use the firewall to secure your network and block exploits and attacks against it. However, there might be times when you want to gather information about attacks, rather than to block them. JUNOS software with enhanced services provides the alarm-without-drop statement that allows an attack to take place. You can then study what has occurred to try to understand the attacker's method, strategy, and objectives. Gathering this information helps you to understand the kinds of attacks that threaten your network. With this information, you can better fortify your defenses. For details, see the *JUNOS Software Security Configuration Guide*.

Testing the Firewall

Organizations take different approaches to testing their configurations before deploying their firewalls. Some organizations create a firewall configuration and test the product under lab conditions, using a packet generator. Some organizations contend that packet generators do not introduce network traffic noise and dirty packets that are part of normal network activity, and, therefore, cannot accurately indicate the response of the firewall to actual events. To best test the firewall in your lab, capture your normal traffic and replay it in the lab before you deploy the firewall.

Because true network conditions cannot be emulated, some companies use a split optical line to test the firewall under real-time traffic conditions. Others shut down business temporarily at off hours by calling an outage, and then test the firewall during that period.

Chapter 5

Implementing Route-based IPSec VPNs for Branch Offices

This chapter describes how to implement route-based IPSec VPNs for branch offices. It uses a deployment scenario for a mid-size branch office. The deployment plans includes two IPSec VPNs to provide better throughput and to ensure that one VPN to headquarters is available if the other fails. Apart from the interfaces and the service providers used, the VPNs are the same. For this reason, the chapter explains the configuration for only one of them.

First, the chapter describes how to implement a basic IPSec VPN with the minimum requirements. Then it builds on the basic VPN, adding features that enhance its capability.

NOTE: Chassis clustering is supported for route-based IPSec VPNs. For details, see the *JUNOS Software Security Configuration Guide.*

For background information and design concepts, see “Designing IPSec VPNs for Branch Offices” on page 29.

This chapter includes the following sections:

About Route-based IPSec VPNs

Unlike for policy-based IPSec VPNs, for route-based IPSec VPNs a policy refers to a destination address not an IPSec VPN tunnel. When the JUNOS software with enhanced services looks up a route to find the interface to use to send traffic to the packet's destination address, it finds a route via a secure tunnel interface (st-x). The tunnel interface is bound to a specific IPSec VPN tunnel, and the traffic is routed to the tunnel if the policy action is permit. The policy dictates whether a packet from a source address containing a payload of a certain application type (or service type) is to be delivered to the destination secure tunnel interface and on through the IPSec VPN tunnel.

With route-based IPSec VPNs, dynamic routing information can be exchanged through the VPN tunnel. You can enable an instance of a dynamic routing protocol, such as Border Gateway Protocol (BGP), on a tunnel interface that is bound to a VPN tunnel. The local routing instance exchanges routing information through the tunnel with a neighbor routing instance enabled on a tunnel interface bound to the other end.

About the Route-based IPSec VPN Branch Deployment

This scenario describes how to deploy route-based IPSec VPNs for the Albuquerque mid-size branch office of the New Bank of the Southwest. The branch users will use the VPNs to connect to corporate headquarters (HQ) at the Los Angeles site.

The deployment includes two IPSec VPNs with two separate service providers. A primary and a secondary route are configured. The router maintains both routes with the secondary VPN route configured as the less preferred one. Consequently, traffic is routed through the secondary tunnel automatically only if the primary tunnel fails. Because configuration of the primary and secondary VPNs is similar, this chapter describes how to configure the primary one only.

Figure 12 shows the topology for this deployment.

Figure 12: Topology for the Albuquerque Branch Office IPSec VPNs

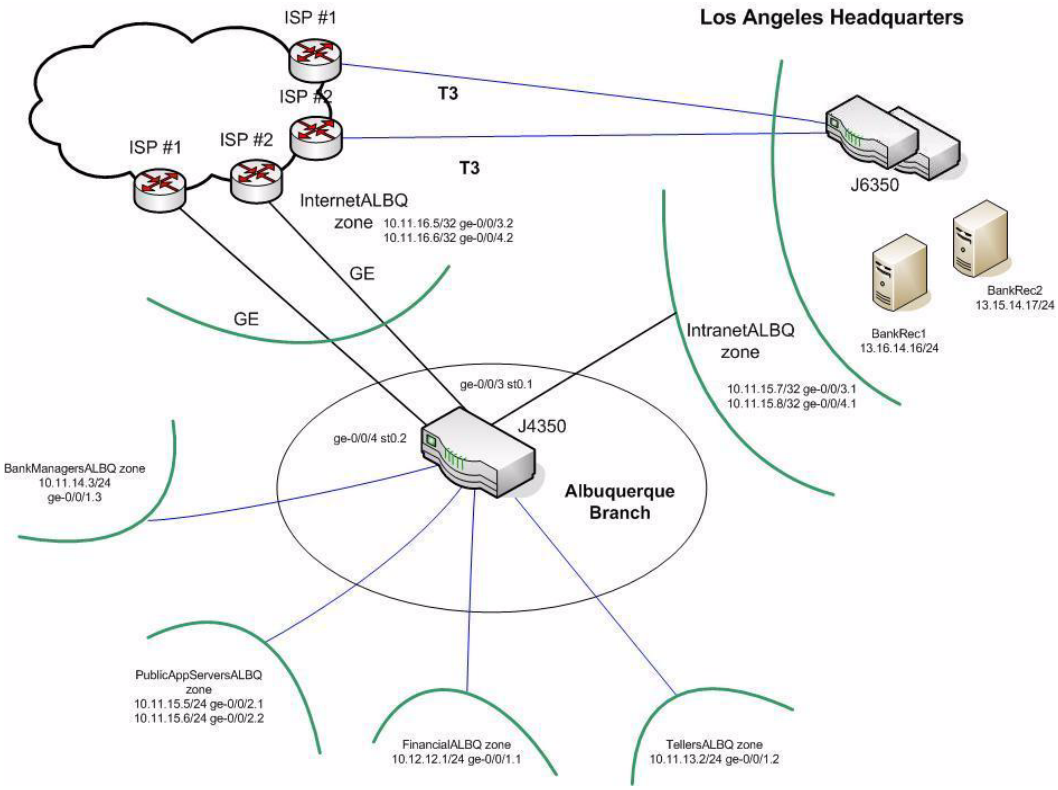

Configuring Interfaces for the VPNs

For route-based IPSec VPNs, you must first configure the tunnel interface and assign it to a zone. Then you can bind the VPN to the tunnel interface in the IPSec VPN configuration.

To configure the interface to be used for the VPN, enter the following statement in configuration mode:

```
user@host# set interfaces st0 unit 1
```

To add the interface to the IntranetALBQ zone, enter the following statement in configuration mode:

```
user@host# set security zones security-zone IntranetALBQ
    interfaces st0.1
```

Tasks for Implementing Basic Route-based IPSec VPN

To create a basic IPSec VPN consisting of required features only, you complete the tasks shown in Table 42.

Table 42: Procedure for Creating a Basic Route-based IPSec VPN

Task	Instructions
1. Configure the required security zones and interfaces.	See “Implementing Firewall Deployments for Branch Offices” on page 37.
2. Configure addresses for inbound and outbound policies to be used in proxy IDs.	See “Implementing Firewall Deployments for Branch Offices” on page 37.
3. Configure static routes.	See “Implementing Firewall Deployments for Branch Offices” on page 37.
4. IKE Phase 1: Configure the peer gateway information.	See “Configuring the Peer Gateway and Destination Peer Information” on page 109.
5. IKE Phase 1: Configure a proposal.	See “Configuring IKE Phase 1 Proposals” on page 113.
6. IKE Phase 1: Configure a policy.	See “Configuring IKE Phase 1 Policies” on page 110.
7. IPSec Phase 2: Configure the IPSec VPN.	See “Configuring the IPSec VPN” on page 117.
8. IPSec Phase 2: Configure a proposal.	See “Configuring IPSec Phase 2 Proposals” on page 119.
9. IPSec Phase 2: Configure a policy.	See “Configuring IPSec Phase 2 Policies” on page 123.
10.Configure firewall security policies.	See “Implementing Firewall Deployments for Branch Offices” on page 37.

The deployment example shows how to configure the VPN using statements entered in configuration mode. You can also enter the statements at the [edit security] hierarchy, in which case you do not need to specify “security” as part of the statement.

Configuring a Basic IPSec VPN

Employees at the Albuquerque branch office require access to servers at corporate headquarters in Los Angeles. This section describes how to create a basic VPN with minimum requirements to be used for that purpose.

Route-based IPSec VPN Statements

This section shows the CLI commands used to configure the basic IPSec VPN. The complete configuration is shown after the CLI statements.

CAUTION: To use this sample configuration, you must commit the configuration after all the statements are entered. If you intend to commit the statements individually, you must reverse the order of statements within each phase because of dependencies among the statements.

Configuring Basic Route-based IPSec VPN Configuration Using the CLI

IKE Phase 1

Destination peer gateway name and peer address. (See "Configuring the Peer Gateway Name and Destination Peer Address" on page 109.)	user@host# set security ike gateway los-angeles-gw3 address 5.2.3.4
Phase 1 IKE policy reference. (See "Configuring the IKE Phase 1 Policy Reference" on page 109.)	user@host# set security ike gateway los-angeles-gw3 ike-policy pol1
Phase 1 gateway external interface. (See "Configuring the IKE Phase 1 External Interface" on page 110.)	user@host# set security ike gateway los-angeles-gw3 external-interface ge-0/0/3.2
Phase 1 IKE policy. (See "Configuring IKE Phase 1 Policies" on page 110.)	user@host# set security ike policy pol1 pre-shared-key ascii-text 1234pskey5678 user@host# set security ike policy pol1 proposals prop1
Phase 1 IKE proposal. (See "Configuring IKE Phase 1 Proposals" on page 113.)	user@host# set security ike proposal prop1 authentication-method pre-shared-keys dh-group group2 encryption-algorithm aes-256-cbc authentication-algorithm md5

Phase 2 IPSec VPN

IPSec VPN name. (See "Naming the IPSec VPN Gateway" on page 117.)	user@host# set security ipsec vpn albuquerque-vpn1
IPSec Phase 2 Peer gateway name (See "Identifying the Peer Gateway" on page 118.)	user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3
IPSec Phase 2 policy reference (See "Configuring the Reference to the IPSec Phase 2 Policy" on page 118.)	user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3 ipsec-policy ipsec-pol1
Binding the IPSec VPN to a tunnel interface (See "Binding the IPSec VPN to a Tunnel Interface" on page 119.)	user@host# set security ipsec vpn albuquerque-vpn1 bind-interface st0.1

IPSec Phase 2 proposal (See “Configuring IPSec Phase 2 Proposals” on page 119.)	user@host# set security ipsec proposal ipsec-prop1 protocol esp authentication-algorithm hmac-md5-96 encryption-algorithm aes-256-cbc
IPSec Phase 2 policies. (See “Configuring IPSec Phase 2 Policies” on page 123.)	user@host# set security ipsec policy ipsec-pol1 proposals ipsec-prop1

Basic Route-based IPSec VPN Configuration

The following example shows the same configuration shown in the set statements.

```
security {
         ike {
            gateway los-angeles-gw3 {
                address 5.2.3.4;
                ike-policy pol1
                external-interface ge-0/0/3.2
            }
            policy pol1
                pre-shared-key ascii-text 1234pskey5678;
                proposal prop1;
            }
            proposal prop1 {
                authentication-method pre-shared-keys;
                dh-group group2;
                encryption-algorithm aes-256-cbc;
                authentication-algorithm md5;
            }
        }
        ipsec {
            vpn albuquerque-vpn1 {
                ike {
                    gateway los-angeles-gw3;
                    ipsec-policy ipsec-pol1;
                    }
                }
                bind-interface st0.1
            }
            proposal ipsec-prop1 {
                protocol esp;
                authentication-algorithm hmac-md5-96;
                encryption-algorithm aes-256-cbc;
            }
            policy ipsec-pol1 {
                proposals ipsec-prop1;
            }
        }
    }
```

Configuring the Peer Gateway and Destination Peer Information

As part of Phase 1 for an IPSec VPN, you configure a gateway between the local peer and its destination (or partner) peer. This configuration identifies the destination peer so that the local one can communicate with it.

Configuring the Peer Gateway Name and Destination Peer Address

```
security —-> ike —->gateway
            <destination-peer-gateway-name>
            address <destination-peer-addr>
```

NOTE: The destination peer gateway name specified in Phase 1 must match the destination peer gateway name specified in Phase 2.

To identify the destination peer gateway (los-angeles-gw3) for the basic IPSec VPN configuration and specify the peer's address, enter the following command in configuration mode:

```
user@host# set security ike gateway los-angeles-gw3 address 5.2.3.4
```

Table 43 defines the statement parameters.

Table 43: IPSec VPN Peer Gateway Name and Address Configuration Statement

Destination Peer Gateway Identity	Syntax
Destination peer gateway name and address	`set security ike gateway <destination-peer-gateway-name> address (<destination-peer-addr> | <hostname> . . .)` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `address`: Address of the destination peer gateway, specified as either an IPv4 address or hostname.

Configuring the IKE Phase 1 Policy Reference

```
security —-> ike —-> gateway
                    <destination-peer-gateway-name>
                        ike-policy <ike-policy-name>
```

For the destination peer, you configure a reference to the name of the IKE policy to be used for it.

You refer to the policy to be used in the ike-policy statement of the peer gateway configuration. IKE policies are configured under the IKE hierarchy, not the gateway hierarchy, because there can be more than one of them.

To specify the IKE policy to use for the destination peer gateway (los-angeles-gw3), enter the following command in configuration mode:

```
user@host# set security ike gateway los-angeles-gw3 policy pol1
```

Table 44 defines the statement parameters.

Table 44: IKE Phase 1 Policy Reference Configuration Statement

Destination Peer Policy Reference	Syntax
Refers to the IKE policy to be used for communication with the destination peer. In turn, the policy refers to the proposal containing cryptography parameters to present to the peer for Phase 1 negotiation.	`set security ike gateway <destination-peer-gateway-name> ike-policy <ike-policy-name>` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `ike-policy-name`: Name of the policy specified as an alphanumeric string.

Configuring the IKE Phase 1 External Interface

```
security—ike—gateway
                    <destination-peer-gateway-name>
                    external-interface <interface-name>
```

You must specify the external interface to be used for the IPSec tunnel.

To specify the external interface to be used to communicate with the Los Angeles gateway (los-angeles-gw3), enter the following statement in configuration mode:

```
user@host# set security ike gateway los-angeles-gw3 external-interface
ge-0/0/3.1
```

Table 45 defines the statement parameters.

Table 45: External Interface Name Configuration Statement

External Interface Name	Syntax
Outgoing interface to be used for the IPSec VPN.	`set security ike gateway <destination-peer-gateway-name> external-interface <interface-name>` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `interface-name`: Name of the outgoing interface to be used for the VPN.

Configuring IKE Phase 1 Policies

An IKE policy specifies the authentication method to be used between the peers to identify themselves to each other during their Phase 1 exchanges (either certificate or preshared keys). It also refers to one or more proposals to be presented in succession to the peer for negotiation. Proposals contain parameters whose values are used to secure the Phase 2 tunnel.

The authentication method can be either preshared keys or certificates. Use of certificates is more scalable than preshared keys, but certificates are more complicated to implement.

Configuring the IKE Policy Name

```
security ——> ike -->
            policy <policy-name>
```

An IKE policy has a name, or label, which is referred to from within the ike-policy statement of the destination peer gateway configuration.

You can specify the IKE policy name as part the statements that belong to it. However, this section shows how to configure it alone to distinguish it from the policy's parameters:

```
user@host# set security gateway los-angeles-gw3 ike-policy pol1
```

Table 46 defines the policy name parameter.

Table 46: IKE Policy Name Configuration Statement

IKE Policy Name	Syntax
Name of an IKE policy. A peer gateway configuration refers to the policy to be used for it in its ike-policy statement.	`set security ike policy <policy-name>` ■ `policy-name`: An alphanumeric string identifying the policy.

Configuring IKE Policy Authentication Parameters

```
security —-> ike —->
            policy <policy-name>
                (pre-shared-key ascii-text <key> | certificate
                                                    <certificate-information>)
```

During Phase 1 of the IKE negotiations, the peers authenticate themselves to each other. For this purpose they use either preshared keys or certificates. You configure the type of authentication and its parameters to be used in the IKE policy.

NOTE: In the proposal referred to from within the policy you specify the authentication method. For example, if you specify certificate and its parameters from within the policy, the proposal that the policy refers to can specify the type of certificate. If you specify the preshared key from within the policy, the proposal states that preshared keys are used.

To specify the type of authentication to be used for the local peer, enter the following statement in configuration mode:

```
user@host# set security ike policy pol1 pre-shared-key ascii-text 1234pskey5678
```

Table 47 defines the authentication method statement parameters.

Table 47: IKE Policy Authentication Method Configuration Statement

IKE Policy Authentication Method	Syntax
Method to be used for the peers to authenticate themselves to each other during Phase 1 negotiations.	`set security ike policy <policy-name> certificate local-certificate <certificate ID> trusted-ca (<CA ID> | use-all) peer-cert-type (pkcs7 | x509-signature) | preshared key` ■ `certificate ID`: The local certificate to use. If no value is specified, the JUNOS software with enhanced services chooses from among all configured certificates. ■ `CA ID`: The preferred certificate authority for the destination peer to for its certificate. The use-all keyword instructs the device to use all configured certificate authorities. Otherwise, you must specify a configured certificate authority ID. If no value is specified, then no certificate request is sent (although incoming certificates are still accepted). ■ `peer-cert-type`: The preferred type of certificate to request from the peer: either PKCS7 or X509-Signature. (default: X509-Signature).

Configuring a Reference to the IKE Proposal

```
security ——> ike —->
            policy <policy-name>
                        proposals <proposal-name>
```

The IKE policy includes a reference to one or more IKE proposals whose parameters are to be negotiated with the destination peer. You can configure proposals or refer to one of the predefined proposal sets that are identified in Table 48 on page 113.

The agreed upon parameters, resulting from the negotiation, are used to secure a tunnel with the destination peer. The proposals are presented to the destination peer in the order in which they are specified until the destination peer accepts one of them. If you intend to configure and specify more than one proposal for negotiation, it is a good idea to specify first the most restrictive policy, that is, the one that offers the greatest security. For the negotiation to work, the destination peer must also have configured and referenced at least one of the proposals presented by the local gateway.

NOTE: The tunnel, which is the outcome of Phase 1, is used by the peers for Phase 2 negotiations. It is not used for transit traffic. The outcome of Phase 2 determines how data traffic is to be secured.

To refer to the IKE proposal whose parameters are to be sent to the Los Angeles peer for negotiation, enter the following statement in configuration mode:

```
user@host# set security ike policy pol1 proposals prop1
```

Table 48 defines the proposal reference statement parameters.

Table 48: IKE Phase 1 Proposal Configuration Statement

IKE Phase 1 Proposal Reference	Syntax
A reference to one or more IKE Phase 1 proposals whose parameters are to be negotiated with the destination peer. The IKE policy for the peer gateway refers to the proposals to be used by name.	`set ike policy <policy-name> proposals [<proposal1> <proposal2>...] proposal-set <proposal-set-name>` ■ `policy-name`: ASCII text string. ■ `proposal1, proposal2`: A reference to the name(s) of one or more proposals to be used for negotiation with the peer. ■ `proposal-set-name`: A reference to a predefined proposal set to be used for negotiation with the peer.
Predefined Phase 1 proposals sets (for details, see the *JUNOS Software CLI Reference*. ■ Basic: ■ Proposal 1: preshared key, g1, des, sha1 ■ Proposal 2: preshared key, g1, des, md5 ■ Compatible: ■ Proposal 1: preshared key, g2, 3des, sha1 ■ Proposal 2: preshared key, g2, 3des, md5 ■ Proposal 3: preshared key, g2, des, sha1 ■ Proposal 4: preshared key, g2, des, md5 ■ Standard: ■ Proposal 1: preshared key, g2, 3des, sha1 ■ Proposal 2: preshared key, g2, aes128, sha1	

Configuring IKE Phase 1 Proposals

During Phase 1 of the IKE negotiations, the local peer proposes to the destination peer the cryptography to be used to secure the IPSec Phase 2 tunnel. The Phase 2 tunnel is then used to negotiate values to be used to secure transit traffic across the IPSec VPN.

The result of Phase 1 negotiations determines the security of the tunnel to be used for Phase 2. The outcome of Phase 2 is the method to be used to secure data that traverses the IPSec VPN.

Phase 1 proposals include parameters, such as the authentication and encryption algorithms to be used. The local peer presents proposals to its destination peer one after another, beginning with the first one specified. For the negotiation to work, the destination peer must also have configured and referenced at least one of the proposals presented by the local gateway.

Configuring the IKE Proposal Name

```
security —> ike —> proposal <proposal-name>
```

You can configure one or more proposals within a Phase 1 configuration. You provide a name, or label, for an IKE proposal, to identify it.

You specify the name for a proposal when you configure any of its parts. You can configure the entire proposal at one time, as shown in the following statement:

```
user@host# set security ike proposal prop1 authentication-method pre-shared-keys
dh-group group2 encryption-algorithm aes-256-cbc authentication-algorithm md5
```

The following sections explain each part.

Table 49 gives the proposal name syntax.

Table 49: IKE Proposal Name Configuration Statement

IKE Proposal Name	Syntax
The name, or label, of the IKE Phase 1 proposal that follows. The gateway policy statement refers to one or more proposals to be used.	`set security ike proposal <proposal-name>` ■ `proposal-name`: Name of the proposal to be used to identify it.

Configuring the IKE Proposal Authentication Method

```
security —> ike —> proposal <proposal-name>
                    authentication-method
                       (pre-shared-keys | RSA signatures)
```

Within a proposal you can specify the authentication method to be used. If, within the policy, you specified that certificates are to be used, in the proposal you can specify the type of certificate to use.

For example, if you specified the certificate parameter and its values within the policy statement, the policy must refer to a proposal that specifies the type of certificate to use, such as RSA.

If you specified a preshared key within the policy, the proposal must refer to a proposal that specifies pre-shared-key as the authentication method to be used.

NOTE: The proposal option pre-shared-keys is different from the policy statement pre-shared-key in that it is plural. You configure a specific key from within the policy. From within the proposal, you indicate that you want to use pre-shared keys.

To configure the authentication method for prop1 for IKE Phase 1, enter the following statement in configuration mode:

```
user@host# set security ike proposal prop1 authentication-method pre-shared-keys
```

Table 50 defines the IKE proposal authentication method parameter.

Table 50: IKE Proposal Authentication Method Configuration Statement

IKE Proposal Authentication Method	Syntax
The kind of authentication method to use for the peers to identify themselves to each other. ■ For a preshared key, the *policy* gives the specific key. The *proposal* states that pre-shared keys are to be used. ■ For a certificate. the policy certificate parameter identifies which certificate is to be used. The proposal tells the kind of certificate to use.	`set security ike proposal <proposal-name> authentication-method <method>` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(authentication-method) method`: Specifies the authentication method, which corresponds to the policy specification. Specify one of the following options: ■ `pre-shared-keys` ■ `rsa-signatures`

Configuring the IKE DH Group

```
security ——> ike ——> proposal <proposal-name>
                     dh-group (group1 | group2 | group5)
```

A Diffie-Hellman group performs key exchange to generate keying material for per-packet encryption and authentication.

To configure the Diffie-Hellman group to be specified as part of the proposal, enter the following statement in configuration mode:

```
user@host# set security ike proposal prop1 dh-group group2
```

Table 51 gives an overview of use of the Diffie-Hellman group parameter and its syntax.

Table 51: IKE DH Group Configuration Statement

IKE Proposal DH Group	Syntax
Diffie-Hellman group to use to increase the security of the tunnel. A Diffie-Hellman exchange allows the peers to produce a shared secret over an unsecured medium without passing the secret through the network. The size of the prime modulus used in each group differs and determines the strength of its security.	`set ike proposal <proposal-name> dh-group (group1 | group2 | group5)` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(dh-group)`: Specify one of the following keywords to indicate the Diffie-Hellman group to be used. ■ `group1` (768-bit modulus) ■ `group2` (1024-bit modulus) ■ `group5` (1536-bit modulus) Because the modulus for each Diffie-Hellman group is a different size, the peers must agree to use the same group.

Configuring the IKE Proposal Encryption Algorithm

```
security ——> ike ——> proposal <proposal-name>
                encryption-algorithm (des-cbc |3des-cbc | aes-128-cbc |
                                                aes-192-cbc| aes-256-cbc)
```

During IKE Phase 1, the peers negotiate the cryptographic parameters to be used for their IPSec IKE Phase 2 exchanges. These parameters are specified as part of a proposal.

The encryption algorithm agreed on by the two peers is used to encrypt parameters that they exchange and negotiate in IKE Phase 2.

A proposal can contain only one encryption algorithm, but because a policy can refer to more than one proposal, several different encryption methods can be proposed to the peer.

To configure the encryption algorithm, enter the following command in configuration mode:

```
user@host# set security ike proposal prop1 encryption-algorithm aes-256-cbc
```

Table 52 defines the encryption algorithm configuration.

Table 52: IKE Encryption Algorithm Configuration Statement

IKE Proposal Encryption Algorithm	Syntax
Encryption algorithm specified within a proposal to be presented to the peer for negotiation. The agreed-upon algorithm is used to secure exchanges between the peers during IKE Phase 2.	`set ike proposal <proposal-name> encryption-algorithm (des-cbc | 3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc)` ■ `proposal-name`: Name of the proposal, which is referred to in the policy statement of peer gateways to specify use of it. ■ `(encryption-algorithm)`: Specify one of the five encryption algorithms to use in the proposal. In order from strongest to weakest, the options are ▪ `aes-256-cbc` ▪ `aes-192-cbc` ▪ `aes-128-cbc` ▪ `3des-cbc` ▪ `des-cbc`

Configuring the IKE Authentication Algorithm

```
security ——> ike —-> proposal <proposal-name>
                        authentication-algorithm (md5 | sha1)
```

During IKE Phase 1, the peers negotiate the cryptographic parameters to be used for their IKE IPSec Phase 2 exchanges. These parameters are specified as part of a proposal.

The authentication algorithm agreed on by the two peers is used to encrypt messages that they exchange and negotiate in IKE Phase 2 over the tunnel create during Phase 1 (IKE Phase 1 SA).

A proposal can contain one authentication algorithm, but because a policy can refer to more than one proposal, several different authentication methods can be proposed to the peer.

To configure the authentication algorithm for the IKE Phase 1 proposal, enter the following statement in configuration mode:

```
user@host# set security ike proposal prop1 authentication-algorithm md5
```

Table 53 defines the IKE proposal authentication algorithm.

Table 53: IKE Proposal Authentication Algorithm Configuration Statement

IKE Proposal Authentication Algorithm	Syntax
Authentication algorithm to propose to the peer to be used to authenticate exchanges between the peers during IKE Phase 2. A proposal contains the authentication algorithm, and the peer must accept the entire proposal.	`set ike proposal <proposal-name> authentication-algorithm (md5 | sha1)` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(authentication-algorithm)`: Specify one of the two authentication algorithms to use in the proposal. Options are: `sha1` and `md5`.

Configuring the IPSec VPN

After the peer gateway has been defined and the IKE Phase 1 proposals have been configured, you configure the IPSec VPN gateway. The IPSec VPN gateway configuration brings together the peer gateway configuration, the IPSec proposal, and the tunnel interface to be used for the IPSec VPN.

Naming the IPSec VPN Gateway

```
security —-> ipsec --> vpn
             vpn <vpn-name>
```

In defining an IPSec VPN, you first give it a name. This is the name of the VPN that you are creating overall. It differs from the name for the peer gateway, which identifies the configuration parameters for the peer.

To configure the name for the IPSec VPN, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn albuquerque-vpn1
```

Table 54 defines the IPSec VPN name syntax.

Table 54: IPSec VPN Name Configuration Statement

IPSec VPN name	Syntax
Name of the IPSec VPN being configured for the local gateway. The IPSec VPN gateway name is specified outside the IKE hierarchy. This is because it identifies the entire IPSec VPN, not just the peer gateway.	`set security ipsec vpn <vpn-name>` ■ `vpn-name`: Name of the IPSec VPN.

Identifying the Peer Gateway

```
security —-> ipsec --> vpn
        vpn <vpn-name>
            ike
                gateway <destination-peer-gateway-name>
```

The destination peer gateway name introduces the configuration for it. This name must be the same as the one specified for the gateway configuration for IKE Phase 1. For Phase 2, you configure a peer gateway under the Phase 2 IKE hierarchy.

To specify the name of the destination peer (`los-angeles-gw3`), enter the following statement in configuration mode:

```
user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3
```

Table 55 defines the peer gateway name parameter for Phase 2.

Table 55: IPSec Phase 2 Destination Peer Gateway Name Configuration Statement

IPSec Phase 2 Destination Peer Gateway Name	Syntax
Name of the destination peer whose gateway you are configuring.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name>` ■ `vpn-name`: IPSec VPN name. ■ `destination-peer-gateway-name`: Name of the destination peer whose gateway information you also configured during IKE Phase 1.

Configuring the Reference to the IPSec Phase 2 Policy

```
security --> ipsec--> vpn <vpn-name>
                ike
                    ipsec-policy <ipsec-policy-name>
```

As you did for Phase 1, for Phase 2 you specify a reference to the policy to be used for the VPN. You refer to the IPSec policy in the ipsec-policy parameter of the IKE configuration.

To refer to the IPSec policy to be used, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn albuquerque-vpn1 ike ipsec-policy ipsec-pol1
```

Table 56 gives the syntax for Phase 2 policy reference parameter.

Table 56: IPSec Policy Reference Configuration Statement

Reference to the IPSec Policy to Use	Syntax
Refers to the Phase 2 policy to be used. In turn, the policy refers to one or more proposals to be negotiated with the destination peer.	`set security ipsec vpn <vpn-name> ike ipsec-policy <ipsec-policy-name>` ■ `vpn-name`: A string identifying the IPSec VPN itself. ■ `ipsec-policy-name`: Name of the policy to be used for the peer gateway.

Binding the IPSec VPN to a Tunnel Interface

```
security —-> ipsec --> vpn
            bind-interface <interface-name>
```

A tunnel interface serves as an opening to a tunnel. Traffic enters and exits a VPN tunnel through it. For route-based VPNs, you bind an interface to be used for the VPN to a tunnel interface. You can reference the tunnel interface in a route and then reference that destination in one or more policies.

You bind a route-based VPN tunnel to a tunnel interface so that the router can route traffic to and from it. You can bind a route-based VPN tunnel to a tunnel interface (with or without an IP address/netmask).

To bind the tunnel interface to the IPSec VPN, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn albuquerque-vpn1 bind-interface st0.1
```

Table 57 gives the syntax for this parameter.

Table 57: IPSec Bind Interface Configuration Statement

Reference to the IPSec Policy to Use	Syntax
Binds the tunnel interface to the IPSec VPN.	`set security ipsec vpn <vpn-name> bind-interface <interface-name>` ■ `vpn-name`: A string identifying the IPSec VPN itself. ■ `interface-name`: The tunnel interface to be used for the VPN.

Configuring IPSec Phase 2 Proposals

During IPSec Phase 2, the peers negotiate cryptographic parameters that determine how data traffic that traverses the IPSec VPN is to be secured. These parameters identify the kind of authentication and encryption to be used. They are specified in a proposal to which the gateway policy refers.

You configure IPSec proposals under the IPSec hierarchy, not under the peer gateway configuration.

Configuring the Phase 2 IPSec Proposal Name

```
security --> ipsec
            proposal <proposal-name>
```

The proposal name, or label, identifies a proposal so that it can be referred to from within a policy. You can specify a proposal name when you configure its parameters. For clarity, this section shows how to specify the name alone.

To configure the name for the IPSec VPN proposal, enter the following statement in configuration mode:

```
user@host# set security ipsec proposal ipsec-prop1
```

Table 58 gives the syntax for the IPSec proposal name specification.

Table 58: IPSec Proposal Name Configuration Statement

IPSec Proposal Name	Syntax
This name is referred to in an IPSec policy to specify which proposal is to be used.	`set security ipsec proposal <proposal-name>` ■ `proposal-name`: Name of the IPSec proposal whose protocol you are configuring

Configuring the Phase 2 IPSec Protocol

```
security --> ipsec
            proposal <proposal-name>
                protocol (esp | ah)
```

Data traffic traversing IPSec VPNs is protected by either of two protocols that treat the data as a payload by encapsulating it and encrypting and authenticating it.The IPSec protocols are Encapsulating Security Payload (ESP) and Authentication Header (AH).

To configure the IPSec protocol for the VPN, enter the following statement in configuration mode:

```
user@host# set security ipsec proposal ipsec-prop1 protocol esp
```

Table 59 gives an overview of the IPSec protocol syntax.

Table 59: IPSec Protocol Definition Configuration Statement

IPSec Protocols	Syntax

One of two IPSec protocols to be used to authenticate and/or encrypt data traffic traversing the IPSec VPN.	`set security ipsec proposal <proposal-name> protocol (esp | ah)`
	■ `proposal-name`: Name of the IPSec proposal whose protocol you are configuring
	■ `(protocol)`: Either of the following protocols:
	  ■ `esp`: Encapsulating Security Payload
	  ■ `ah`: Authentication Header

IPSec Protocols	Syntax
■ ESP provides the following services: ■ Encrypts data, rendering it undecipherable. ■ Provides data origin authentication ■ Provides data integrity ■ Supports antireplay protection ■ What ESP does not provide: ■ It authenticates and encrypts only the data portion of the packet, excluding the IP header.	
■ AH provides the following services: ■ Provides data origin authentication ■ Supports antireplay protection for most of the packet ■ What AH does not provide: ■ Antireplay protection of the fields in the IP header, which can change during transit. ■ Encryption ■ Confidentiality	

Configuring the IPSec Proposal Authentication Algorithm

```
security --> ipsec
             proposal <proposal-name>
                     authentication algorithm (hmac-md5-96 | hmac-sha1-96)
```

Each of the IPSec protocols—ESP and AH—uses an algorithm to verify the authenticity and integrity of the content of the packet and its origin. You specify the algorithm to use as part of a proposal to be presented to the destination peer for negotiation.

A packet is authenticated through use of a checksum calculated via a hash-based message authentication code (HMAC) such as MD5 or SHA-1.

For ipsec-prop1, MD5 is used.

To configure the authentication algorithm for the proposal, enter the following statement in configuration mode.

```
user@host# set security ipsec proposal ipsec-prop1 hmac-md5-96
```

Table 60 defines the authentication algorithm.

Table 60: IPSec Proposal Authentication Algorithm Configuration Statement

IPSec proposal authentication algorithm	Syntax

Authentication algorithm to be included in the specified proposal.	`set security ipsec proposal <proposal-name> authentication-algorithm (hmac-md5-95 | hmac-sha1-96)`
	■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal whose authentication algorithm you are configuring.
	■ `(authentication-algorithm)`: One of the following options:
	  ■ `hmac-md5-95`
	  ■ `hmac-sha1-96`

Configuring the IPSec Proposal Encryption Algorithm for ESP

```
security --> ipsec
            proposal <proposal-name>
                    encryption algorithm (3des-cbc| aes-128-cbc | aes-192-cbc |
                                                aes-256-cbc | des-cbc)
```

The IPSec ESP protocol provides for encryption of data traffic as well as authentication. As part of a proposal that specifies ESP as its IPSec protocol, you specify the encryption algorithm.

To configure the encryption algorithm to be used for ESP in ipsec-prop1, enter the following statement in configuration mode.

```
user@host# set security ipsec proposal ipsec-prop1 encryption-algorithm
aes-256-cbc
```

Table 61 defines the encryption algorithm.

Table 61: IPSec Proposal Encryption Algorithm Configuration Statement

IPSec Proposal Encryption Algorithm	Syntax

Identifies the encryption algorithm to be used for encryption via ESP.	`set security ipsec proposal <proposal-name> encryption-algorithm (3des-cbc| aes-128-cbc | aes-192-cbc |aes-256-cbc | des-cbc)`
	■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal whose encryption algorithm you are configuring.
	■ `(encryption-algorithm)`: Specified as a keyword:
	  ■ `3des-cbc`
	  ■ `aes-128-cbc`
	  ■ `aes-192-cbc`
	  ■ `aes-256-cbc`
	  ■ `des-cbc`

Configuring IPSec Phase 2 Policies

Similar to IKE Phase 1, for Phase 2 an IPSec VPN policy statement contains a reference to the proposal to be presented to the destination peer for negotiation. (The proposal specifies the parameters to be used to secure data traversing the IPSec VPN.)

Configuring the IPSec Policy Name

```
security ipsec
            policy <policy-name>
```

Each Phase 2 IPSec policy has a name, or label, which is referred to from with the policy statement of the destination peer gateway configuration to identify the policy to be used for it.

You can specify the IPSec policy name as part of the statements that compose it. However, this section shows how to configure it alone for clarity:

```
user@host# set security ipsec policy ipsec-pol1
```

Table 62 defines the policy name parameter.

Table 62: IPSec Policy Name Configuration Statement

IPSec Policy Name	Syntax
Name of an IPSec policy. A peer gateway configuration refers to the policy to be used for it in its policy statement.	`set security ipsec policy <policy-name>` ■ `policy-name`: An alphanumeric string identifying the policy.

Configuring a Reference to the IPSec Proposals

```
security ——> ike ——>
            policy <policy-name>
                        proposals <proposal-name>
```

An IPSec policy includes a reference to one or more IPSec proposals whose parameters are negotiated with the destination peer.

The agreed-upon parameters, resulting from the negotiation, are used to secure data traffic that transits the VPN. The local peer presents the proposals in order to the destination peer until the destination peer accepts one of them.

For the policy to refer to the IPSec proposal whose parameters are to be sent to the Los Angeles peer for negotiation, enter the following statement in configuration mode:

```
user@host# set security ipsec policy ipsec-pol1 proposal ipsec-prop1
```

Table 63 defines the IPSec proposal reference syntax.

Table 63: IPSec Phase 2 Proposal Reference Configuration Statement

IKE Phase 2 Proposal Reference	Syntax
A reference to one or more IPSec Phase 2 proposals or proposal sets to be used to negotiate security for the tunnel.	`set security ipsec policy <policy-name> proposal [<proposal1> <proposal2> ...] proposal-set <proposal-set-name>` ■ `policy-name`: ASCII text string. ■ `proposal1, proposal2`: A reference to the name or names of one or more proposals to be used for negotiation with the peer. ■ `proposal-set-name`: A reference to a predefined proposal set to be used for negotiation with the peer.
Predefined Phase 2 proposals sets (For details, see the *JUNOS Software CLI Reference*. ■ Basic: ■ Proposal 1: esp-no pfs, des, sha1 ■ Proposal 2: esp-no pfs, des md5 ■ Compatible: ■ Proposal 1: esp-no pfs, 3des, sha1 ■ Proposal 2: esp-no pfs, 3des, md5 ■ Proposal 3: esp-no pfs, des, sha1 ■ Proposal 4: no pfs, des, md5 ■ Standard: ■ Proposal 1: esp, g2 (pfs), 3des, sha1 ■ Proposal 2: esp, g2, aes128, sha1	

Tasks for Implementing an Advanced a Route-based IPSec VPN

This section lists the tasks entailed in creating a route-based IPSec VPN with additional features beyond minimum requirements.

To create a route-based IPSec VPN, with advanced features, you complete the tasks shown in Table 64.

Table 64: Procedure for Creating an Advanced Route-based IPSec VPN

Task	Instructions
1. Configure the required security zones and interfaces.	See “Configuring Interfaces for the VPNs” on page 105. See “Designing Firewalls for Branch Offices” on page 11.
2. Configure addresses for inbound and outbound policies to be used in proxy IDs.	See “Designing Firewalls for Branch Offices” on page 11.
Configure Phase 1 of the IPSec VPN	
3. Specify that the VPN should check for and respond to an invalid Security Parameter Index (SPI) in peer packet headers.	See “Configuring the VPN to Respond to Invalid Security Parameter Indexes” on page 129.
4. Configure the peer gateway (IKE Phase 1) and destination peer information.	See “Configuring the Peer Gateway and Destination Peer Information” on page 109.
5. Configure the local external interface.	See “Configuring the IKE Phase 1 External Interface” on page 110.
6. Configure an IKE Phase 1 policy.	See “Configuring IKE Phase 1 Policies” on page 110.
7. Configure one or more IKE Phase 1 proposals.	For this VPN, two IKE Phase 1 proposals are configured. See “Configuring IKE Phase 1 Proposals” on page 113 for a description that is also applicable when two or more proposals are configured.
Configure Phase 2 of the IPSec VPN	
8. Configure the global VPN monitoring statement.	See “Enabling VPN Monitoring” on page 130.
9. Configure the IPSec VPN name.	See “Configuring the IPSec VPN” on page 117.
10.Configure the peer gateway (IPSec Phase 2).	See “Identifying the Peer Gateway” on page 118.
11.Disable the antireplay feature for testing purposes.	See “Disabling Antireplay Checking” on page 131.
12.Configure the gateway policy reference.	See “Configuring the Reference to the IPSec Phase 2 Policy” on page 118
13.Configure the proxy identity.	See “Configuring the Proxy Identity” on page 132.
14.Bind the IPSec VPN to the tunnel interface.	See “Binding the IPSec VPN to a Tunnel Interface” on page 119.
15.Configure the VPN monitoring options.	See “Configuring VPN Monitoring Options” on page 133.
16.Specify when tunnels should be set up.	See “Determining When to Establish Tunnels” on page 134.

Task	Instructions
17.Configure one or more IPSec (IKE Phase 2) proposals.	See “Configuring IPSec Phase 2 Proposals” on page 119.
18.Configure an IPSec Phase 2 policy.	See “Configuring IPSec Phase 2 Policies” on page 123.
19.Configure firewall policies.	See “Implementing Firewall Deployments for Branch Offices” on page 37.

The configuration includes the following additional features:

- Invalid SPI check
- Two IKE Phase 1 proposals
- IPSec Phase 2 global VPN monitoring
- Antireplay checking disabled (Antireplay checking is enabled by default.)
- Proxy identity configuration

Configuring an Advanced Route-based IPSec VPN

This section builds on the basic IPSec VPN to show how to configure additional IPSec VPN features. Because this VPN is an enhancement of the basic one, steps and features used for the basic configuration are identified but not described.

CAUTION: To use this sample configuration, you must commit the configuration after all of the statements are entered. If you intend to commit the statements individually, you must reverse the order of statements within each IKE phase because of dependencies among the statements.

IPSec VPN Statements Summary

This section shows the CLI statements used to configure additional features for the IPSec VPN. The entire configuration is provided, with the additional statements shown in bold.

The complete configuration is shown after the set of CLI statements.

IKE Phase 1: Advanced Route-based IPSec VPN Statements

Responding to invalid SPIs. (See “Configuring the VPN to Respond to Invalid Security Parameter Indexes” on page 129.)

```
user@host# set security ike respond-bad-spi 10
```

Destination peer information. (See "Configuring the Peer Gateway Name and Destination Peer Address" on page 109.)	user@host# set security ike gateway los-angeles-gw3 address 5.2.3.4
Phase 1 IKE policy reference. (See "Configuring the IKE Phase 1 Policy Reference" on page 109.)	user@host# set security ike gateway los-angeles-gw3 ike-policy pol1
Phase 1 gateway external interface. (See "Configuring the IKE Phase 1 External Interface" on page 110.)	user@host# set security ike gateway los-angeles-gw3 external-interface ge-0/0/3.2
Phase 1 IKE policy. (See "Configuring IKE Phase 1 Policies" on page 110.)	user@host# set security ike policy pol1 pre-shared-key ascii-text 1234pskey5678 user@host# set security ike policy pol1 proposals prop1
Phase 1 IKE proposal. (See "Configuring IKE Phase 1 Proposals" on page 113.)	user@host# set security ike proposal prop1 authentication-method authentication-method pre-shared-keys dh-group group2 encryption-algorithm aes-256-cbc authentication-algorithm md5 user@host set security ike proposal prop2 authentication-method dsa-signature dh-group group5 encryption-algorithm aes-192-cbc authentication-algorithm sha1

Configuring the Phase 2 IPSec VPN

IPSec VPN monitoring global statement. (See "Enabling VPN Monitoring" on page 130.)	user@host# set security ipsec vpn-monitor-options interval 10 threshold 5
IPSec VPN name. (See "Naming the IPSec VPN Gateway" on page 117.)	user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3
IPSec Phase 2 Peer gateway name. (See "Identifying the Peer Gateway" on page 118.)	user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3
Antireplay disabled. (See "Disabling Antireplay Checking" on page 131.)	user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3 no-anti-replay

IPSec Phase 2 policy reference. (See "Configuring the Reference to the IPSec Phase 2 Policy" on page 118.)	user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3 ipsec-policy ipsec-pol1
Binding the IPSec VPN to a tunnel interface. (See "Binding the IPSec VPN to a Tunnel Interface" on page 119.)	user@host# set security ipsec vpn albuquerque-vpn1 bind-interface st0.1
VPN monitoring options. (See "Configuring VPN Monitoring Options" on page 133.)	user@host# set security ipsec vpn albuquerque-vpn1 vpn-monitor optimized source-interface ge-0/0/3.2 destination-ip 5.2.3.4
Tunnel establishment. (See "Determining When to Establish Tunnels" on page 134.)	user@host# set security ipsec vpn albuquerque-vpn1 establish-tunnels immediately
IPSec Phase 2 proposal. (See "Configuring IPSec Phase 2 Proposals" on page 119.)	user@host# set security ipsec proposal ipsec-prop1 protocol esp authentication-algorithm hmac-md5-96 encryption-algorithm aes-256-cbc
IPSec Phase 2 policies. (See "Configuring IPSec Phase 2 Policies" on page 123.)	user@host# set security ipsec policy ipsec-pol1 proposal ipsec-prop1

Advanced Route-based IPSec VPN Configuration with Additional Features

The following example shows the same configuration presented as a config excerpt.

```
security {
     ike {
            respond-bad-spi 10
            gateway los-angeles-gw3 {
                address 5.2.3.4;
                ike-policy pol1
                external-interface ge-0/0/3.2
            }
            policy pol1 {
                    preshared-key
                    ascii-text 1234pskey5678
            }
            proposal prop1 {
                authentication-method pre-shared-keys;
                dh-group group2;
                encryption-algorithm aes-256-cbc;
                authentication-algorithm md5;
            }
```

```
            proposal prop2 {
                authentication-method pre-shared keys;
                dh-group group5;
                encryption-algorithm aes-192-cbc;
                authentication-algorithm sha1;
            }

        }
        ipsec {
            vpn-monitor-options {
                interval 10;
                threshold 5;
            }
            vpn albuquerque-vpn1 {
                ike {
                    gateway los-angeles-gw3;
                    no-anti-replay-check;
                    ipsec-policy route-ipsec-pol1;
                    proxy-identity {
                        local 1.3.4.7/24;
                        remote 5.2.3.4/24;
                    }
                }
                bind-interface st0.0
                vpn-monitor {
                    optimized;
                    source-interface ge-0/0/3.2
                    destination 5.2.3.4
                establish-tunnels immediately
                proposal ipsec-prop1 {
                    protocol esp;
                    authentication-algorithm hmac-md5-96;
                    encryption-algorithm aes-256-cbc;
                }
                policy ipsec-pol1 {
                    proposal ipsec-prop1;
                }
            }
        }
```

Configuring the VPN to Respond to Invalid Security Parameter Indexes

```
security ——> ike
            respond-bad-spi
```

The security associations (SAs) between two IPSec VPN peers can become unsynchronized for some reason. You can configure the invalid SPI response feature to recognize and respond to this condition. When the feature is configured, the local peer can recognize an invalid or unknown SPI in a packet header from its destination peer. The JUNOS software with enhanced services responds to the condition by resetting the state of the destination peer so that the two peers are synchronized again.

To configure the IPSec VPN for the Albuquerque branch to respond to invalid SPIs from its Los Angeles peer 10 times before engaging the peer in a new IKE session, enter the following statement in configuration mode:

```
user@host# set security ike respond-bad-spi 10
```

Table 65 describes the Invalid SPI Response statement parameters.

Table 65: Invalid SPI Response Configuration Statement

SPI	Syntax
Invalid SPI Response Enables one peer to recognize and handle an invalid IPSec security parameter index (SPI) value in the packet header of the other peer and synchronize the state of the two peers again by initiating a new IKE session. (The IKE session generates new IPSec SA keying material.)	set security ike respond-bad-spi <number> ■ number: Number of times to respond to an invalid SPI from the destination peer before engaging it in a new IKE session.

Enabling VPN Monitoring

```
security —-> ike —-> ipsec-vpn
            vpn-monitor-options
                interval <seconds>
                threshold <number>
```

If a destination peer's IPSec SA is deleted and the local peer's IPSec SA is not, the local peer will continue to send traffic, incurring associated encryption and authentication costs, until it expires or discovers that its peer IPSec SA no longer exists.

The VPN monitoring feature provides a way for a peer to monitor the state of its destination peer to determine if it exists. If the VPN monitoring feature is configured, when the router detects that the destination peer is no longer available, it deletes its own IPSec SA, returning the two peers to a symmetrical state and stopping transmission of its own IPSec traffic. If the establish-tunnels immediately option is enabled, the router renegotiates the IKE and IPSec SAs after it deletes its own set.

At the global level, you can configure the interval at which echo messages are sent to the peer. You can also configure the threshold for the number of messages sent after which the router determines if the peer is available.

VPN monitoring operates at a finer granularity than Dead Peer Detection (DPD) because it checks individual SAs.

Enter the following statement in configuration mode to configure the Albuquerque branch VPN to send ICMP echo requests to its peer at a 10-second interval for five consecutive times without a response before engaging the peer in a new IKE session:

```
user@host# set security ipsec vpn-monitor-options interval 10 threshold 5
```

Table 66 describes the VPN monitoring options statement parameters.

Table 66: VPN Monitoring Global Options Configuration Statement

VPN monitoring	Syntax
Set VPN monitoring options to specify when to send ICMP echo requests to the peer and when to determine if the peer is available.	`set security ipsec vpn-monitoring-options interval <seconds> threshold <number>` ■ `(interval) seconds`: Interval expressed in seconds for when the router should send echo requests to the destination peer to monitor its liveliness and determine if it is reachable. ■ `(threshold) number`: Number of consecutive unsuccessful ICMP echo requests (pings) to send to the peer before it is declared unavailable.

Disabling Antireplay Checking

```
security --> ipsec --> vpn <vpn-name>
                ike
                 gateway <destination-peer-gateway-name>
                    no-anti-replay
```

Antireplay checking makes it impossible for an intruder to replay a packet. Attackers can intercept packets and insert altered ones into the data stream. They can use intercepted packets later to flood the system and cause a denial-of-service (DoS) attack, and they can use them to gain entry to a trusted system. Antireplay checking detects packets that match the sequence numbers of those that have already been received, and it discards any replayed packets.

By default, the antireplay feature is enabled. It is a good idea to leave the feature enabled when you are running the system online. However, to avoid processing incurred by use of the feature, you might want to disable it when you are testing or debugging the system. Also, you might want to disable antireplay checking to resolve compatibility issues with third-party peers. However, in such cases, you weaken the security of your system.

For the Los Angeles destination peer gateway configuration, antireplay checking remains enabled. However, if you wanted to disable it for the branch, you would enter the following statement in configuration mode:

```
user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3
no-anti-replay
```

Table 67 gives an overview of this parameter and its syntax.

Table 67: Antireplay Checking Configuration Statement

Antireplay Checking	Syntax
Disables or enables the feature that detects packets matching the sequence numbers of those that have already been received. Discards replayed packets.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name> no anti-replay` ■ `vpn-name`: Name of the IPSec VPN created for the local gateway. ■ `destination-peer-gateway-name`: Name of the destination peer. This is the same name for the destination gateway as specified in the IKE Phase 1 configuration. ■ `no anti-replay`: Disables antireplay checking ■ `set anti-replay`: Enables antireplay checking

Configuring the Proxy Identity

```
security ipsec vpn <vpn-name>
                ike
                 gateway <destination-peer-gateway-name>
                    proxy-identity
                        local <ipv4-prefix>
                        remote <ipv4-prefix>
                        service <service-name>
```

A proxy identity is a kind of agreement between IKE peers to permit traffic through a tunnel if the traffic matches the specified match criteria of local address, remote address, and service. By default, the router uses the identities and service information specified in the firewall policy if you do not configure them.

CAUTION: The proxy id for both peers must match. 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.

Enter the following statement in configuration mode to configure the proxy ID for the Albuquerque branch VPN:

```
user@host# set security ipsec vpn albuquerque-vpn1 ike gateway los-angeles-gw3
proxy-identity local 1.3.4.7/24 remote 5.2.3.4/24
```

Table 68 gives an overview of this parameter and its syntax.

Table 68: Proxy ID Parameters Configuration Statement

Proxy ID	Syntax
Specifies the local address, remote address, and service that traffic must match. By default, the router uses information in the policy. The Proxy IDs of the peers must match. One must have the other's address as its remote address.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name> proxy-identity local <ipv4-prefix> remote <ipv4-prefix> service <service-name>` ■ `vpn-name`: Name of the IPSec VPN created for the local gateway. ■ `destination-peer-gateway-name`: Name of the destination peer. This is the same name for the destination gateway as specified in the IKE Phase 1 configuration. ■ `(local) ipv4-prefix`: Local address and netmask. ■ `(remote) ipv4-prefix:` Remote address and netmask. ■ `service-name`: Name of the service (port and protocol combination) to protect.

Configuring VPN Monitoring Options

```
security --> ipsec --> vpn <vpn-name>
                vpn-monitor-options
                    optimized
                    source-interface <interface-name>
                    destination <destination-IP-address>
```

If you enable VPN monitoring, you can set options for it, including whether to optimize the process by relying on evidence of peer liveliness based on traffic patterns rather than through transmission of ICMP request messages. You can also specify the interface to be used to send ICMP messages and the destination IP address for them.

Enter the following statement in configuration mode to configure the VPN monitoring options for the Albuquerque branch VPN:

```
user@host# set security ipsec vpn albuquerque-vpn1 vpn-monitor optimized
source-interface ge-0/0/3.2 destination-ip 5.2.3.4
```

Table 69 identifies the monitoring options parameters.

Table 69: VPN Monitoring Options Configuration Statement

VPN Monitoring Options	Syntax
VPN monitoring options to disable transmission of ICMP messages and use traffic patterns instead. You also specify source and destination information.	`set security ipsec vpn <vpn-name> vpn-monitor optimized source-interface <interface-name> destination <IPv4-address>` ■ `optimized`: Directs the router to use traffic patterns as evidence of peer liveliness. In this case, the router does not send ICMP requests to verify the liveliness of the peer, disabled by default. ■ `(source-interface) interface-name`: The source interface from which the ICMP requests are sent. If you do not specify one, the router uses the local tunnel endpoint interface. ■ `(destination) IPv4-address`: Address of the destination of the ICMP request messages. If you do not specify one, the router uses the destination peer address.

Determining When to Establish Tunnels

```
security ipsec
            vpn <vpn-name>
                establish-tunnels (immediately | on-traffic)
```

You can direct the router to establish a VPN tunnel immediately after the configuration is committed instead of waiting until data traffic is transmitted and IKE must be negotiated with the destination peer gateway.

To configure the local gateway to establish a tunnel immediately after the configuration is committed, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn albuquerque-vpn1 establish-tunnels
immediately
```

NOTE: For the VPN monitor option, if establish-tunnels immediately is specified, the router renegotiates the IKE and IPSec SAs after it deletes its own set.

Table 70: Establish Tunnels Configuration Statement

Establish Tunnels Immediately	Syntax
Establish tunnels when the configuration is committed, not when data traffic triggers it. If VPN monitoring is configured, the router renegotiates the IKE and IPSec SAs after it deletes its own.	`set security ipsec vpn <vpn-name> establish-tunnels (immediately | on-traffic)` ■ `vpn-name`: The name of IPSec VPN created for the local gateway. ■ `immediately` or `on-traffic`: You must specify either option.

Common IKE IPSec Problems

Phase 1: Common Problems

This section identifies problems that administrators have encountered which pertain to IKE Phase 1 configuration:

- Proposal mismatch: The IKE Phase 1 proposals configured on each side do not agree.

  In this case, the local gateway (initiator) sees retransmissions and a retransmission limit indicator. The problem is evident at the destination gateway (responder): All proposals sent by the initiator were rejected.

- Preshared key mismatch: The keys do not match.

- No route information is configured. To establish a gateway connection, you must either configure an explicit route or a default one that is used to reach the gateway,

- It may happen that the destination gateway (responder) does not recognize the incoming request as originating from a valid peer gateway. Any of three misconfigurations at the destination peer could cause this problem:

  - The peer ID could be configured incorrectly.

  - The outgoing interface address is incorrect.

  - The peer addresses could be misconfigured. Each peer's source address must be configured as the other's destination address.

Phase 2 Common Problems

Here are some problems that can occur if errors exist in the configurations at either of the peer gateways:

- Proposal mismatch: In this case, both the local (initiator) and the destination (responder) gateways report the problem.

- For a Proxy ID mismatch, destination gateway possesses the error information that reports the problem.

  At the destination gateway, output in the system log shows address book entries configured at the local peer gateway. The administrator at the destination gateway can compare this information with the local address book, and correct it on the appropriate system.

  If you are the destination gateway (the receiver), the Proxy ID sent by your peer is displayed in the event log. In this case, the mismatch is obvious: The destination peer's local settings are not those specified in the ProxyID received from the local gateway (the sender).

Chapter 6

Implementing Policy-based IPSec VPNs Branch Offices

This chapter describes how to implement a site-to-site policy-based IPSec VPN for a branch office connecting to corporate headquarters. For background information on IPSec VPNs and design concepts, see “Designing IPSec VPNs for Branch Offices” on page 29.

NOTE: Chassis clustering is supported for policy-based IPSec VPNs. For details, see the *JUNOS Software Security Configuration Guide*.

This chapter includes the following sections:

About the Policy-based IPSec VPN Branch Deployment

The scenario described in this chapter explains how to configure a site-to-site policy-based IPSec VPN for the San Antonio branch office of the New Bank of the Southwest connecting to corporate headquarters at the Los Angeles site. Figure 13 shows the topology for the deployment.

Figure 13: Topology for a Policy-based IPSec VPN from the San Antonio Branch to Los Angeles Headquarters

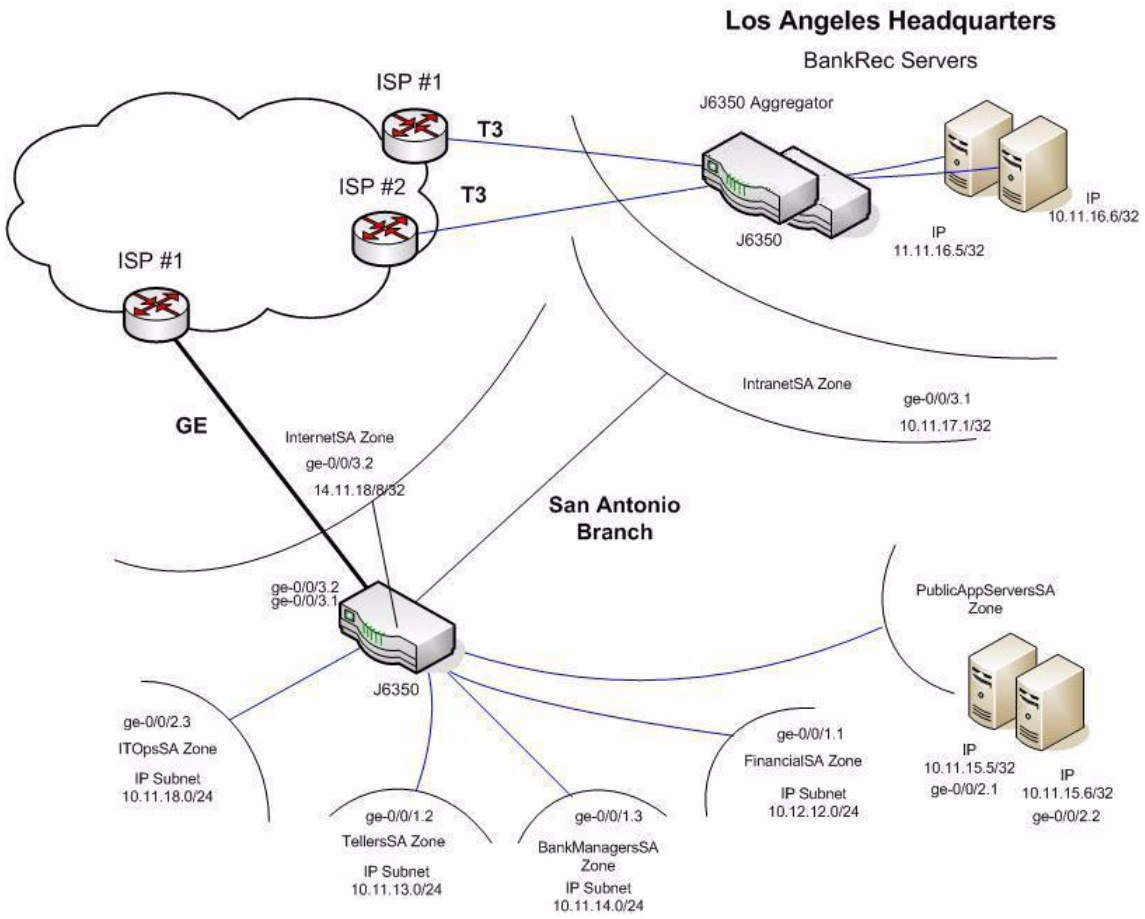

About Configuring a Policy-based IPSec VPN

This section provides background information on configuring a site-to-site policy-based IPSec VPN. Then it provides an overview of the steps the process entails. Before you begin to create an IPSec VPN, see IPSec VPN concepts in “Designing IPSec VPNs for Branch Offices” on page 29.

NOTE: Network Address Translation (NAT) is not supported for policy-based IPSec VPNs.

To implement a site-to-site policy-based IPSec VPN, you configure two fundamental parts:

- The IPSec VPN configuration itself, which identifies the relationship with the peer gateway and defines the VPN
- The VPN policy statement that specifies the VPN to be used and allows traffic to traverse it

A site-to-site IPSec VPN exists between two peers, or endpoints. To implement an IPSec VPN, administrators at each endpoint configure parameters that the peers use to negotiate cryptography used to generate security associations (SAs), or keys.

To create an IPSec VPN, you configure two sets of parameters, one for Phase 1 and another for Phase 2. The peers negotiate these parameters for different purposes.

- Phase 1 results in establishment of a secure tunnel which is used for Phase 2. The Phase 1 negotiated and agreed on cryptography is used to secure the message exchanges during Phase 2. (It is not used to transmit transit traffic securely across the IPSec VPN.)
- Using the secure tunnel generated as a result of Phase 1, during Phase 2 the peers negotiate the cryptography to be used to secure transmission of transit traffic across the VPN.

You can use a manual method to generate keys—that is, SAs—or you can use a method called Internet Key Exchange (IKE), which manages key generation for you. IKE is the recommended and most commonly used method. The sample deployment uses IKE.

Configuration of IPSec VPNs can seem confusing because you specify similar parameters for both phases. For example, for both IKE Phase 1 and Phase 2, you configure policies and proposals, which play similar roles in relation to their part of the negotiation process, but whose negotiated parameters are used in the end for different purposes.

Table 71 gives an overview of the complete procedure for creating an IPSec VPN.

Table 71: Creating a Policy-based IPSec VPN

Task	Reference
1. Configure the required security zones and interfaces.	See “Implementing Firewall Deployments for Branch Offices” on page 37. The chapter describes how to implement a firewall for a route-based IPSec, but the process is largely the same.
2. Configure a high-level IPSec parameter (invalid SPI).	See “Configuring the Security Parameter Index Check” on page 145.
3. Configure the Phase 1 peer gateway and destination peer information.	See “Configuring the IKE Phase 1 Peer Gateway and Peer Information” on page 147.
4. Configure one or more Phase 1 proposals.	See “Configuring IKE Phase 1 Proposals” on page 152.
5. Configure one or more Phase 1 policies	See “Configuring IKE Phase 1 Policies” on page 150.
6. Configure the IPSec VPN, peer gateway, and destination peer information.	See “Configuring the Phase 2 IPSec VPN, Destination Peer Gateway, and Peer Information” on page 157.
7. Configure one or more Phase 2 IPSec proposals.	See “Configuring Phase 2 Proposals” on page 160.
8. Configure one or more Phase 2 IPSec policies.	See “Configuring Phase 2 IPSec Policies” on page 164.
9. Configure firewall security policies.	See “Implementing Firewall Deployments for Branch Offices” on page 37 for an explanation of how to create security policies.

Summary of Steps for Configuring a Policy-based IPSec VPN

This section summarizes the steps you follow to configure a policy-based IPSec VPN for the San Antonio branch office communicating with its peer gateway at the Los Angeles headquarters central site. Each step refers to the section that describes its procedure.

Table 72: Steps for Implementing a Policy-based IPSec VPN

Step	Process
Configure the Peer Gateway and IKE Phase 1	
■ Configure the local gateway to handle an invalid security parameter index from its destination peer.	■ See Phase 1:"Configuring the Security Parameter Index Check" on page 145.
■ Configure IKE Phase 1.	■ See "Configuring the IKE Phase 1 Peer Gateway and Peer Information" on page 147. This section includes: ■ "Configuring the Peer Gateway Name and Destination Peer Address" on page 147. ■ "Determining the State of a Peer with Dead Peer Detection" on page 147. ■ "Configuring a Reference to the IKE Policy for the Destination Peer" on page 148. ■ "Configuring the Local (Outgoing) Interface" on page 149.
■ Configure IKE Phase 1 policies.	■ See Phase 1: "Configuring IKE Phase 1 Policies" on page 150. This section includes: ■ "Configuring the IKE Policy Name" on page 150. ■ "Configuring the Peer Authentication Parameters" on page 151. ■ "Configuring a Reference to the IKE Proposal" on page 151.
■ Configure IKE Phase 1 proposals.	■ See "Configuring IKE Phase 1 Proposals" on page 152. This section includes: ■ "Configuring the IKE Proposal Name" on page 153. ■ "Configuring the IKE Proposal Authentication Method" on page 153. ■ "Configuring the IKE DH Group" on page 154. ■ "Configuring the IKE Encryption Algorithm" on page 155. ■ "Configuring the IKE Authentication Algorithm" on page 156.
Configure the VPN and IPSec Phase 2	
■ Configure the IPSec VPN and Phase 2 destination peer gateway.	■ See "Configuring the Phase 2 IPSec VPN, Destination Peer Gateway, and Peer Information" on page 157. This section includes: ■ "Configuring the IPSec VPN and Peer Gateway Name" on page 157. ■ "Configuring the Reference to the Phase 2 Policy" on page 158. ■ "Using Antireplay Checking" on page 158. ■ "Configuring the IPSec Proposal Name" on page 160. ■ "Tunnels Establishment" on page 160.

Step	Process
■ Configure IPSec Phase 2 proposals.	■ See "Configuring Phase 2 Proposals" on page 160. This section includes: ■ "Configuring Phase 2 IPSec Protocols" on page 161. ■ "Configuring the IPSec Proposal Authentication Algorithm" on page 162. ■ "Configuring the IPSec Proposal Encryption Algorithm" on page 163.
■ Configure IPSec Phase 2 policies.	■ See "Configuring Phase 2 IPSec Policies" on page 164. This section includes: ■ "Configuring the IPSec Policy Name" on page 164. ■ "Configuring a Reference to the IPSec Proposals" on page 164. ■ "Configuring the Perfect Forward Secrecy Parameter" on page 165.

Configuring a Policy-based IPSec VPN

This section describes how to configure the primary policy-based IPSec VPN to be used for transit traffic from the San Antonio branch office to corporate headquarters at the Los Angeles site. The site uses a secondary VPN for redundancy, but because the configuration is largely similar, it is not shown here.

Policy-based IPSec VPN Configuration Set Statements Summary

This section shows the CLI commands used to configure a policy-based IPSec VPN. The complete configuration is shown after the CLI statements.

CAUTION: To use this configuration, you must commit the configuration after all of the statements are entered. If you intend to commit the statements individually, you must reverse the order of statements within each phase because of dependencies among the statements.

Configuring IKE Phase 1

Global IKE parameter. (See "Configuring the Security Parameter Index Check" on page 145.)

```
user@host# set security ike respond-bad-spi 10
```

Peer gateway name and address. (See "Configuring the Peer Gateway Name and Destination Peer Address" on page 147.)

```
user@host# set security ike gateway los-angeles-gw1 address 10.11.16.20
```

Dead peer detection (DPD). (See "Determining the State of a Peer with Dead Peer Detection" on page 147.)

```
user@host# set security ike gateway los-angeles-gw1 dead-peer-detection
always-send interval 10 threshold 3
```

IKE policy reference. (See “Configuring a Reference to the IKE Policy for the Destination Peer” on page 148.)	user@host# set security ike gateway los-angeles-gw1 ike-policy pol1
Local outgoing interface. (See “Configuring the Local (Outgoing) Interface” on page 149.)	user@host# set security ike gateway los-angeles-gw1 external-interface ge-0/0/3.1
IKE policy. (See “Configuring the IKE Policy Name” on page 150 and “Configuring the Peer Authentication Parameters” on page 151.	user@host# set security ike policy pol1 pre-shared-key ascii-text 1234pskey5678
IKE policy proposal reference. (See “Configuring a Reference to the IKE Proposal” on page 151.)	user@host# set security ike policy pol1 proposals prop1
IKE proposals. (See “Configuring IKE Phase 1 Proposals” on page 152.)	user@host# set security ike proposal prop1 authentication-method pre-shared-keys authentication-algorithm sha1 encryption-algorithm 3des-cbc user@host# set security ike proposal prop2 authentication-method rsa-signatures encryption-algorithm aes-128-cbc authentication-algorithm md5

Configuring the VPN and IPSec Phase 2

IPSec VPN and peer gateway. (See “Configuring the IPSec VPN and Peer Gateway Name” on page 157.)	user@host# set security ipsec vpn san-antonio-vpn1 ike gateway los-angeles-gw1
IPSec Phase 2 policy reference. (“Configuring the Reference to the Phase 2 Policy” on page 158.)	user@host# set security ipsec vpn san-antonio-vpn1 ike gateway los-angeles-gw1 ipsec-policy ipsec-pol1
Antireplay checking. (“Using Antireplay Checking” on page 158.)	user@host# set security ipsec vpn san-antonio-vpn1 ike gateway los-angeles-gw1 no-anti-replay-check

Don't fragment bit cleared. (See "Configuring the Don't Fragment Bit" on page 159.)	user@host# set security ipsec vpn san-antonio-vpn1 df-bit clear
Establish tunnels immediately. (See "Tunnels Establishment" on page 160.)	user@host# set security ipsec vpn san-antonio-vpn1 establish-tunnels immediately
IPSec Phase 2 proposals. (See "Configuring Phase 2 Proposals" on page 160.)	user@host# set security ipsec proposal prop1 protocol esp authentication-algorithm hmac-sha1-96 encryption-algorithm 3des-cbc user@host# set security ipsec proposal prop2 protocol ah authentication-algorithm hmac-md5-96
IPSec Phase 2 policy, proposal reference. (See "Configuring the IPSec Policy Name" on page 164 and "Configuring a Reference to the IPSec Proposals" on page 164.)	user@host# set security ipsec policy ipsec-pol1 proposal prop1
IPSec Phase 2 policy, perfect forward secrecy (PFS). (See "Configuring the Perfect Forward Secrecy Parameter" on page 165.)	user@host# set security ipsec policy ipsec-pol1 perfect-forward-secrecy keys group2

Policy-based IPSec VPN Configuration

The following example shows the same deployment presented as a configuration excerpt.

```
security {
    ike {
        respond-bad-spi 10;
        gateway los-angeles-gw1 {
            address 10.11.16.20;
            dead-peer-detection
                always-send;
                interval 10;
                threshold 3;
        }
            ike-policy pol1;
            external-interface ge-0/0/3.1;
        }
        policy pol1
            pre-shared-key ascii-text 1234pskey5678;
            proposal prop1;
        }
        proposal prop1 {
            authentication-method pre-shared-keys;
            authentication-algorithm sha1;
```

```
                encryption-algorithm 3des-cbc;
            }
            proposal prop2 {
                authentication-method pre-shared-keys;
                encryption-algorithm aes-128-cbc;
                authentication-algorithm md5;
            }

        }
        ipsec
            vpn san-antonio-vpn1 {
                ike {
                    gateway los-angeles-gw1;
                    no-anti-replay-check;
                    ipsec-policy ipsec-pol1;
                    }
                df-bit clear;
                establish-tunnels immediately
        }
        proposal prop1 {
            protocol esp;
            authentication-algorithm hmac-sha1-96;
            encryption-algorithm 3des-cbc;
        }
        proposal prop2 {
            protocol ah;
            authentication-algorithm hmac-md5-96;
        }
        policy ipsec-pol1 {
            proposal prop1;
            perfect-forward-secrecy keys group2;
        }
    }
}
```

Configuring the Security Parameter Index Check

You configure the invalid SPI feature as a global IKE statement.

Responding to Invalid Security Parameter Indexes (SPI)

```
security—ike
            respond-bad-spi <number>
```

The SAs between two peers of an IPSec VPN can become unsynchronized for some reason. To handle this situation, JUNOS software with enhanced services software supports use of the Invalid SPI response feature. If this feature is configured for a peer gateway, it can recognize and respond to an invalid or unknown SPI in a packet header from its destination peer. In this case, the router resets the state of the peer so that the two peers are synchronized again.

To configure the IPSec VPN for the San Antonio branch to respond to bad SPIs from its peer ten times before engaging the peer in a new IKE session, enter the following statement in configuration mode:

Enter following statement in configuration mode:

```
user@host# set security ike respond-bad-spi 10
```

Table 73 gives the set statement syntax for the Invalid SPI Response feature.

Table 73: Invalid SPI Response (respond-bad-spi) Configuration Statement

SPI	Syntax
Invalid SPI Response Enables one peer to recognize and handle an invalid Security Parameter Index (SPI) value in the packet header of the other peer and synchronize the state of the two peers again by initiating a new IKE session. (The IKE session generates new IPSec SA keying material.)	`set security ike respond-bad-spi <number>` ■ `number`: Number of times to respond to an invalid SPI, per gateway, before engaging the peer in a new IKE session.

How the Invalid SPI Response Feature Works

This section summarizes how the respond-bad-spi feature works, then it gives an example.

Table 74: How the respond-bad-spi Feature Works

How It Works
Peer A with respond-bad-spi configured detects IPSec packets with invalid SPI information from Peer B.
Peer A responds to the packets by initiating establishment of another IKE SA with Peer B.
After IKE SA established, Peer A deletes the old IPSec SA from Peer B
Peer A and Peer B are synchronized.

Even when dead peer detection (DPD) is enabled, the IPSec SA state of two IKE peers can become unsynchronized. In the following example, initially both Peer A at the San Antonio branch and Peer B at the Los Angeles central office have the same set of SAs:

San Antonio Peer A	Los Angeles Peer B
IKE1 SA	IKE1 SA
IPSec1 SA	IPSec1 SA

A serious problem occurs in the San Antonio network that causes all its SAs to be deleted. As a result, peer A is without SAs. Because Peer B at the Los Angeles site has deleted its IKE SA only, the two peers fall into an unsynchronized state:

San Antonio Peer A	Los Angeles Peer B
	IPSec1 SA

Peer B continues to send IPSec packets. The invalid SPI response feature allows Peer A to recognize and respond to these packets by establishing a new IKE SA with Peer B and then sending a delete message to delete Peer B's old IPSec1 SA:

San Antonio Peer A	Los Angeles Peer B
IKE2 SA	IKE2 SA
delete (IPSec1 SA) ———— >	IPSec1 SA

NOTE: A malicious attacker could spoof IPSec packets with invalid SPIs, causing the San Antonio peer to initiate an IKE process to its defined peers. To protect against this misuse, the invalid SPI response feature has the following characteristics:

- The router responds to SPIs only from statically configured peers.
- You can use the following operational mode command to clear the count of bad SPIs per gateway:

  user@host> clear security ike respond-bad-spi-count <gateway-name>

Configuring the IKE Phase 1 Peer Gateway and Peer Information

As part of IKE Phase 1, you configure information about the gateway between the local and remote peers. You configure a name for the gateway and the IP address of the destination peer. You can also configure the dead peer detection (DPD) feature to allow the local peer to determine if the remote peer is active. DPD can be used to determine the liveliness of the peer when routing protocols are not used.

Configuring the Peer Gateway Name and Destination Peer Address

```
security-->ike-->gateway
                    <destination-peer-gateway-name>
                    address <destination-peer-addr>
```

To identify the destination peer, you specify the name of the gateway to be used for it and its IP address or hostname. To configure information for Los Angeles destination peer, enter the following statement in configuration mode:

user@host# set security ike gateway los-angeles-gw1 address 10.11.16.20

Table 75 gives an overview of the gateway name and destination peer address configuration and its syntax.

Table 75: IPSec VPN Peer Gateway Name and Peer Address Configuration Statement

Destination Peer Gateway Identity	Syntax
Destination peer gateway name and peer address	`set security ike gateway <destination-peer-gateway-name> address <destination-peer-addr>` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `destination-peer--addr`: IPv4 address or hostname of the destination peer gateway.

Determining the State of a Peer with Dead Peer Detection

```
security—->ike—->gateway <destination-peer-gateway-name>
                    dead-peer-detection
                            always-send
                            interval <seconds>
                            threshold <number>
```

The deployment design for the San Antonio gateway with its Los Angeles destination peer includes use of a feature called Dead Peer Detection (DPD). This feature allows the San Antonio peer to communicate with its destination Los Angeles peer to determine if the destination peer is available.

DPD functions by sending a hello message to a destination peer if the local peer has not received traffic from it during the specified time. If it has received traffic, the peer is considered available and no hello message is sent (which conserves on CPU resources).

Problems can ensue if a peer continues to send packets to a partner that is no longer active. If a peer no longer exists, sending packets incurs the unnecessary processing expense associated with encryption, authentication, and route lookup only to result in packet loss. DPD addresses these problems by determining the state of the peer and ending the connection if it determines that the peer is unavailable.

To configure the San Antonio-to-Los Angeles gateway to include use of DPD, enter the following statement in configuration mode:

```
user@host# set security ike gateway los-angeles-gw1 dead-peer-detection
always-send interval 10 threshold 3
```

Table 76 gives an overview of the dead peer detection feature and its syntax.

Table 76: Dead Peer Detection Configuration Statement

Dead Peer Detection	Syntax
Dead Peer Detection (DPD) Enables the system to determine if its peer is active and to declare the peer dead if it does not respond after a specified interval, If a destination peer is known to be dead, the local peer stops transmitting data.	`set security ike gateway <destination-peer-gateway-name> dead-peer-detection always-send interval <seconds> threshold <number>` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `always-send`: Send DPD requests regardless of traffic patterns. ■ `(interval) seconds`: Number of seconds that the peer waits for traffic from its destination peer before sending a DPD request packet. (10-60) (default 10) ■ `(threshold) number`: Maximum number of unsuccessful DPD requests sent after which the system declares the peer dead. (1-5) (default 5)

Configuring a Reference to the IKE Policy for the Destination Peer

```
security—>ike—>gateway
                    <destination-peer-gateway-name>
                        ike-policy <ike-policy-name>
```

For a destination peer, you configure a reference to the name of the IKE policy to be used. You refer to the policy in the ike-policy statement of the gateway configuration.

To specify the IKE policy to use for the Los Angeles peer, enter the following statement in configuration mode:

```
user@host# set security ike gateway los-angeles-gw1 ike-policy pol1
```

Table 77: IKE Phase 1 Peer Policy Reference Configuration Statement

IKE Phase 1 Policy Reference	Syntax
Refers to the IKE policy to be used for communication with the destination peer gateway. In turn, the policy refers to the proposal containing cryptography parameters to present to the peer for Phase 1 negotiation.	`set security ike gateway <destination-peer-gateway-name> ike-policy <ike-policy-name>` ■ `destination-peer-gateway-name`: Name of the gateway, specified as an alphanumeric string. ■ `ike-policy-name`: Name of the policy to be used for the gateway.

Configuring the Local (Outgoing) Interface

```
security—>ike—->gateway
                    <destination-peer-gateway-name>
                            external-interface <IP-address>
                            local-identity <local-peer-address>
```

You must configure an interface to be used for the IPSec VPN. To configure the local interface for the San Antonio gateway, enter the following statement in configuration mode:

```
user@host# set security ike gateway los-angeles-gw1 external-interface
ge-0/0/3.1
```

You can also specify a local identity to send to the destination peer for it to use to refer to the local peer. However, this parameter is optional. If you do not specify one, the IP address of the external interface is used.

Table 78 gives an overview of these parameters and their syntax.

Table 78: External Interface and Local Peer Identity Configuration Statement

Local Peer Interface and Identity	Syntax
The external-interface parameter specifies the local outgoing interface to be used for the IPSec VPN. The local identity parameter specifies the identity of the local peer so that its partner destination gateway can communicate with it. If you do not specify a local identity, the system defaults to using the IP address of the external interface. You do not need to specify both the external interface and the local identity.	`set security ike gateway <destination-peer-gateway-name> external-interface <interface-name> local-identity (inet IPv4-address | hostname <fully-qualified-domain-name> | user-at-hostname <email-address> | distinguished-name)` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `interface-name`: Name of the interface to be used for the IPSec VPN. ■ `local-identity`: One of the following means of identifying the local peer, to be used by the destination peer: ■ `(inet) IPv4-address`: Valid IPv4 address ■ `(hostname) fully qualified domain name`: String representing a hostname. ■ `(user-at-hostname) e-mail-address`: String representing the email address. ■ `distinguished-name`: Defaults to the distinguished name (DN) from the certificate. Normal certificate rules apply.

Configuring IKE Phase 1 Policies

An IKE policy specifies the authentication method to be used by the peers to identify themselves to each other during their Phase 1 exchanges: either a certificate or a preshared key. It also refers to the Phase 1 proposal to be presented to the peer for negotiation.

Configuring the IKE Policy Name

```
security —> ike—>
            policy <policy-name>
```

Each IKE policy has a name, or label, which is referred to by the ike-policy statement of the peer gateway configuration.

To specify the IKE policy name for the San Antonio gateway, enter the following statement in configuration mode:

```
user@host# set security ike policy pol1
```

Table 79 gives an overview and syntax for the IKE policy name.

Table 79: IKE Policy Name Configuration Statement

IKE Policy Name	Syntax
Name of the IKE policy to be configured. An IKE gateway configuration refers to the policy to be used for it in its ike-policy parameter.	`set security ike policy <policy-name>` `policy-name`: ASCII text string.

Configuring the Peer Authentication Parameters

```
security—>ike—->
            policy <policy-name>
                (pre-shared-key ascii-text <key> |
                            certificate <certificate-information>)
```

During Phase 1 of the IKE negotiations, the peers authenticate themselves to each other. For this purpose, they use either preshared keys or certificates. You configure the authentication type and its parameters in the IKE policy.

In the proposal, referred to from within the policy, you specify the authentication method. For example, if you specify a certificate and its arguments in the policy, the proposal that the policy refers to can specify the kind of certificate to be used, such as RSA.

To specify the type of authentication to be used for the Los Angeles peer gateway, enter the following statement in configuration mode:

```
user@host# set security ike policy pol1 pre-shared-key ascii-text 1234pskey5678
```

Table 80 gives an overview of the authentication method specification.

Table 80: IKE Policy Authentication Method Configuration Statement

Peer Authentication Method and Parameters	Syntax

Authentication method to use to secure a tunnel to be used for Phase 2.
For details on use of certificates, see the *JUNOS Software Security Configuration Guide*.
	`set security ike policy <policy-name> pre-shared-key ascii-text <key> | certificate <certificate-information>)`
	■ `policy-name`: ASCII text string.
	■ (If pre-shared-key) `key`: Unreadable string
	■ (If certificate) `certificate-information`: Certificate configuration parameters
	  ■ `(local-certificate) certificate ID`
	  ■ `(trusted-ca) CA-ID | use-all`
	  ■ `(peer-certificate-type) pkcs7 | x509-signature`

Configuring a Reference to the IKE Proposal

```
security—>ike—>
            policy <policy-name>
                        proposals <proposal-name>
```

The IKE policy includes a reference to one or more IKE proposals whose parameters are presented to the destination peer for negotiation. The agreed-upon values are then used to secure a tunnel with the destination peer to be used for Phase 2 negotiations.

To refer to the IKE proposal to be negotiated with the Los Angeles peer, enter the following statement in configuration mode:

```
user@host# set security ike policy pol1 proposal prop1
```

Table 81 gives an overview of the parameter and its syntax.

Table 81: IKE Phase 1 Proposal Reference Configuration Statement

IKE Phase 1 Proposal Reference	Syntax
A reference to one or more IKE Phase 1 proposals to be used for Phase 1 negotiations.	`set ike policy <policy-name> proposal [<proposal1> <proposal2>...] proposal-set <proposal-set-name>` ■ `policy-name`: ASCII text string. ■ `proposal`: Reference to the name(s) of one or more proposals to be used. ■ `proposal-set-name`: reference to a predefined proposal set to be used.
Predefined Phase 1 proposals sets: ■ Basic: ■ Proposal 1: preshared keys, g1, des, sha1 ■ Proposal 2: preshared keys, g1, des, md5 ■ Compatible: ■ Proposal 1: preshared keys, g2, 3des, sha1 ■ Proposal 2: preshared keys, g2, 3des, md5 ■ Proposal 3: preshared keys, g2, des, sha1 ■ Proposal 4: preshared keys, g2, des, md5 ■ Standard: ■ Proposal 1: preshared keys, g2, 3des, sha1 ■ Proposal 2: preshared keys, g2, aes128, sha1	

Configuring IKE Phase 1 Proposals

During Phase 1, the peers negotiate cryptography parameters that determine how the tunnel for Phase 2 negotiations is to be secured. These parameters, such as authentication and encryption algorithms, are specified in proposals. If more than one proposal is specified, they are presented to the destination peer, one after another. If you intend to specify more than one proposal for negotiation, it is a good idea to specify the most restrictive one first.

For the negotiation to work, the destination peer also must have configured at least one of the proposals presented to it by the local gateway. A successful Phase 1 negotiation concludes successfully when both peers agree to accept at least one set of the Phase 1 proposal parameters, and then process them.

Configuring the IKE Proposal Name

```
security ——>ike ——> proposal <proposal-name>
```

You provide a name, or label, for an IKE proposal to identify it so that it can be referred to from within the policy configuration used for a gateway.

To specify the names for the two proposals for the San Antonio IPSec VPN, enter the following statements in configuration mode:

```
user@host# set security ike proposal prop1

user@host# set security ike proposal prop2
```

You can configure the proposal name as part of the components that compose the proposal. This section shows how to configure the name alone for clarity.

Table 82 shows the IKE proposal name configuration and syntax.

Table 82: IKE Proposal Name Configuration Statement

IKE Proposal Name	Syntax
The name, or label, of the IKE Phase 1 proposal that follows.	`set security ike proposal <proposal-name>` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. The negotiated and agreed-upon parameters of a proposal produce a secure tunnel for Phase 2 negotiations.

Configuring the IKE Proposal Authentication Method

```
security ——> ike ——> proposal <proposal-name>
                     authentication-method
                        (pre-shared-keys | rsa-signatures)
```

Within a proposal, you specify the type of authentication to be used. If you specified in the policy that certificates are to be used, in the proposal you specify the type of certificate to use.

If you specified a certificate within the policy, the policy must refer to a proposal that specifies the method to use, such as RSA. If you specified a preshared key within the policy, the policy must refer to a proposal that specifies preshared keys as the authentication method.

For the San Antonio gateway, two proposals are configured. To specify their authentication methods, enter the following statements in configuration mode:

```
user@host# set security ike proposal prop1 authentication-method pre-shared-keys

user@host# set security ike proposal prop2 authentication-method rsa-signatures
```

NOTE: The proposal option pre-shared-keys is plural, because it refers to a method. The policy specification pre-shared-key is singular because it refers to a specific key. You configure a specific key from within the policy. From within the proposal, you indicate that you want to use preshared keys.

Table 83 shows the IKE proposal authentication method.

Table 83: IKE Proposal Authentication Method Configuration Statement

IKE Proposal Authentication Method	Syntax
The kind of authentication method to use. ■ For a preshared key, the *policy* gives the specific key. The *proposal* states that pre-shared-keys are to be used. ■ For a certificate. the policy certificate parameter identifies which certificate is to be used, but the proposal tells the kind of certificate to use.	`set security ike proposal <proposal-name> authentication-method <method>` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(authentication-method) method`: Specifies the authentication method and corresponds to the policy specification. For the method, specify one of the following options: ■ `pre-shared-keys` ■ `rsa-signatures`

Configuring the IKE DH Group

```
security—>ike—> proposal <proposal-name>
                    dh-group (group1 | group2 | group5)
```

NOTE: The San Antonio IKE Phase 1 proposal configurations do not contain Diffie-Hellman parameters. This section is provided for completeness.

A Diffie-Hellman group performs key exchange to generate keying material for per-packet encryption and authentication. It allows the peers to generate a shared secret value over an unsecured medium without passing the secret value through the wire. There are five Diffie-Hellman groups. The JUNOS software with enhanced services software supports the following three groups:

- DH Group 1: 768-bit modulus
- DH Group 2: 1024-bit modulus
- DH Group 5: 1536-bit modulus

A larger modulus indicates a more secure key. However, the larger the modulus the longer it takes to generate the key. More CPU cycles are utilized for generation of longer keys.

Because the modulus for each Diffie-Hellman group is a different size, the peers must agree to use the same group.

Table 84 gives an overview of the Diffie-Hellman group parameter and its syntax.

Table 84: IKE DH Group Configuration Statement

IKE Proposal DH Group	Syntax
Species the Diffie-Hellman group to use to increase the security of the tunnel. A Diffie-Hellman exchange allows the peers to produce a shared secret over an unsecured medium without passing the secret through the wire. The size of the prime modulus used in each group differs. The larger the size, the more secure the key.	`set ike proposal <proposal-name> dh-group (group1 | group2 | group5)` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(dh-group)`: Specify one of the following options to indicate the Diffie-Hellman group to be used. ■ `group1` (768-bit modulus) ■ `group2` (1024-bit modulus) ■ `group5` (1536-bit modulus) Because the modulus for each Diffie-Hellman group is a different size, the peers must agree to use the same group.

Configuring the IKE Encryption Algorithm

```
security—>ike—>proposal <proposal-name>
                encryption-algorithm (des-cbc |3des-cbc | aes-128-cbc |
                                          aes-192-cbc| aes-256-cbc)
```

During IKE Phase 1, the peers negotiate the cryptography parameters to be used for their Phase 2 exchanges. These parameters are specified as part of a proposal. The encryption algorithm agreed on by the two peers is used to encrypt parameters that they exchange and negotiate in IKE Phase 2.

A proposal can contain one encryption algorithm. However, because a policy can refer to more than one proposal, several different encryption methods can be proposed to the peer.

To configure the encryption algorithms for the two IKE Phase 1 proposals for the San Antonio IPSec VPN, enter the following statements in configuration mode:

```
user@host# set security ike proposal prop1 encryption-algorithm 3des-cbc

user@host# set security ike proposal prop2 encryption-algorithm aes-128-cbc
```

Table 85 gives an overview of the parameter and its syntax.

Table 85: IKE Encryption Algorithm Configuration Statement

IKE Proposal Encryption Algorithm	Syntax
Encryption algorithm to be presented to the destination peer as part of the proposal. The agreed-upon algorithm is used to secure exchanges between the peers during IKE Phase 2.	`set ike proposal <proposal-name> encryption-algorithm (des-cbc | 3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc)` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(encryption-algorithm)`: Specify one of the five encryption algorithms to use in the proposal. Here are the options, in order from strongest to weakest: ▪ `aes-256-cbc` ▪ `aes-192-cbc` ▪ `aes-128-cbc` ▪ `3des-cbc` ▪ `des-cbc`

Configuring the IKE Authentication Algorithm

```
security—>ike—>proposal <proposal-name>
                        authentication-algorithm (md5 | sha1)
```

During IKE Phase 1, the peers negotiate the cryptography parameters to be used for their IKE IPSec Phase 2 exchanges. These parameters are specified as part of a proposal.

The authentication algorithm agreed on by the two peers is used to authenticate messages that they exchange and negotiate in IKE Phase 2 over the IKE Phase 1 SA, or tunnel.

A proposal can contain one authentication algorithm, but because a policy can refer to more than one proposal, several different authentication methods can be proposed to the peer.

To configure the authentication algorithms for the two IKE Phase 1 proposals for the San Antonio IPSec VPN, enter the following statements in configuration mode:

```
user@host# set security ike proposal prop1 authentication-algorithm sha1

user@host# set security ike proposal prop2 authentication-algorithm md5
```

Table 86 shows the configuration and syntax for the IKE proposal authentication algorithm.

Table 86: IKE Proposal Authentication Algorithm Configuration Statement

IKE Proposal Authentication Algorithm	Syntax
Authentication algorithm to propose to the peer to be used to authenticate exchanges between the peers during IKE Phase 2. A proposal contains the authentication algorithm. The peer must accept the entire proposal.	`set ike proposal <proposal-name> authentication-algorithm (md5 | sha1)` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(authentication-algorithm)`: Specify one of the two authentication algorithms to use in the proposal. Options are: `sha1` and `md5`.

Configuring the Phase 2 IPSec VPN, Destination Peer Gateway, and Peer Information

After you configure IKE Phase 1 parameters, you configure the IPSec VPN for the local gateway. Within the IPSec VPN specification, you identify the interface to be used for the gateway and the address of the local peer.

Configuring the IPSec VPN and Peer Gateway Name

```
security-->ipsec-->vpn <vpn-name>
               gateway <destination-peer-name>
```

You can configure the IPSec VPN name and its destination peer gateway name together. The destination peer gateway name introduces the configuration for it.The gateway name used for Phase 2 must be the same as that used for Phase 1.

To specify the name of the IPSec VPN and its destination peer, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn san-antonio-vpn1 ike gateway los-angeles-gw1
```

Table 87 gives an overview of these two parameters and the statement syntax.

Table 87: VPN Name and Destination Peer Gateway Name Configuration Statement

IPSec VPN name and Destination Peer Gateway Name	Syntax
The name of the IPSec VPN being configured for the local gateway and the name of the destination peer whose gateway you are configuring.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name>` ■ `ipsec-vpn-name`: Name of the overall IPSec VPN for the local peer. ■ `destination-peer-gateway-name`: Name of the destination peer. This name must be the same as the one that you specified for the gateway in the IKE Phase 1 configuration.

Configuring the Reference to the Phase 2 Policy

```
security-->ipsec-->vpn <vpn-name>
                ike
                    gateway <destination-peer-gateway-name>
                    ipsec-policy <ipsec-policy-name>
```

For each destination peer gateway, you configure a reference to the name of the IPSec policy to be used for that gateway. You refer to the IPSec policy to be used in the ipsec-policy parameter of a gateway configuration.

To refer to the IPSec policy to use for the Los Angeles gateway, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn san-antonio-vpn1 ike gateway los-angeles-gw1
ipsec-policy ipsec-pol1
```

Table 88 gives the syntax for this parameter.

Table 88: IPSec Policy Reference Configuration Statement

Reference to the IPSec Policy to Use	Syntax
Refers to the IPSec policy to be used for communication with the specified destination peer gateway.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name> ipsec-policy <ipsec-policy-name>` ■ `vpn-name`: A string identifying the IPSec VPN created for the local gateway. ■ `destination-peer-gateway-name`: The name of the destination peer. This is the same name for the destination gateway as specified in the IKE Phase 1 configuration. ■ `ipsec-policy-name`: Name of the policy, configured outside the destination peer to be used for the peer gateway.

Using Antireplay Checking

```
security-->ipsec-->vpn <vpn-name>
                ike
                 gateway <destination-peer-gateway-name>
                    no-anti-replay
```

Antireplay checking makes it impossible for an intruder to replay a packet. Attackers can intercept packets and insert altered ones into the data stream. They can use intercepted packets later to flood the system causing a denial-of-service (DoS) attack. They can also use them to gain entry to a trusted system. Anti-reply checking detects packets that match the sequence numbers of those that have already been received, and it discards any replayed packets.

By default, the antireplay feature is enabled. It is a good idea to leave the feature enabled when you are running the system online. However, to avoid processing incurred by use of the feature, you might want to disable it when you are testing or debugging the system. Also, you might want to disable antireplay checking to resolve compatibility issues with third-party peers. However, in such cases, you weaken the security of your system.

For the Los Angeles destination peer gateway configuration, antireplay checking remains enabled.

If you wanted to disable it for the Los Angeles peer, you would enter the following statement in configuration mode:

```
user@host# set security ipsec vpn san-antonio-vpn1 ike gateway los-angeles-gw1
no-anti-replay
```

Table 89 gives an overview of this parameter and its syntax.

Table 89: Antireplay Checking Configuration Statement

Antireplay Checking	Syntax

Disables or enables the feature that detects packets that match the sequence numbers of those that have already been received, and helps to discard replayed ones.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name> no anti-replay`
	
	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name> set anti-replay`
	■ `vpn-name`: Name of the IPSec VPN created for the local gateway.
	■ `destination-peer-gateway-name`: The name of the destination peer. This is the same name for the destination gateway as specified in the IKE Phase 1 configuration.
	■ `no anti-replay`: Disables antireplay checking.
	■ `set anti-replay`: Enables antireplay checking.

Configuring the Don't Fragment Bit

```
security-->ipsec-->vpn <vpn-name>
                    df-bit clear
```

The Don't Fragment (DF) bit feature allows you to set or clear the DF bit. Based on your configuration, the DF bit is constructed in the outer IP header. By default, the bit is cleared. The San Antonio site accepts the default value for the DF bit. This statement description is included for completeness.

Table 90 gives an overview of this parameter configuration.

Table 90: IPSec VPN Don't Fragment Bit Configuration Statement

IPSec VPN DF Bit	Syntax
Specifies that packets are not to be fragmented. Clearing the DF bit specifies that packet fragmentation is allowed, to accommodate smaller MTU sizes.	`set security ipsec vpn <vpn-name> df-bit clear` ■ `vpn-name`: Name of the IPSec VPN created for the local gateway. ■ `df-bit clear`: Disables use of packet fragmentation.

Tunnels Establishment

```
security-->ipsec-->vpn <vpn-name>
                    establish-tunnels (immediately | on-traffic)
```

This parameter allows you to specify whether IKE is triggered after the VPN configuration is committed to bring up the tunnel immediately or when traffic destined for the VPN is generated.

To configure the San Antonio gateway to initiate the IKE process immediately after the IPSec VPN configuration is committed, enter the following statement in configuration mode:

```
user@host# set security ipsec vpn san-antonio-vpn1 establish-tunnels
immediately
```

Table 91 gives an overview of the feature and its syntax.

Table 91: Establish Tunnels Immediately Configuration Statement

Establish Tunnels Immediately	Syntax
Specifies that the IPSEc VPN tunnel is to be brought up either immediately after the VPN configuration is committed or when traffic that requires it is generated.	`set security ipsec vpn <vpn-name> establish-tunnels (immediately | on-traffic)` ■ `vpn-name`: The name of IPSec VPN created for the local gateway. ■ `immediately` or `on-traffic`: You must specify either option.

Configuring Phase 2 Proposals

During IPSec Phase 2, the peers negotiate cryptography parameters that determine how transit traffic that traverses the IPSec VPN is to be secured. These parameters identify the kind of authentication and encryption to be used.

A proposal is referred to within an IPSec policy. An IPSec policy is referred to in the ipsec-policy statement of a peer gateway configuration.

Configuring the IPSec Proposal Name

```
security-->ipsec
           proposal <proposal-name>
```

Two proposals are configured for the San Antonio IPSec VPN. To configure the names for the San Antonio IPSec VPN proposals, enter the following statements in configuration mode:

```
user@host# set security ipsec proposal prop1

user@host# set security ipsec proposal prop2
```

You can configure the proposal name as part of the components that compose it. It is shown here separately for clarity.

Table 92 gives an overview of the IPSec proposal name specification and its syntax.

Table 92: IPSec Proposal Name Configuration Statement

IPSec proposal name	Syntax
IPSec proposal name.	`set security ipsec proposal <proposal-name>` ■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal.

Configuring Phase 2 IPSec Protocols

```
security-->ipsec
            proposal <proposal-name>
                protocol (esp | ah)
```

Transit traffic traversing an IPSec VPN is protected by either of two protocols that treat the data as a payload by encapsulating it and then encrypting and authenticating it or merely authenticating it, depending on the protocol. They are Encapsulating Security Payload (ESP) or Authentication Header (AH).

Two IPSec proposals are configured for the San Antonio IPSec VPN. To configure their IPSec protocols, enter the following statements in configuration mode:

```
user@host# set security ipsec proposal prop1 protocol esp

user@host# set security ipsec proposal prop2 protocol ah
```

Table 93 gives an overview of the IPSec protocols configuration and syntax.

Table 93: IPSec Protocol Specification Configuration Statement

IPSec Protocols	Syntax

IPSec protocol used to authenticate and encrypt data traffic traversing the IPSec VPN.	`set security ipsec proposal <proposal-name> protocol (esp | ah)`
	■ `proposal-name`: Name of the IPSec proposal whose protocol you are configuring
	■ `(protocol)`: Either of the following protocols:
	  ■ `esp`: Encapsulating Security Payload
	  ■ `ah`: Authentication Header

IPSec Protocols	Syntax
■ esp: IPSec protocol that provides the following services: ■ Encrypts data, rendering it undecipherable. ■ Data origin authentication ■ Data integrity ■ Antireplay protection ■ What it doesn't provide: ■ It authenticates and encrypts only part of the packet, excluding the IP header. Therefore, It protects only the data portion of the packet.	
■ ah: IPSec protocol that provides the following services: ■ Data origin authentication ■ Antireplay protection for most of the packet ■ What it doesn't provide: ■ Antireplay protection of the fields in the IP header which can change during transit. ■ Encryption ■ Confidentiality	

Configuring the IPSec Proposal Authentication Algorithm

```
security-->ipsec
            proposal <proposal-name>
                    authentication algorithm (hmac-md5-96 | hmac-sha1-96)
```

Each of the IPSec protocols—ESP and AH—uses an algorithm to verify the authenticity and integrity of the content of the packet and its origin. You specify the algorithm to use as part of proposal to be presented to the destination peer for negotiation.

A packet is authenticated through use of a checksum calculated via a hash-based message authentication code (HMAC) and either Message Digest version 5 (MD5) or Secure Hash Algorithm-1 (SHA-1) hash functions, which are defined as follows:

- MD5—This algorithm 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 to verify content and source authenticity and integrity.

- SHA-1—This algorithm produces a 160-bit hash from a message of arbitrary length and a 20-byte key. It is regarded as more secure than MD5 because of the larger hashes it produces. However, because the computational processing is done in the ASIC, the additional performance cost is negligible.

For the example scenario, two proposals are configured, one of which uses ESP and one of which uses AH. For the ESP proposal HMAC SHA-1 is used. For the AH proposal, HMAC MD5 is used.

To configure the authentication algorithm for proposal 1, which uses ESP, enter the following statement in configuration mode:

```
user@host# set security ipsec proposal prop1 authentication-algorithm
    hmac-sha1-96
```

To configure the authentication algorithm for proposal 2, which uses AH, enter the following statement in configuration mode:

```
user@host# set security ipsec proposal prop2 authentication-algorithm
    hmac-md5-96
```

Table 94 shows an overview of the authentication algorithm specification and syntax.

Table 94: IPSec Proposal Authentication Algorithm Configuration Statement

IPSec proposal authentication algorithm	Syntax

Authentication algorithm to be used for the IPSec protocol authentication process.	`set security ipsec proposal <proposal-name> authentication-algorithm (hmac-md5-95 | hmac-sha1-96)`
	■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal containing the IPSec protocol for which you are specifying an authentication algorithm.
	■ `(authentication-algorithm)`: One of the following options:
	  ■ `hmac-md5-95`
	  ■ `hmac-sha1-96`

Configuring the IPSec Proposal Encryption Algorithm

```
security-->ipsec
            proposal <proposal-name>
                    encryption algorithm (3des-cbc| aes-128-cbc | aes-192-cbc |
                                                aes-256-cbc | des-cbc)
```

The IPSec ESP protocol encrypts transit traffic as well as authenticating it. If encryption is to be used, you specify the encryption algorithm as part of a proposal that specifies ESP as its IPSec protocol.

For proposal 1, you specify the encryption algorithm to use by entering the following statement in configuration mode:

```
user@host# set security ipsec proposal prop1 encryption-algorithm 3des-cbc
```

Table 95 gives an overview of the parameter and its syntax.

Table 95: IPSec Proposal Encryption Algorithm Configuration Statement

IPSec proposal encryption algorithm	Syntax
Identifies the encryption algorithm to be used for the IPSec ESP protocol.	`set security ipsec proposal <proposal-name> encryption-algorithm (hmac-md5-95 | hmac-sha1-96)` ■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal. ■ `(encryption-algorithm)`: Specified as a keyword: ■ 3des-cbc ■ aes-128-cbc ■ aes-192-cbc ■ aes-256-cbc ■ des-cbc

Configuring Phase 2 IPSec Policies

An IPSec policy refers to the proposals to be negotiated with the peer. If perfect forward secrecy is used, the policy also contains its specification.

Configuring the IPSec Policy Name

```
security-->ipsec
            policy <policy-name>
```

Each Phase 2 IPSec policy has a name, or label. The destination peer gateway configuration species a policy name to identify the policy to be used.

You can specify the IPSec policy name as part of the statements that compose it. However, this section shows how to configure it separately for clarity:

```
user@host# set security ipsec policy ipsec-pol1
```

Table 96 defines the policy name parameter.

Table 96: IPSec Policy Name Configuration Statement

IPSec Policy Name	Syntax
Name of an IPSec policy. A peer gateway configuration refers to the policy to be used for it in its policy statement.	`set security ipsec policy <policy-name>` ■ `policy-name`: An alphanumeric string identifying the policy.

Configuring a Reference to the IPSec Proposals

```
security-->ipsec
            policy <policy-name>
                proposals <proposal-name>| proposal-set <proposal-set-name>
```

Similar to IKE Phase 1, for Phase 2 an IPSec VPN policy statement contains a reference to the proposal to be presented to the destination peer for negotiation.

A proposal specifies parameters to be negotiated with the partner. The agreed-upon values are used to secure transit traffic traversing the IPSec VPN.

To configure the proposal reference for the San Antonio VPN, enter the following statement in configuration mode.

```
user@host# set security ipsec policy ipsec-pol1 proposal prop1
```

Table 97 defines the IPSec proposal reference statement syntax.

Table 97: IPSec Phase 2 Proposal Reference Configuration Statement

IKE Phase 2 Proposal Reference	Syntax
A reference to one or more IPSec Phase 2 proposals or proposal sets to be used to negotiate security for the IPSec tunnel.	`set security ipsec policy <policy-name> proposal [<proposal1> <proposal2>...] | proposal-set <proposal-set-name>` ■ `policy-name`: ASCII text string. ■ `proposal1, proposal2`: A reference to the name(s) of one or more proposals to be negotiated with the peer. ■ `proposal-set-name`: A reference to a predefined proposal set to be negotiated with the peer.
Predefined Phase 2 proposals sets: ■ Basic: ■ Proposal 1: esp-no pfs, des, sha1 ■ Proposal 2: esp-no pfs, des md5 ■ Compatible: ■ Proposal 1: esp-no pfs, 3des, sha1 ■ Proposal 2: esp-no pfs, 3des, md5 ■ Proposal 3: esp-no pfs, des, sha1 ■ Proposal 4: no pfs, des, md5 ■ Standard: ■ Proposal 1: esp, g2 (pfs), 3des, sha1 ■ Proposal 2: esp, g2, aes128, sha1	

Configuring the Perfect Forward Secrecy Parameter

```
security-->ipsec
            policy <policy-name>
                    perfect-forward-secrecy keys <Diffie-Hellman-group>
```

Perfect Forward Secrecy (PFS) is a method for deriving Phase 2 keys independent of and unrelated to keys used for previously generated IPSec SAs. In IKE Phase 2, if PFS is configured, a second Diffie-Hellman exchange is performed. As a consequence, keying material for each IPSec SA is completely independent of every other SA, including the IKE SA.

Alternatively, without the second DH exchange of PFS, the Phase 1 proposal creates the key (the SKEYID_d key) from which all Phase 2 keys are derived. Because these keys are derived from a single key, processing time is faster than it is when PFS is used. However, the process is less secure—if an unauthorized person obtains the original key, keys derived from it are compromised. Because a second DH exchange occurs with PFS, processing is slower but the keys are more secure.

If you configure PFS, you must specify the Diffie-Hellman group to be used.

- Diffie-Hellman groups:
  - DH group 1: 768-bit modulus
  - DH group 2: 1024-bit modulus
  - DH group 5: 1536-bit modulus

Because the modulus for each DH group is a different size, the peers must agree to use the same group. The larger the modulus, the stronger the key.

To configure PFS for the ipsec-pol1 policy, enter the following statement in configuration mode.

```
user@host# set security ipsec policy ipsec-pol1 perfect-forward-secrecy keys
group2
```

Table 98 defines the PFS syntax.

Table 98: Perfect Forward Secrecy Configuration Statement

Perfect Forward Secrecy	Syntax
A method for deriving Phase 2 keys independent of and unrelated to keys used for previously generated IPSec SAs.	`user@host# set ipsec policy <policy-name> perfect-forward-secrecy keys (group1 | group2 | group5)` ■ `policy-name`: Name of the policy to be created. ■ `(keys) DH group`: The Diffie-Hellman group to be used for PFS.

About Routing and IPSec VPNs

To establish a gateway connection, you must either configure an explicit route or a default one that is used to reach the gateway. If a route to the gateway is not configured, it is not possible for the packets to travel through the VPN to the destination peer.

NOTE: This chapter does not explore how to configure virtual routing instances and routes. For details on how to configure a static route, see “Implementing Firewall Deployments for Branch Offices” on page 37.

For complete information describing how to configure static and dynamic routes, see the *JUNOS Software Interfaces and Routing Configuration Guide*.

Chapter 7

Implementing Remote Access IPSec VPN for Branch Offices

This chapter describes how to implement an IPSec VPN to support a remote access peer whose IP address is dynamically assigned. It also describes how to implement IPSec VPNs for remote access users who travel or work at home.

IPSec VPNs push out the network's perimeter by including support for remote access users. If they are granted permission by security policies at the remote sites, users working at branch, regional, and corporate sites can connect to servers and hosts across the IPSec VPN to gain access to services and resources securely and transparently.

Users working at home or traveling can check their e-mail, compose and read documents, and run other applications as if they were at their branch offices or connecting from the branch office to the central site. Unlike dialup connectivity in the past, which was characterized by the high cost of leased lines and hardware maintenance, IPSec VPN remote access connections offer enterprise corporations a cost-effective way to provide these users with secure connectivity. For example, employees on the road can access corporate resources worldwide and incur only local charges to connect to a service provider.

This chapter includes the following sections:

About Dynamically Assigned IP Addresses and IPSec VPNs

When a participant in an IPSec VPN has an IP address that is assigned dynamically, it changes the relationship between the two participants. Here are two cases in which one of the IPSec VPN participants has a dynamic IP address.

- At many sites, Dynamic Host Configuration Protocol (DHCP) servers assign IP addresses to hosts when those hosts join the protected network. Because these addresses are not permanent, they cannot be used directly to identify the host. Therefore, they cannot be used in an IPSec VPN configuration to specify the address of a gateway.
  - The configuration for an IPSec VPN gateway must specify the IP address of its destination peer.
  - If a peer has a dynamically assigned IP address, it must initiate establishment of the tunnel with its partner because it cannot determine the peer's IP address.
- When a user working at home or traveling uses an Internet Service Provider (ISP) to connect to a branch office or corporate headquarters across an IPSec VPN, the user is assigned an IP address dynamically.

  In this scenario, there is no peer gateway at the user end of the IPSec VPN. The VPN extends directly to the client software. The user launches the IPSec VPN client software, which begins the IKE identification and setup process.

About Dynamic Endpoint IPSec VPNs and Remote Access Connections

Network administrators at either end of a potential IPSec VPN configure the gateways to negotiate parameters used to establish a secure tunnel. Different conditions apply to the configuration and tunnel setup initiation, depending on the type of site-to-site VPN. Here are the two types of site-to-site IPSec VPNs, their configuration requirements and tunnel setup initiation processes:

Peer-to-Peer IPSec VPN

- For each peer, the physical interface or subinterface used as the outgoing interface has a fixed IP address.
- The configurations for both peers specify the other's address as an IP address.

  For background on IPSec VPN tunnels, see "Designing IPSec VPNs for Branch Offices" on page 29.

Dynamic Endpoint IPSec VPNs

- The physical interface or subinterface used as the outgoing interface of one peer has a dynamically assigned IP address.
- The physical interface or subinterface used as the outgoing interface of the other peer has a static IP address.
- The configuration for the peer with the dynamically assigned IP address specifies its remote peer's address as an IP address.

- The configuration for the peer with the static IP address specifies its remote peer's address as a dynamic address using a method other than an IP address. See "Configuring the Peer Gateway Name and Destination Address" on page 174.

Using the JUNOS software with enhanced services, you can configure the local IPSec peer to respond to negotiations initiated by a remote peer with a dynamically assigned IP address. "Tasks for Implementing a Dynamic Endpoint Site-to-Site IPSec VPN" on page 169 gives an example of this configuration.

Remote Access Connections

IPSec VPNs initiated from a remote peer with an Internet connection to a JUNOS software with enhanced services J-series Services Router providing IPSec VPN services at a branch office are considered end-to-site tunnels, or remote access connections.

- When a user working at home or traveling connects to an Internet Service Provider (ISP), they are assigned an IP address dynamically.
- When a user launches client software to begin an IPSec VPN session, the software connects to the gateway at the site, and the site gateway software prompts the user for authentication information.
- If the user satisfies all of the authentication requirements, a remote access connection across an IPSec VPN is initiated, negotiated, and established.

Using the JUNOS software with enhanced services, you can configure a local gateway to support remote access IPSec VPN connections on a per-user or group basis.

Tasks for Implementing a Dynamic Endpoint Site-to-Site IPSec VPN

Managers at the Los Angeles site require access to servers at the San Antonio branch office. For this purpose, the network administrator plans to configure a remote access IPSec VPN.

The IPSec VPN deployment for this scenario is premised on the following assumptions:

- IP addresses at the Los Angeles headquarters site are dynamically assigned to hosts.
- The outgoing interface on the host that is used for the IPSec VPN gateway at the Los Angeles headquarters has a dynamically assigned IP address.
- The outgoing interface on the JUNOS software with enhanced services J-series Services Router at the San Antonio branch site that is used for the IPSec VPN has a static IP address.

To create an IPSec VPN to respond to tunnel setup (IKE) negotiations initiated by a peer gateway with a dynamically assigned IP address, you complete the tasks shown in Table 99 and Table 100.

Table 99: Task for Configuring a Dynamic Endpoint IPSec VPN

Task	Instruction
1. Configure the required security zones and assign interfaces to them.	See "Implementing Firewall Deployments for Branch Offices" on page 37.

Table 100: Procedure for Configuring a Dynamic Endpoint IPSec VPN

Task	Instructions
2. IKE Phase 1: Configure the peer gateway information, including a dynamic address specification.	See "Configuring the Peer Gateway for a Dynamic Endpoint IPSec VPN" on page 174. ■ "Configuring the Peer Gateway Name and Destination Address" on page 174. ■ "Configuring the IKE Phase 1 Policy Reference" on page 175. ■ "Configuring the Local Peer's Outgoing Interface" on page 175.
3. IKE Phase 1: Configure a Phase 1 proposal.	See "Configuring IKE Phase 1 Proposals" on page 177.
4. IKE Phase 1: Configure an IKE policy.	See "Configuring IKE Phase 1 Policies" on page 176.
5. IPSec Phase 2: Configure the IPSec VPN name and gateway information.	See "Configuring the IPSec VPN" on page 180.
6. IPSec Phase 2: Configure a Phase 2 proposal.	See "Configuring IPSec Phase 2 Proposals" on page 182.
7. IPSec Phase 2: Configure a policy.	See "Configuring IPSec Phase 2 Policies" on page 185.
8. Configure firewall security policies. For this scenario, configure firewall policies to give the managers at the Los Angeles site access to the server at the San Antonio branch.	See "Implementing Firewall Deployments for Branch Offices" on page 37. For this scenario, see also the example policy statement following the IPSec VPN configuration in "IPSec VPN Statements Summary" on page 171.

Configuring a Dynamic Endpoint Site-to-Site IPSec VPN

This section describes how to configure a dynamic endpoint site-to-site IPSec VPN for the San Antonio peer. The local interface used as the outgoing interface for the IPSec VPN has a static IP address.

IPSec VPN Statements Summary

This section shows the CLI commands used to configure a policy-based dynamic endpoint IPSec VPN for the San Antonio branch office peer responding to its Los Angeles dynamic endpoint peer. The complete configuration for the VPN is shown after the set of CLI statements.

After the complete configuration, a policy statement is shown that permits traffic from any of the users at hosts behind the Los Angeles peer gateway to a specific server at the San Antonio branch office through use of a specific set of applications.

CAUTION: To use this sample configuration, you must commit the configuration after all the statements are entered. If you intend to commit the statements individually, you must reverse the order of statements within each IKE phase because of dependencies among the statements.

Configuring Dynamic Endpoint IPSec VPN: IKE Phase 1

Dynamic endpoint peer gateway name and dynamic address. (See "Configuring the Peer Gateway Name and Destination Address" on page 174.)	user@host > set security ike gateway los-angeles-remote-peer dynamic hostname corp2.newbank.com
Phase 1 IKE policy reference. (See "Configuring the IKE Phase 1 Policy Reference" on page 175.)	user@host> set security ike gateway los-angeles-remote-peer ike-policy pol1
Phase 1 gateway external interface. (See "Configuring the Local Peer's Outgoing Interface" on page 175.)	user@host# set security ike gateway los-angeles-remote-peer external-interface ge-0/0/3.2
Phase 1 IKE policy. (See "Configuring IKE Phase 1 Policies" on page 176.)	user@host > set security ike policy pol1 pre-shared-key ascii-text 1234pskey5678proposals prop1

Phase 1 IKE proposals. (See “Configuring IKE Phase 1 Proposals” on page 177.)

```
user@host > set security ike proposal prop1 authentication-method
    pre-shared-keys dh-group group2 authentication-algorithm sha1
    encryption-algorithm 3des-cbc
```

Configuring the Dynamic Endpoint IPSec VPN: Phase 2

IPSec VPN and peer gateway. (See “Configuring the IPSec VPN” on page 180.)

```
user@host > set security ipsec vpn LA-to-SA-Remote ike gateway
    los-angeles-remote-peer
```

IPSec Phase 2 policy reference. (See “Configuring the Reference to the IPSec Phase 2 Policy” on page 181.)

```
user@host > set security ipsec vpn LA-to-SA-Remote ike gateway
    los-angeles-remote-peer ipsec-policy ipsec-pol1
```

IPSec Phase 2 proposal. (See “Configuring IPSec Phase 2 Proposals” on page 182.)

```
user@host > set security ipsec proposal prop1 protocol esp
    authentication-algorithmhmac-md5-96 encryption-algorithm 3des-cbc
```

IPSec Phase 2 policy, proposal reference. (See “Configuring a Reference to the IPSec Proposals” on page 185.)

```
user@host > set security ipsec policy ipsec-pol1 proposals prop1
```

IPSec Phase2 policy, perfect forward secrecy. (See “Configuring Perfect Forward Secrecy” on page 186.)

```
user@host > set security ipsec policy ipsec-pol1
    perfect-forward-secrecy keys group2
```

Dynamic Endpoint IPSec VPN Configuration

The following example shows the same deployment presented as a configuration excerpt.

```
security {
        ike {
            gateway los-angeles-remote-peer {
                dynamic hostname corp2.newbank.com
                ike-policy pol1;
                external-interface ge-0/0/3.1;
            }
            policy pol1 {
                pre-shared-key 1234pskey5678;
                proposals prop1;
            }
            proposal prop1 {
                authentication-method pre-shared-keys;
                dh-group group 2;
                authentication-algorithm sha1;
                encryption-algorithm 3des-cbc;
            }
```

```
            }
            ipsec
                vpn LA-to-SA-Remote {
                    ike {
                        los-angeles-remote-peer;
                        ipsec-policy ipsec-pol1;
                        }
                    }
                    proposal prop1 {
                        protocol esp;
                        authentication-algorithm hmac-md5-96;
                        encryption-algorithm 3des-cbc;
                    }
                    policy ipsec-pol1 {
                        proposal prop1;
                        perfect-forward-secrecy keys group2;
                }
            }
        }
```

Policy Configuration for the LA-to-SA-Remote Peer

According to the parameters of the following policy statement, any of the bank managers at the Los Angeles site (source-address any) can initiate a session with a specific server (destination-address 10.11.14.0/24) at the San Antonio site using the applications included in the application set BankAppSet (application-set BankAppSet). Their data is transmitted securely across the IPSec VPN established between the peer gateways (ipsec-vpn LA-to-SA-Remote).

The server with the IP address 10.11.14.0 that the bank managers at the Los Angeles site wish to access is in the BankManagersSA zone. The IntranetSA zone encompasses the hosts at corporate headquarters accessed via the IPSec VPN.

```
security {
        policies from-zone IntranetSA to-zone BankManagersSA {
            policy BankManIn {
                match {
                    source-address {
                        any;
                    } destination-address {
                        10.11.14.0/24;
                    } application-set {
                        BankAppSet;
                    }
                }
                then {
                    permit {
                        tunnel {
                            ipsec-vpn LA-to-SA-Remote;
                        }
                    }
                }
            }
        }
    }
```

Configuring the Peer Gateway for a Dynamic Endpoint IPSec VPN

As part of Phase 1, you configure a connection between the local and remote peers. For a dynamic endpoint IPSec VPN, you must identify the remote peer using information that is constant. This is the only respect in which remote access IPSec VPNs differ from other site-to-site IPSec VPNs.

Configuring the Peer Gateway Name and Destination Address

```
security-–> ike-–>gateway
            <destination-peer-gateway-name>
                dynamic (hostname <fqdn> |
                    user-at-home <user-fqdn>
                    distinguished-name container <cn> wildcard <wc>|
                    inet <ipv4-addr-1>)
```

TIP: The destination peer gateway name specified in Phase 1 must match the destination peer gateway name specified in Phase 2.

Table 101 defines the statement parameters used for dynamic endpoint IPSec VPNs.

Table 101: IPSec VPN Peer Gateway Name and Address Configuration Statement

Destination Peer Gateway Identity	Syntax
Destination peer gateway name and address	`set security ike gateway <destination-peer-gateway-name> dynamic (hostname <fqdn> | user-at-home <user-fqdn> | distinguished-name container <cn> wild-card <wc> | inet <ipv4-addr-1>)` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `(dynamic)`: ■ `(hostname) fqdn`: Reference to the destination peer gateway, specified as a fully qualified domain name. ■ `(user-at-home) user-fqdn:` e-mail address ■ (`distinguished-name container`) cn (wildcard) wc: Software uses the DN from the certificate. ■ `(inet) ipv4-addr-1`

Configuring the IKE Phase 1 Policy Reference

```
security-–> ike-–>gateway
            <destination-peer-gateway-name>
                 dynamic...
                ike-policy <ike-policy-name>
```

You configure a reference to the name of the IKE policy to be used for the destination peer gateway. Table 102 defines the statement parameters.

Table 102: IKE Phase 1 Policy Reference Configuration Statement

Destination Peer Policy Reference	Syntax
Refers to the IKE policy to be used for communication with the peer gateway. In turn, the policy refers to the proposal containing cryptography parameters to present to the peer for Phase 1 negotiation.	`set security ike gateway <destination-peer-gateway-name> dynamic... ike-policy <ike-policy-name>` Note: For a definition of the `destination-peer-gateway-name` parameter, see Table 101. ■ `ike-policy-name`: Name of the policy specified as an alphanumeric string.

Configuring the Local Peer's Outgoing Interface

```
security–-> ike-–>gateway
            <destination-peer-gateway-name>
                dynamic....
                    external-interface <interface-name>
```

Table 103 defines the syntax for the external interface statement.

Table 103: External Interface Name Configuration Statement

External Interface Name	Syntax
Gateway's outgoing interface.	`set security ike gateway <destination-peer-gateway-name> external-interface <interface-name>` For a definition of the `destination-peer-gateway-name` parameter, see Table 101. ■ `interface-name`: Name of the outgoing interface to be used for the VPN.

Configuring IKE Phase 1 Policies

Configuring the IKE Policy Name

```
security --> ike -->
        policy <policy-name>
```

Table 104 defines the policy name parameter. You can specify the policy name as part of the entire policy statement. This section shows the specification alone for clarity.

Table 104: IKE Policy Name Configuration Statement

IKE Policy Name	Syntax
Name of an ike policy. A peer gateway configuration refers to the policy to be used for it in its ike-policy statement.	`set security ike policy <policy-name>` ■ `policy-name`: An alphanumeric string identifying the policy.

Configuring Phase 1 IKE Policy Authentication Parameters

```
security --> ike -->
            policy <policy-name>
                    (pre-shared-keys <key> | certificate <values>)
```

Table 105 defines the authentication method statement syntax.

Table 105: Peer Authentication Type Configuration Statement

Peer Authentication Type and Parameters	Syntax

The type of authentication to be used to secure the tunnel for Phase 2. Specifies the method and value.
 	
For details about certificate configuration, see the *JUNOS Software Security Configuration Guide*.
	`set security ike policy <policy-name> (pre-shared-key <key> | certificate <certificate-information>)`
	■ `policy-name`: ASCII text string.
	■ If pre-shared-key `key`: Unreadable string
	■ If certificate `certificate-information`: Certificate configuration parameters
	  ■ `(local-certificate) certificate ID`
	  ■ `(trusted-ca) CA-ID | use-all`
	  ■ `(peer-certificate-type) pkcs7 | x509-signature`

Configuring a Reference to the IKE Proposal

```
security --> ike -->
            policy <policy-name>
                        proposals <proposal-name>
```

Table 106 defines the syntax.

Table 106: IKE Phase 1 Proposal Configuration Statement

IKE Phase 1 Proposal Reference	Syntax
A reference to one or more IKE Phase 1 proposals whose parameters are to be negotiated with the destination peer. The IKE policy for the peer gateway refers to the proposals by name.	`set security ike policy <policy-name> proposals [<proposal1> <proposal2>...] proposal-set <proposal-set-name>` ■ `policy-name`: ASCII text string. ■ `proposal1, proposal2`: A reference to the name(s) of one or more proposals to be used for negotiation with the peer. ■ `proposal-set-name:` A reference to a predefined proposal set to be used for negotiation with the peer.
Predefined Phase 1 proposals sets: ■ basic: ■ Proposal 1: pre-shared key, g1, des, sha1 ■ Proposal 2: pre-shared key, g1, des, md5 ■ compatible: ■ Proposal 1: pre-shared key, g2, 3des, sha1 ■ Proposal 2: pre-shared key, g2, 3des, md5 ■ Proposal 3: pre-shared key, g2, des, sha1 ■ Proposal 4: pre-shared key, g2, des, md5 ■ standard: ■ Proposal 1: pre-shared key, g2, 3des, sha1 Proposal 2: pre-shared key, g2, aes128, sha1	

Configuring IKE Phase 1 Proposals

During Phase 1 of the IKE negotiations, the local peer proposes to the destination peer the cryptography to be used to secure the IPSec Phase 2 tunnel. The Phase 2 tunnel is not used for transit data. Rather, it is used to provide a secure means for the two peers to negotiate the values to be used to secure transit data.

Configuring the IKE Proposal Name

```
security--> ike--> proposal <proposal-name>
```

Table 107 gives the proposal name statement syntax.

Table 107: IKE Proposal Name Configuration Statement

IKE Proposal Name	Syntax
The name, or label, of the IKE Phase 1 proposal that follows. The gateway policy statement refers to the proposal name. The policy statement can refer to one or more proposals to be used.	`set security ike proposal <proposal-name>` ■ `proposal-name`: Name of the proposal.

Configuring the IKE Proposal Authentication Method

```
security --> ike --> proposal <proposal-name>
                     authentication-method
                        (pre-shared-keys | RSA signatures)
```

Table 108 shows the IKE proposal authentication method statement syntax.

Table 108: IKE Proposal Authentication Method Configuration Statement

IKE Proposal Authentication Method	Syntax
The kind of authentication method to use for the peers to identify themselves to each other. ■ For a preshared key, the *policy* gives the key. The *proposal* states that preshared keys are to be used. ■ For a certificate. the policy certificate parameter identifies which specific certificate is to be used. The proposal tells the kind of certificate to use.	`set security ike proposal <proposal-name> authentication-method <method>` ■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations. ■ `(authentication-method) method`: Specifies the authentication method, which corresponds to the policy specification. Specify one of the following keywords: ■ `pre-shared-keys` ■ `rsa-signatures`

Configuring the IKE DH Group

```
security --> ike --> proposal <proposal-name>
                     dh-group (group1 | group2 | group5)
```

Table 109 gives an overview of use of the Diffie-Hellman group parameter and its syntax.

Table 109: IKE DH Group Configuration Statement

IKE Proposal DH Group	

DH group to use to increase the security of the tunnel. A DH exchange allows the peers to produce a shared secret over an unsecured medium without passing the secret through the wire.
The size of the prime modulus used in each group differs and determines the strength of its security.
	`set security ike proposal <proposal-name> dh-group (group1 | group2 | group5)`
	■ `proposal-name`: Name of the proposal to use for the IKE Phase 1 negotiations.
	■ `(dh-group) option`: Specify one of the following options to indicate the Diffie-Hellman group to be used.
	  ■ group1 (768-bit modulus)
	  ■ group2 (1024-bit modulus)
	  ■ group5 (1536-bit modulus)
	Because the modulus for each DH group is a different size, the peers must agree to use the same group.

Configuring the IKE Proposal Encryption Algorithm

```
security --> ike --> proposal <proposal-name>
                 encryption-algorithm (des-cbc |3des-cbc | aes-128-cbc |
                                       aes-192-cbc| aes-256-cbc)
```

Table 110 defines the encryption algorithm configuration.

Table 110: IKE Encryption Algorithm Configuration Statement

IKE Proposal Encryption Algorithm	Syntax

Encryption algorithm specified within a proposal to be presented to the peer for negotiation. The agreed-upon algorithm is used to secure exchanges between the peers during IKE Phase 2.	`set security ike proposal <proposal-name> encryption-algorithm (des-cbc | 3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc)`
	■ `proposal-name`: Name of the proposal, which is referred to in the policy statement to specify use of it.
	■ `(encryption-algorithm)`: Specify one of the five encryption algorithms to use in the proposal. In order from strongest to weakest, the options are:
	  ■ `aes-256-cbc`
	  ■ `aes-192-cbc`
	  ■ `aes-128-cbc`
	  ■ `3des-cbc`
	  ■ `des-cbc`

Configuring the IKE Authentication Algorithm

```
security --> ike --> proposal <proposal-name>
                        authentication-algorithm (md5 | sha1)
```

Table 111 defines the IKE proposal authentication algorithm.

Table 111: IKE Proposal Authentication Algorithm Configuration Statement

IKE Proposal Authentication Algorithm	Syntax
Authentication algorithm to propose to the peer to be used to authenticate exchanges between the peers during IKE Phase 2. A proposal contains the authentication algorithm, and the peer must accept the entire proposal.	`set security ike proposal <proposal-name> authentication-algorithm (md5 | sha1)` ■ `proposal-name`: Name of the proposal to be negotiated during IKE Phase 1. `(authentication-algorithm)`: Specify one of the two authentication algorithms to use in the proposal. Options are: SHA1 and MD5.

Configuring the IPSec VPN

After the peer gateway has been defined and the IKE Phase 1 proposals have been configured, you configure the IPSec VPN. The IPSec VPN gateway configuration brings together the peer gateway configuration, the IPSec proposal, and the tunnel interface to be used for the IPSec VPN.

Naming the IPSec VPN Gateway

```
    security -—> ipsec --> vpn
            vpn <vpn-name>
```

Table 112 defines the IPSec VPN name syntax.

Table 112: IPSec VPN Name Configuration Statement

IPSec VPN Name	Syntax
Name of the IPSec VPN being configured for the local gateway.	Name of the IPSec VPN being configured for the local gateway.

Identifying the Peer Gateway

```
security --> ipsec --> vpn
            vpn <vpn-name>
                ike
                    gateway <destination-peer-gateway-name>
```

Table 113 defines the peer gateway name parameter for Phase 2.

Table 113: IPSec Phase 2 Destination Peer Gateway Name Configuration Statement

IPSec Phase 2 Destination Peer Gateway Name	Syntax
Name of the destination peer whose gateway you are configuring.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name>` ■ `vpn-name`: IPSec VPN name. ■ `destination-peer-gateway-name`: Name of the destination peer. This name must be the same as the one you specified in the IKE Phase 1 configuration.

Configuring the Reference to the IPSec Phase 2 Policy

```
security --> ipsec --> vpn <vpn-name>
            ike
                ipsec-policy <ipsec-policy-name>
```

Table 114 gives the syntax for Phase 2 policy reference parameter.

Table 114: IPSec Policy Reference Configuration Statement Syntax

Reference to the IPSec Policy to Use	Syntax
Refers to the IPSec policy to be used for communication with the specified destination peer gateway.	`set security ipsec vpn <vpn-name> ike gateway <destination-peer-gateway-name> ipsec-policy <ipsec-policy-name>` ■ `vpn-name`: A string identifying the IPSec VPN created for the local gateway. ■ `destination-peer-gateway-name`: The name of the destination peer. This is the same name for the destination gateway as specified in the IKE Phase 1 configuration. ■ `ipsec-policy-name`: Name of the policy to be used for the peer gateway.

Configuring the Gateway to Set Up an IPSec VPN When Traffic Triggers It

```
security --> ipsec --> vpn <vpn-name>
            establish-tunnels (immediately | on-traffic)
```

This feature enables you to specify whether IKE is triggered immediately after the IPSec VPN configuration is committed, or only when data traffic passes and IKE must be negotiated with the peer gateway. This scenario uses the default value of `on-traffic`.

CAUTION: For dynamic endpoint IPSec VPNs and remote access connections, you must use on-traffic, which is the default value. This section is shown for completeness, not because you would need to change the default value.

Table 115 gives the establish-tunnels statement syntax.

Table 115: When to Establish Tunnels Parameter

Establish Tunnels Parameter	Syntax
Refers to the IPSec policy to be used for communication with the specified destination peer gateway.	`set security ipsec vpn <vpn-name> establish-tunnels immediately | on-traffic` ■ `vpn-name`: A string identifying the IPSec VPN to be created for the local gateway. ■ `destination-peer-gateway`-name: The name of the destination peer. This is the same name for the destination gateway as specified in the IKE Phase 1 configuration. `(establish-tunnels)`. Either of the following keywords: ■ `immediately`: The IPSec VPN tunnel is set up immediately after the configuration is committed. In this case, IKE negotiations begin immediately. ■ `on-traffic`: The IPSec VPN tunnel is set up when transit traffic triggers it.

Configuring IPSec Phase 2 Proposals

During IPSec Phase 2, the peers negotiate cryptographic parameters that determine how data traffic that traverses the IPSec VPN is to be secured. These parameters identify the kind of authentication and encryption to be used. They are specified in a proposal to which the gateway policy refers by name.

Configuring the Phase 2 IPSec Proposal Name

```
security -> ipsec
            proposal <proposal-name>
```

Table 116 gives the IPSec proposal name syntax.

Table 116: IPSec Proposal Name Configuration Statement

IPSec Proposal Name	Syntax
Name referred to in an IPSec policy to specify which proposal is to be used.	`set security ipsec proposal <proposal-name>` ■ `proposal-name`: Name of the IPSec proposal whose protocol you are configuring

Configuring the Phase 2 IPSec Protocol

```
security -> ipsec
            proposal <proposal-name>
                protocol (esp | ah)
```

Table 117 gives an overview of the IPSec protocol syntax.

Table 117: IPSec Protocol Definition Configuration Statement

IPSec Protocols	Syntax
One of two IPSec protocols used to authenticate and encrypt data traffic traversing the IPSec VPN.	`set security ipsec proposal <proposal-name> protocol (esp | ah)` ■ `proposal-name`: Name of the IPSec proposal whose protocol you are configuring ■ `(protocol)`: Either of the following protocols: ■ `esp`: Encapsulating Security Payload ■ `ah`: Authentication Header
■ `ESP` provides the following services: ■ Encrypts data, rendering it undecipherable. ■ Provides data origin authentication ■ Provides data integrity ■ Supports antireplay protection What it doesn't provide: ■ It authenticates and encrypts only the data portion of the packet, excluding the IP header.	
■ `AH` provides the following services: ■ Provides data origin authentication ■ Supports antireplay protection for most of the packet What it doesn't provide: ■ Antireplay protection of the fields in the IP header which can change during transit. ■ Encryption ■ Confidentiality	

Configuring the IPSec Proposal Authentication Algorithm

```
security --> ipsec
            proposal <proposal-name>
                    authentication algorithm (hmac-md5-96 | hmac-sha1-96)
```

Table 118 defines the authentication algorithm.

Table 118: IPSec Proposal Authentication Algorithm Configuration Statement

IPSec Proposal Authentication Algorithm	Syntax

Authentication algorithm to be included in the specified proposal.	`set security ipsec proposal <proposal-name> authentication-algorithm (hmac-md5-95 | hmac-sha1-96)`
	■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal whose authentication algorithm you are configuring.
	■ `(authentication-algorithm)`: One of the following keywords:
	  ■ `hmac-md5-95`
	  ■ `hmac-sha1-96`

Configuring the IPSec Proposal Encryption Algorithm for ESP

```
security --> ipsec
            proposal <proposal-name>
                    encryption algorithm (3des-cbc| aes-128-cbc | aes-192-cbc |
                                          aes-256-cbc | des-cbc)
```

Table 119 defines the encryption algorithm.

Table 119: IPSec Proposal Encryption Algorithm Configuration Statement

IPSec Proposal Encryption Algorithm	Syntax
Identifies the encryption algorithm to be used for encryption via ESP.	`set security ipsec proposal <proposal-name> encryption-algorithm (3des-cbc| aes-128-cbc | aes-192-cbc |aes-256-cbc | des-cbc)` ■ `proposal-name`: Name of the IPSec VPN Phase 2 proposal whose encryption algorithm you are configuring. ■ `(encryption-algorithm)`: Specified as a keyword.

Configuring IPSec Phase 2 Policies

Configuring the IPSec Policy Name

```
security --> ipsec
            policy <policy-name>
```

Table 120 defines the policy name parameter.

Table 120: IKE Policy Name Configuration Statement

IKE Policy Name	Syntax
Name of an IPSec policy. The configuration refers to the policy to be used by its name.	`set security ipsec policy <policy-name>` ■ `policy-name`: An alphanumeric string identifying the policy.

Configuring a Reference to the IPSec Proposals

```
security—ipsec
            policy <policy-name>
                proposals <proposal-name>| proposal-set <proposal-set-name>
```

Table 121 defines the IPSec proposal reference syntax.

Table 121: IPSec Phase 2 Proposal Reference Configuration Statement

IKE Phase 2 Proposal Reference	Syntax
A reference to one or more IPSec Phase 2 proposals or proposal sets to be used to negotiate security for the tunnel.	`set security ipsec policy <policy-name> proposal [<proposal1> <proposal2>...] | proposal-set <proposal-set-name>` ■ `policy-name`: ASCII text string. ■ `proposal1, proposal2`: A reference to the name(s) of one or more proposals to be used for negotiation with the peer. ■ `proposal-set-name`: A reference to a predefined proposal set to be used for negotiation with the peer.
Predefined Phase 2 proposals sets: Basic: Proposal 1: esp-no pfs, des, sha1 Proposal 2: esp-no pfs, des md5 Compatible: Proposal 1: esp-no pfs, 3des, sha1 Proposal 2: esp-no pfs, 3des, md5 Proposal 3: esp-no pfs, des, sha1 Proposal 4: no pfs, des, md5 Standard: Proposal 1: esp, g2 (pfs), 3des, sha1 Proposal 2: esp, g2, aes128, sha1	

Configuring Perfect Forward Secrecy

```
security ipsec
            policy <policy-name>
                    perfect-forward-secrecy keys <Diffie-Hellman-group>
```

Perfect Forward Secrecy (PFS) is a method for deriving Phase 2 keys independent of and unrelated to keys used for previously generated IPSec SAs. In IKE Phase 2, if PFS is configured, a second DH exchange for each IPSec SA is performed. As a consequence, keying material for each IPSec SA is completely independent of every other SA, including its IKE SA.

If you configure PFS, you must specify the Diffie-Hellman group to be used.

- Diffie-Hellman groups include:
  - DH Group 1: 768-bit modulus
  - DH Group 2: 1024-bit modulus
  - DH Group 5: 1536-bit modulus

Because the modulus for each DH group is a different size, the peers must agree to use the same group.

About IPSec VPNs for Remote Access Connections

The New Bank of the Southwest provides support for its employees to work remotely, either at home or when they travel. When a user launches his IPSec VPN client software to begin a remote session, the client software attempts to connect to the remote peer gateway at the branch or corporate site.

The client software can reach the peer gateway because the peer has a static IP address. For employees to connect remotely to the branch or corporate site, you must configure an IPSec VPN to support the remote connection.

The user seeking a remote access IPSec connection must be able to identify and authenticate herself or himself to the gateway before IKE negotiations can begin. For this purpose, you must configure a user profile.

You can configure profiles on a per-user basis or a group basis. If a corporation has many employees, configuring connections for individual users can be tedious and error prone. To address this problem, IPSec provides two methods you can use to configure group user identities: shared IKE ID and group IKE ID. They are explained in the following sections.

Shared IKE ID Authentication

Shared IKE ID authentication is a method of configuring a group user profile to be used by individual users as a means of authenticating themselves to a site's gateway when they attempt to begin an IPSec VPN session.

If you configure the IPSec VPN gateway to use the shared IKE ID method, a user who wants to establish a remote access connection must first satisfy the shared IKE ID profile identification requirements. Shared IKE ID authentication requires use of eXtended Authentication (XAuth). Therefore, you must specify use of it in the configuration.

The following conditions must exist for the user to qualify to use the profile:

- The user must have an authorized pre-shared key.
- The user must have a shared user identity in the form of an E-mail address

An administrator communicates this information to remote access users.

If the user meets the requirements, IKE negotiations begin.

With shared IKE ID authentication, all users requiring remote access connections share the same key and E-mail address. Because this method relies on communication of these values to users, it is considered insecure and prone to illegal entry. Therefore, a second method of identification is required from the user after the IKE Phase 1 authentication completes and before the IKE Phase 2 process begins.

For this purpose, the shared IKE ID method uses the XAuth protocol as the additional means of authentication.

If the IKE process completes successfully and the user passes the XAuth authentication process, the IPSec VPN tunnel is negotiated and established. Users can then access the resources and services they desire to use, if the security policy allows it.

When you configure the XAuth access profile, you must specify RADIUS (radius-server) as the authentication server type and it must be specified using the authentication-order statement in the XAuth profile.

NOTE: If the RADIUS server supports use of an IP address as part of a user entry, the JUNOS software with enhanced services will assign the address to the user as a virtual IP address.

Group IKE ID Authentication

Group IKE ID authentication is a method of configuring a group user profile to be used by individual users as a means of authenticating themselves to the gateway so that IKE negotiations can begin.

If you configure the IPSec VPN to use the group IKE ID method, a user who wants to establish a remote access connection must first satisfy the group IKE ID profile identification requirements before IKE negotiations can begin.

The following conditions must exist for the user to qualify to use the group profile:

- The user must have a pre-shared key seed (The pre-shared key is based on the pre-shared key seed value specified at the VPN gateway and the full identity of the remote access user.)
- The rightmost part of the user's identification must match the group IKE ID. (You can use e-mail address or FQDN for this purpose.)

Tasks for Implementing Remote Access Connection IPSec VPNs Using Shared IKE ID

Because shared IKE ID facilitates the deployment of a large number of remote access users, it is the method that the New Bank of the Southwest uses.

To create an IPSec VPN to allow for dynamic remote access connections for 20 users using Shared IKE ID, you complete the tasks shown in Table 122 and Table 123.

Table 122: Preliminary Procedure for Creating an IPSec VPN for Remote Access Connections

Task	Reference
1. Configure the required security zones and assign their interfaces to them.	See "Implementing Firewall Deployments for Branch Offices" on page 37.
2. Configure firewall policies.	See "Implementing Firewall Deployments for Branch Offices" on page 37.

Table 123: Procedure for Creating an IPSec VPN for Remote Access Connections

Task	Reference
1. IKE Phase 1: Configure the peer gateway information.	See "Configuring the Peer Gateway" on page 192. ■ "Configuring the Peer Gateway Name and Address" on page 192 ■ "Configuring the Group Profile IKE ID Type for Remote Access Users" on page 193 ■ "Configuring the Number of User Connections" on page 194 ■ "Configuring the Reference to the Access Profile to Be Used for XAuth Authentication" on page 195 ■ "Configuring the IKE Phase 1 Policy Reference" on page 196 ■ "Configuring the Local Gateway's Outgoing Interface" on page 196
2. IKE Phase 1: Configure a policy.	See "Configuring IKE Phase 1 Policies" on page 176.
3. IKE Phase 1: Configure a proposal.	See "Configuring IKE Phase 1 Proposals" on page 177.
4. IPSec Phase 2: Configure the IPSec VPN.	See "Configuring the IPSec VPN" on page 180.
5. IPSec Phase 2: Configure a proposal.	See "Configuring IPSec Phase 2 Proposals" on page 182.
6. IPSec Phase 2: Configure a policy.	See "Configuring IPSec Phase 2 Policies" on page 185.

IPSec VPN Statements Description and Configuration Using Shared IKE ID

This section shows the CLI commands used to configure the IPSec VPN to support remote access connections using the shared IKE ID group profile with XAuth authentication. Then it shows the complete configuration for the IPSec VPN.

After the complete configuration, it shows the server specification for the access profile referred to in the configuration. The access profile specifies the server to be used for XAuth user authentication. You must use a RADIUS server for XAuth user authentication.

NOTE: In the sections that follow the configuration, parts that are common to this configuration and the one discussed in "Configuring a Dynamic Endpoint Site-to-Site IPSec VPN" on page 171 are included, but they do not contain the syntax definition for the statement. Instead, they refer to the syntax definition provided in the corresponding section of the dynamic endpoint site-to-site IPSec VPN scenario.

Remote Access Connections IPSec VPN Configuration Using the CLI

Configuring IKE Phase 1

Peer gateway name, address. (See "Configuring the Peer Gateway Name and Address" on page 192.)

```
user@host# set security ike gateway san-antonio-remote-users
        dynamic user-at-home SanAnRemoteUsers@newbankSW.com
```

Phase 1: IKE ID type. (See "Configuring the Group Profile IKE ID Type for Remote Access Users" on page 193.)

```
user@host# set security ike gateway san-antonio-remote-users
    dynamic user-at-home SanAnRemoteUsers@newbankSW.com ike-user-type
    shared-ike-id
```

Phase 1: Maximum number of remote access connections. (See "Configuring the Number of User Connections" on page 194.)

```
user@host# set security ike gateway san-antonio-remote-users dynamic
    user-at-home SanAnRemoteUsers@newbankSW.com connections-limit 20
```

Phase 1: XAuth authentication access profile reference. (See "Configuring the Reference to the Access Profile to Be Used for XAuth Authentication" on page 195.)

```
user@host# set security ike gateway san-antonio-remote-users
    xauth access-profile SA-remote-users
```

Phase 1: IKE policy reference. (See "Configuring the IKE Phase 1 Policy Reference" on page 175.)

```
user@host# set security ike gateway san-antonio-remote-users ike-policy pol2
```

Phase 1: gateway external interface. (See "Configuring the Local Peer's Outgoing Interface" on page 175.)

```
user@host# set security ike gateway san-antonio-remote-users
    external-interface ge-0/0/3.2
```

Phase 1 IKE policy. ("Configuring IKE Phase 1 Policies" on page 176)

```
user@host# set security ike policy pol2 pre-shared-key
    ascii-text 1234pskey5678
```

Phase 1 IKE proposal reference. ("Configuring a Reference to the IKE Proposal" on page 177.)

```
user@host# set security ike proposals prop2
```

Phase 1 IKE proposals. (“Configuring IKE Phase 1 Proposals” on page 177.)

```
user@host# set security ike proposal prop2 authentication-method
    pre-shared-keys authentication-algorithm sha1
    encryption-algorithm 3des-cbc
```

Configuring the Phase 2 IPSec VPN

Phase 2: IPSec VPN and peer gateway. (“Configuring the IPSec VPN” on page 180.)

```
user@host# set security ipsec vpn SA-RemoteAccessUsers ike gateway
    san-antonio-remote-users
```

Phase 2 IPSec policy reference. (“Configuring the Reference to the IPSec Phase 2 Policy” on page 181.)

```
user@host# set security ipsec vpn SA-RemoteAccessUsers ike gateway
    san-antonio-remote-users ipsec-policy ipsec-pol2
```

Phase 2: Establish tunnel when transit traffic requires it. (See“Configuring the Gateway to Set Up an IPSec VPN When Traffic Triggers It” on page 181.)

```
user@host# set security ipsec vpn SA-RemoteAccessUsers establish-tunnels
on-traffic
```

IPSec Phase 2 proposals. (“Configuring IPSec Phase 2 Proposals” on page 182.)

```
user@host# set security ipsec proposal prop1 protocol esp
    authentication-algorithmhmac-md5-96 encryption-algorithm 3des-cbc
```

IPSec Phase 2 policy, proposal reference. (“Configuring IPSec Phase 2 Policies” on page 185.)

```
user@host# set security ipsec policy ipsec-pol1 proposals prop2
```

The following example shows the same configuration presented as a configuration excerpt.

```
security {
     ike {
            gateway {
                san-antonio-remote-users
                    dynamic
                        user-at-home SanAnRemoteUsers@newbankSW.com;
                        ike-user-type shared-ike-id;
                        connections-limit 20;
                    xauth access-profile SA-remote-users
                    ike-policy pol 2
                    external-interface ge-0/0/3.2
                    }
```

```
                policy {
                    pol2
                    pre-shared-key ascii-text 1234pskey5678
                    proposals prop2
                }
                proposal prop2{
                    authentication-method rsa-signature;
                    dh-group group2;
                    encryption-algorithm aes-256-cbc;
                    authentication-algorithm md5;
                }
            }
            ipsec {
                vpn SA-RemoteAccessUsers {
                ike {
                    gateway
                        san-antonio-remote-users
                        ipsec-policy ipsec-pol2;
                    }
                establish-tunnels on-traffic
                }
            }
            proposal ipsec-prop2 {
                protocol esp;
                authentication-algorithm hmac-md5-96;
                encryption-algorithm aes-256-cbc;
            }
            policy ipsec-pol2 {
                proposal ipsec-prop2;
            }
        }
    }
```

The following example shows the statement for specifying the RADIUS server to be used for XAuth authentication when shared IKE ID is used:

```
user@host# set access profile SA-remote-users radius-server 1.2.3.4
```

Configuring the Peer Gateway

Configuring the Peer Gateway Name and Address

```
security—ike—gateway
            <destination-peer-gateway-name>
                dynamic hostname <fqdn> |
                    user-at-home <user-fqdn>
                    ike-user-type (group-ike-id | shared-ike-id)
```

For a remote access connection, because the user does not have a gateway behind him, the peer gateway address must be specified as dynamic.

Because the user's host is assigned an address dynamically, an identifier other than an IP address must be used. In this case, you may use either the user's hostname or an email address.

NOTE: The destination peer gateway name specified in Phase 1 must match the destination peer gateway name specified in Phase 2.

Table 124 shows the syntax for the peer gateway name and address statement.

Table 124: IPSec VPN Peer Gateway Name and Address Configuration Statement

Destination Peer Gateway Identity	Syntax
Destination peer gateway name and dynamic address	`set security ike gateway <destination-peer-gateway-name> dynamic (hostname <fqdn> | user-at-home <user-fqdn> | distinguished-name container <cn> wildcard | inet <ipv4-addr-1>)` ■ `destination-peer-gateway-name`: Name of the destination peer gateway, specified as an alphanumeric string. ■ `dynamic`: ■ `(hostname) fqdn`: Reference to the remote access peer gateway or remote access user, specified as a fully qualified domain name. ■ `(user-at-home) user-fqdn:` email address ■ `(distinguished-name container) cn wildcard <wc>.` ■ `(inet) ipv4-addr-1.`

Configuring the Group Profile IKE ID Type for Remote Access Users

```
security—ike—gateway
            <destination-peer-gateway-name>
                dynamic...
                    ike-user-type (group-ike-id | shared-ike-id)
```

Before IKE negotiations can begin, a user requiring a remote access IPSec VPN connection must be successfully authenticated by the gateway. For this to occur, you must configure an IKE ID user or group profile. JUNOS software with enhanced services provides the following two kinds of group profiles that you can use:

- Group IKE ID

  This method uses a pre-shared key seed and the rightmost part of the user's identification to determine if the user meets the qualifications of the group profile. If the user meets the requirements, IKE negotiations begin.

- Shared IKE ID

  This method uses a complete pre-shared key and a shared user identity in the form of an email address to determine if the user qualifies. If the user meets the requirements, IKE negotiations begin.

  Because all users share the same key and E-mail address and the administrator communicates the information to them, this method is considered insecure. Therefore, a second method of identification is required from the user after the IKE Phase 1 authentication completes and before the IKE Phase 2 process begins. XAuth authentication is used for this purpose.

If the IKE process completes successfully and the user passes the XAuth authentication process, the IPSec VPN tunnel is established.

Table 125 shows the syntax for the IKE ID statement.

Table 125: IKE ID Group Profile Configuration Statement

IKE ID Group Profile	Syntax
Group profile method to use for IKE user authentication for remote access users. There are two types of profiles: ■ shared IKE ID. Requires: ■ A pre-shared key ■ An email address ■ XAuth authentication ■ group-ike-id. Requires: ■ A pre-shared key seed. The rightmost part of the user identification must match the group IKE ID (email address or FQDN)	`set security ike gateway <destination-peer-gateway-name> dynamic....` For the destination peer gateway name and address, see Table 101. (`ike-user-type`): ■ `group-ike-id`: Use the group IKE ID method. ■ `shared-ike-id`: Use the shared IKE ID method.

Configuring the Number of User Connections

```
security—->ike—->gateway
            <destination-peer-gateway-name>
                dynamic...
                connections-limit <number-connections>
```

You can specify the maximum number of concurrent remote access user connections, after which no more requests for IPSec VPN connections will be honored until other users have disconnected from the site's gateway. Table 126 shows the syntax for the statement.

Table 126: Maximum Number of Remote Access Connections Configuration Statement

Maximum Number of Concurrent Remote Access User Connections	Syntax
Maximum number of concurrent remote access connections	`set security ike gateway <destination-peer-gateway-name> dynamic.... connections-limit <number-connections>` For the destination peer gateway name and address, see Table 101. ■ `number-connections`: Maximum number of connections to allow

Configuring the Reference to the Access Profile to Be Used for XAuth Authentication

```
security—->ike—->gateway
                xauth
                    access-profile <profile name>
```

When shared IKE ID is used, all users share the same key and email address. Because this method relies on communication of these values to users, it is considered insecure. Therefore a second method of identification is required from the user after IKE Phase 1 completes execution and before Phase 2 begins. For this purpose, the shared IKE ID method requires use of the eXtended Authentication (XAuth) protocol as the additional means of authentication.

You must use a RADIUS server for XAuth authentication with shared IKE ID. In the configuration, you specify the name of an XAuth access profile which refers to the RADIUS server to be used. You use the authentication-order statement in the profile for this purpose, as illustrated by the following example:

```
profile xauth {
    authentication-order radius;
    radius-server {
        1.2.1.4 {
            port 1812;
            secret "$5uSOLP1dNr40GkmR$5"
```

If the RADIUS server supports use of an IP address as part of the user's entry, JUNOS software with enhanced services will assign the address to the user as a virtual IP address.

TIP: Juniper Networks provides IPSec client software that you can use for remote access called Netscreen Remote.

Table 127 shows the syntax for the statement to use to specify the XAuth authentication profile reference.

Table 127: XAuth Access Profile Configuration Statement

XAuth Access Profile	Syntax
A reference to the XAuth access profile to use for the shared ID profile.	`set security ike gateway <destination-peer-gateway-name> dynamic.... access-profile <profile-name>` For the destination peer gateway name and address, see Table 101. ■ `profile-name`:

Configuring the IKE Phase 1 Policy Reference

```
security—->ike—->gateway
            <destination-peer-gateway-name>
                 dynamic <hostname <fqdn> |....>
                ike-policy <policy-name>
```

For details, see “Configuring the IKE Phase 1 Policy Reference” on page 175.

Configuring the Local Gateway’s Outgoing Interface

```
security—->ike—->gateway
            <destination-peer-gateway-name>
                dynamic <hostname <fqdn> |....>
                    external-interface <interface-number>
```

For details, see “Configuring the Local Peer’s Outgoing Interface” on page 175.

Configuring IKE Phase 1 Policies

Configuring the IKE Policy Name

```
security—->ike—->
        policy <policy-name>
```

For details, see “Configuring the IKE Policy Name” on page 176.

Configuring Phase 1 IKE Policy Authentication Parameters

```
security—->ike—->
            policy <policy-name>
                    (pre-shared-keys <key> | certificate <values>)
```

For details, see “Configuring Phase 1 IKE Policy Authentication Parameters” on page 176.

Configuring a Reference to the IKE Proposal

```
security—->ike—->
            policy <policy-name>
                            proposals <proposal-name>
```

For details, see “Configuring a Reference to the IKE Proposal” on page 177.

Configuring IKE Phase 1 Proposals

Configuring the IKE Proposal Name

```
security—->ike—->proposal <proposal-name>
```

For details, see “Configuring the IKE Proposal Name” on page 178.

Configuring the IKE Proposal Authentication Method

```
security—->ike—->proposal <proposal-name>authentication-method
                    (pre-shared-keys | RSA signatures)
```

For details, see “Configuring the IKE Proposal Authentication Method” on page 178.

Configuring the IKE DH Group

```
security—->ike—->proposal <proposal-name>
                     dh-group (group1 | group2 | group5)
```

For details, see “Configuring the IKE DH Group” on page 179.

Configuring the IKE Proposal Encryption Algorithm

```
security—->ike—->proposal <proposal-name>
                 encryption-algorithm (des-cbc |3des-cbc | aes-128-cbc |
                                                aes-192-cbc| aes-256-cbc)
```

For details, see “Configuring the IKE Proposal Encryption Algorithm” on page 179.

Configuring the IKE Authentication Algorithm

```
security—->ike—->proposal <proposal-name>
                         authentication-algorithm (md5 | sha1 | sha2)
```

For details see, “Configuring the IKE Authentication Algorithm” on page 180.

Configuring the IPSec VPN Gateway

Naming the IPSec VPN Gateway

```
security—->ipsec—->vpn
        vpn <vpn-name>
```

For details, see “Naming the IPSec VPN Gateway” on page 180.

Identifying the Peer Gateway

```
security—->ipsec —->vpn
        vpn <vpn-name>
            ike
                gateway <destination-peer-gateway-name>
```

For details, see “Identifying the Peer Gateway” on page 180.

Configuring the Reference to the IPSec Phase 2 Policy

```
security ipsec—->vpn <vpn-name>
        ike
            ipsec-policy <ipsec-policy-name>
```

For details, see “Configuring the Reference to the IPSec Phase 2 Policy” on page 181.

Configuring IPSec Phase 2 Proposals

Configuring the Phase 2 IPSec Proposal Name

```
security—->ipsec—->vpn
            proposal <proposal-name>
```

For details, see “Configuring the Phase 2 IPSec Proposal Name” on page 182.

Configuring the Phase 2 IPSec Protocol

```
security —->ipsec
            proposal <proposal-name>
                protocol (esp | ah)
```

For details, see “Configuring the Phase 2 IPSec Protocol” on page 183.

Configuring the IPSec Proposal Authentication Algorithm

```
security—->ipsec
            proposal <proposal-name>
                    authentication algorithm (hmac-md5-96 | hmac-sha1-96)
```

For details, see “Configuring the IPSec Proposal Authentication Algorithm” on page 184.

Configuring the IPSec Proposal Encryption Algorithm for ESP

```
security—->ipsec
            proposal <proposal-name>
                    encryption algorithm (3des-cbc| aes-128-cbc | aes-192-cbc |
                                                aes-256-cbc | des-cbc)
```

For details, see “Configuring the IPSec Proposal Encryption Algorithm for ESP” on page 184.

Configuring IPSec Phase 2 Policies

Configuring the IPSec Policy Name

```
security—->ipsec
            policy <policy-name>
```

For details, see “Configuring the IPSec Policy Name” on page 185.

Configuring a Reference to the IPSec Proposals

```
security—->ipsec
            policy <policy-name>
                proposals <proposal-name>
```

For details, see “Configuring a Reference to the IPSec Proposals” on page 185.

Configuring the Server Specification for the XAuth Access Profile for Shared IKE ID

If you use shared IKE ID as the group profile method, you must configure an access profile to be used for XAuth authentication. In the gateway configuration, you refer to the profile to be used. (See “Configuring the Reference to the Access Profile to Be Used for XAuth Authentication” on page 195.) From the access hierarchy, you configure the profile to refer to the authentication server to be used. You must specify RADIUS as the server type (radius-server).

The following example shows the statement for specifying the RADIUS server to be used for XAuth authentication when shared IKE ID is used:

```
user@host# set access profile <profile-name> radius-server <address>
```

Table 128: XAuth Access Profile Configuration Statement

XAuth Access Profile	Syntax
A reference to the XAuth access profile to use for the shared ID profile.	`set access profile <profile-name> radius-server <address>` ■ `profile-name`: Name of the access profile to be used. `address`: Address of the RADIUS server to be used for XAuth authentication.

Chapter 8

Implementing Network Address Translation for Branch Offices

This chapter explains how to implement Network Address Translation (NAT) for the Albuquerque branch office LAN subnetworks and hosts. The branch uses private IP addresses to conserve public IP addresses and to better protect its LAN's subnetworks and hosts from outside attacks.

This chapter contains the following sections:

About the Branch Office Deployment

The Albuquerque branch of the New Bank of the Southwest is a mid-size branch with two hundred employees. For purposes of illustration, this chapter shows the NAT configurations for a representative group of hosts.

If you are familiar with the Albuquerque branch office topology, including its zones and their members, you can skip this section. In that case, go directly to the NAT configuration beginning with “About Address Translation” on page 202.

Figure 14 shows the topology for the Albuquerque branch office, as defined by the network architect. For descriptions of the security zones, see “Albuquerque Branch Security Zones” on page 40.

Figure 14: Albuquerque Branch Deployment Topology

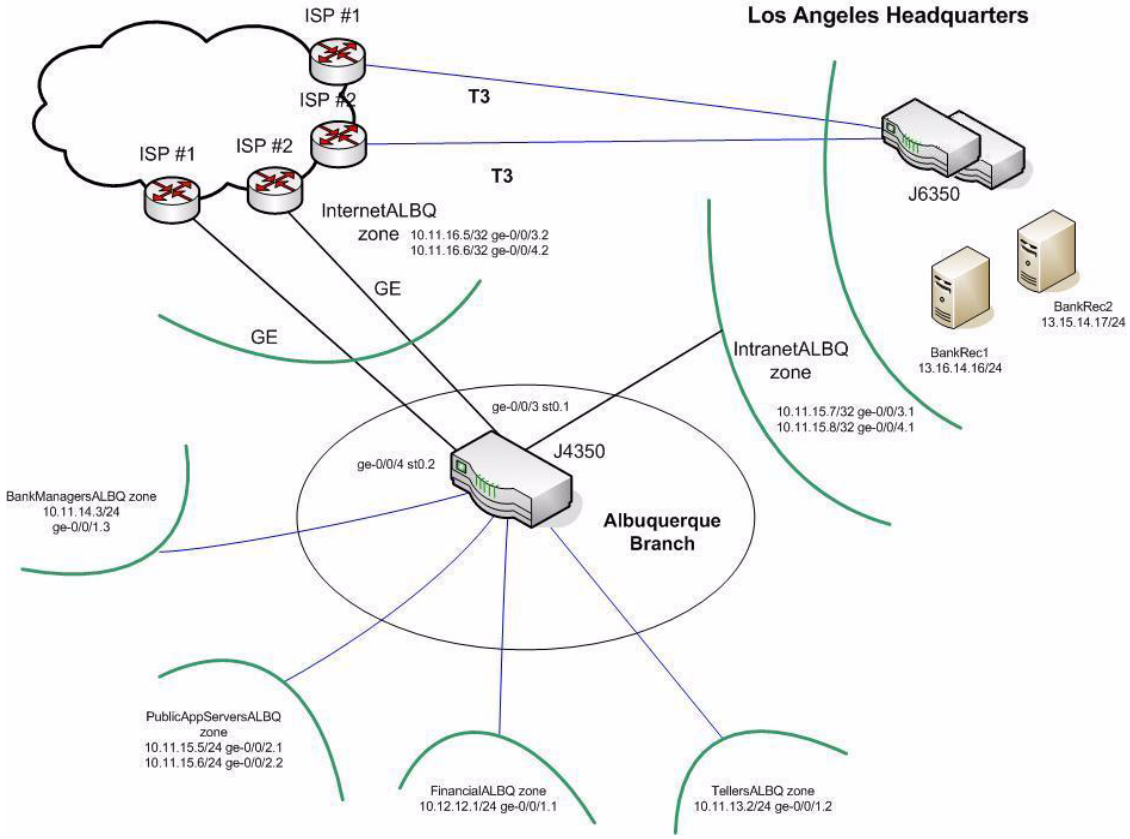

About Address Translation

Many local area networks (LANs) rely on use of private IP addresses for their subnetworks and other hosts. Private addresses allow LANs to conserve use of public IP addresses— to allow multiple hosts to use a single public IP address and to reuse public IP addresses as they are freed up from sessions—and to enhance security by disguising the internal network’s structure and the identity of its hosts.

About Network Address Translation

When hosts with private addresses initiate traffic to a public address space, such as the Internet, the router must translate the private source IP address to a public one recognizable by systems outside the company's network. Also, when sending traffic from one private address space through a VPN to a site that happens to use the same private addresses, the routers at both ends of the tunnel must agree to translate the source and destination IP addresses to mutually neutral addresses.

Network Address Translation (NAT), which is a method of mapping a private IP address in a source packet's header to a public IP address or a public IP address in a destination packet's header to a private IP address, is used in these cases. To a host outside the network, the mapped public IP address appears to be the actual source or destination address of the traffic.

NOTE: Network Address Translation (NAT) is not supported for policy-based IPSec VPNs.

About Port Address Translation and Port Mapping

When Port Address Translation (PAT) is used, the router maintains a list of port numbers to be assigned to a packet to match sessions to the hosts they belong to. With PAT enabled:

- Up to 64,500 hosts can share a single IP address.
- Range 1024 to 65535 is available for port number mapping for each IP address.
- Within the available range:
  - Range 1024 through 63487 is allocated one port a time
  - Range 63488 through 65535 is allocated two consecutive ports at a time.
  - Range 63488 through 65535 is for RTP and CRTP applications, such as SIP, H.323, and RTSP.

You can use PAT for some source NAT configurations. Figure 15 shows a packet header before and after the source IP address and the port address have been translated using source NAT with PAT.

Figure 15: Source IP and Source Port Address Translation

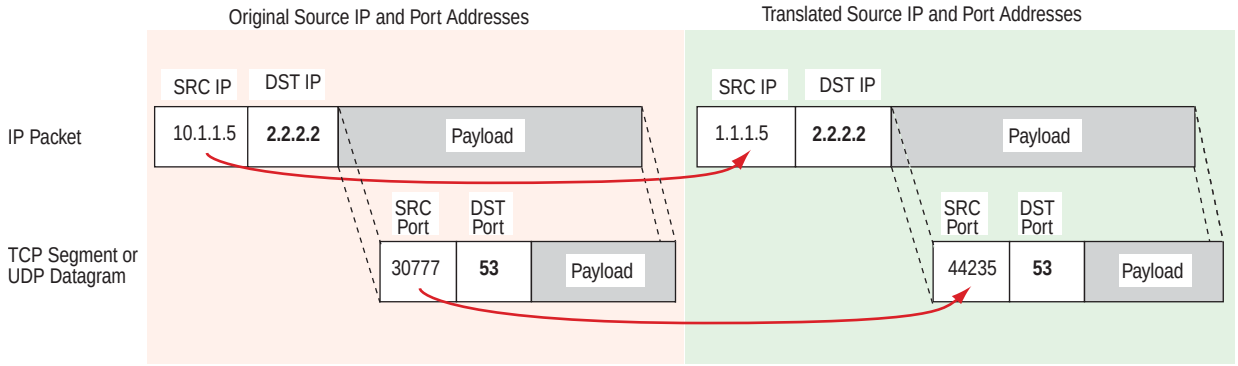

For destination NAT configurations that support port mapping, a different technique is used. Port mapping is the deterministic translation of one public destination port number to another specific private port number. Port mapping differs from Port Address Translation (PAT), which translates any public source port number randomly assigned to the initiating host to another number randomly assigned by the router.

About NAT, PAT, and Port Mapping Implementation

JUNOS software with enhanced services running on a Services Router provides many methods for performing NAT, including translation of port numbers in a packet if Port Address Translation (PAT) or port mapping is also enabled.

You can configure NAT to translate source private IP addresses, destination private IP addresses, or both:

- You can select from many different styles of performing source NAT to translate one or more hosts' private source IP addresses in outgoing packets to public IP addresses.

  For details, see "Configuring Source NAT" on page 205.

- You can use any of several methods to translate destination public IP addresses in incoming packets to their hosts' private IP addresses.

  For details, see "Configuring Destination NAT" on page 219.

Configuring Source NAT

Source Network Address Translation (NAT) entails mapping a private source IP address to a public one. You can dynamically map multiple private source IP addresses to individual public IP addresses. You can map multiple private source IP addresses to a single public IP address. You can also use source NAT to statically map a range of private source IP addresses to a range of public IP addresses so that the one-to-one mapping is maintained across sessions. Not only the private IP address but also the port address can be translated if Port Address Translation (PAT) is enabled. Figure 16 shows a packet header before and after the router performs source NAT translation for the source IP address only.

Figure 16: Source IP Address Translation

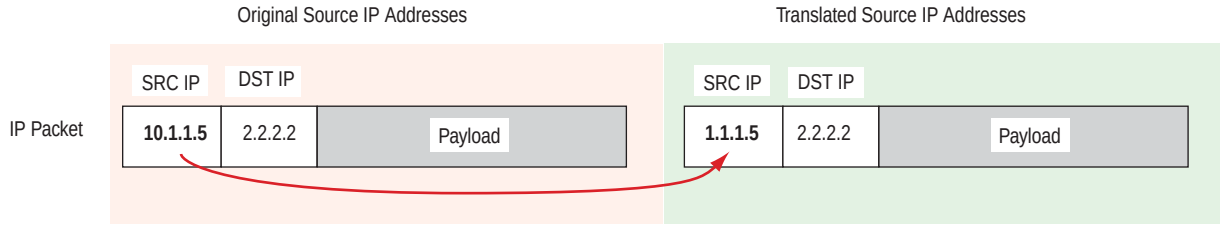

After your zones and address books are set up, to use source NAT, you perform the following tasks:

- Configure source pools on a zone's egress interface.
- Create policies and, from within a policy that stipulates use of source NAT, refer to the source pool to be used.

If you use interface source NAT, the tasks are somewhat different. For details on the specific steps to follow, see "Interface Source NAT" on page 217.

The JUNOS software with enhanced services allows you to configure many different kinds of source pools to satisfy different requirements for source NAT. See "About Source Pools" on page 207 for a description of them.

How you configure a source pool determines its behavior. You use the same policy syntax for specifying various kinds of source NAT. The particular source pool that you specify in the policy determines the method of NAT to be used.

Table 129 gives an overview of the entire policy configuration syntax. The action that applies to source NAT (source-nat) and its attributes is shown in bold.

Table 129: Policy Configuration Statement with Source NAT Action

Policy Configuration	Syntax
Configure a policy.	`set security policies from-zone <zone-name> to-zone <zone-name> policy <policy-name> match source-address (<address-name> | <address-set-name>) destination-address (<address-name> | <address-set-name>) application (<application-name> | <application-set-name>)` `set security policies from-zone <zone-name> to-zone <zone-name> policy <policy-name> then (permit | firewall-authentication | tunnel ipsec-vpn <vpn-name> pair-policy <pair-policy-name> | source-nat (pool <pool-name> | pool-set <pool-set-name> | interface) | destination-nat <name> | deny | reject schedule <scheduler-name> log <session-init session-close> count alarm per-second-threshold <value> per-minute-threshold <value>`  ■ `from-zone zone-name`: Source zone. Name of the zone from which traffic is sent. ■ `to-zone zone-name`: Destination zone. Name of the zone to which traffic is to be sent. ■ `policy-name`: Unique name used to refer to the policy. ■ `source-address address-name`: Name of an address (or address set) as entered in the source zone's address book. ■ `destination-address address-name`: Name of an address (or address set) as entered in the destination zone's address book. ■ `application application-name`: Name of a preconfigured or custom application (or application set). ■ action: Any one of the actions listed in the syntax following the keyword action. For source-nat, specify   ■ `pool` or `pool-set`: The name of the source pool or source pool set from which the public IP address or addresses are to be drawn.   ■ `interface`: If interface source NAT is to be used, the name of the interface whose IP address is to be used as the public source IP address. ■ `scheduler-name`: Optionally, the name of a scheduler whose schedule determines when the policy is active and when it can be used. ■ `log session-init session-close`: Log the traffic that matches the policy ■ count values: Count setters.

This section includes the following source NAT configurations for the Albuquerque branch office:

- "Source NAT with PAT: Dynamically Mapping Multiple Private IP Addresses to Public IP Addresses" on page 209
- "Source NAT with PAT: Mapping Multiple Private IP Addresses to a Single Public IP Address" on page 211
- "Source NAT with PAT: Configuring Address Persistence" on page 219
- "Static Source NAT: Statically Mapping a Private IP Address Range to a Range of Public IP Addresses" on page 212
- "Source NAT Without PAT" on page 215
- "Interface Source NAT" on page 217

About Source NAT

The router supports the following forms of source NAT:

- Source NAT with the source pool defined on the zone's egress interface. The router supports various methods of performing this kind of source NAT depending on the source pool to be used. "About Source Pools" on page 207 describes the kinds of public IP address source pools you can define on interfaces.
- Interface source NAT in which private source IP addresses are translated to the IP address of the egress interface.
- Reverse mapping in static NAT. This kind of source IP address mapping is a result of the bidirectional mapping that is done when destination static NAT is configured. In this case, the policy does not contain an explicit NAT action. If traffic matches a policy and the action is permit and reverse mapping in static NAT is found on the egress interface, the router performs source IP address translation according to reverse mapping. For details, see "Destination NAT: Configuring Static NAT" on page 221.

About Source Pools

Source pools provide the router with a supply of public IP addresses from which to draw when performing source NAT. You configure source pools on egress interfaces and give them names. From within a policy that specifies use of source NAT, you refer by name to the source pool to be used.

The kind of source pool you specify in a policy determines the type of source NAT the router performs on traffic matching that policy. The information you provide when you configure a source pool determines its behavior.

When the router processes outgoing traffic that matches the policy and source NAT is specified, the router rewrites the private IP address in the TCP or UDP packet header. It specifies the new public IP address taken from the specified source pool, and, optionally, a new port address, before it transmits the packet.

You can configure a source pool with or without PAT, a static source pool, an interface source pool, and a source pool set.

Source Pools with PAT

When a policy refers to a NAT source address pool that supports PAT, the router translates both the source IP address and the port number to new values taken from the specified pool. PAT allows for a single public IP address to be shared by multiple hosts.

You can also configure a range of public IP addresses to which private IP addresses can be mapped. In this case, the router translates the private IP source address in a packet to one of the addresses in the range. It is not a one-to-one static mapping of one range to another, as is the case with static source NAT. Rather, the router randomly maps a public IP address from the range to the private source one depending on when the public IP addresses are available.

For example, when a session that uses a translated address is terminated, the public IP address is freed up for use in another mapping. This random assignment might afford another level of security by impeding the attempts of attackers to determine the actual host behind the public IP address.

For sample configurations of this use of NAT for the branch deployment, see the following sections:

- Source NAT with PAT: Dynamically Mapping Multiple Private IP Addresses to Public IP Addresses on page 209
- Source NAT with PAT: Mapping Multiple Private IP Addresses to a Single Public IP Address on page 211

These sections present configurations that create source pools and establish policies to enable source NAT with PAT IP address translation for packets sent from the financial managers' hosts, which belong to the FinancialALBQ zone, to any address in the InternetALBQ zone. For IP address translation to occur, the packets must match all criteria of the policies. PAT is performed automatically in these cases.

For source NAT with PAT, the security device might assign to a single host different addresses for different concurrent sessions unless the source pool or the router is configured for address persistence. (For details on address persistence, see "Source NAT with PAT: Configuring Address Persistence" on page 219.)

Source Pools Without PAT

The router dynamically selects a public source IP address from the source pool specified by a policy, maps it to the host's private IP address, replacing the private one in the packet headers, but it does not translate the original port address.

This type of source pool supports the requirements of applications that must maintain a fixed source port number. For source pools without PAT, the router assigns one translated source address to the same host for all its concurrent sessions. For example, a host may have two Telnet sessions active concurrently. In this case, the same translated IP address is used for both of them. For details on this configuration, see "Source NAT Without PAT" on page 215.

Static Source Pools

The router defines a one-to-one mapping between an original private source IP address to a translated public IP address for a range of IP addresses. The static mapping ensures that the router always translates a particular source IP address from within the range to the same translated address provided by the source pool. For details on this configuration, see "Destination NAT: Configuring Static NAT" on page 221.

Interface Source Pools

When interface source pool is specified in a policy, the router translates the private source IP address of a packet that matches the policy to the address of the egress interface. The router always applies PAT for interface source pool to distinguish sessions from one another.

For details, see "Interface Source NAT" on page 217

Source Pool Sets

NAT is configured on interfaces, and policies are configured in relation to zones. Source and destination zones are specified in a policy. If a zone has multiple egress interfaces, you can define a source pool on each egress interface of a zone and group the single pools into a pool set, then refer to the pool set in a policy.

The router translates the source IP address and port number using the correct source pool within the pool set for traffic matching the policy and existing on a a particular egress interface in the zone. For details on source pools, see the *JUNOS Software Security Configuration Guide*.

Source NAT with PAT: Dynamically Mapping Multiple Private IP Addresses to Public IP Addresses

This sample configuration creates a source NAT with PAT address pool containing a range of public IP addresses. It specifies a policy that translates the private IP addresses of any packets from the financial managers' hosts that match the policy criteria to one of the public IP addresses in the range. Port addresses are also translated. Public IP addresses are assigned dynamically. That is, they are assigned randomly. For details on the source pool, see "Source Pools with PAT" on page 207.

In the following sample configuration, the private IP addresses assigned to the hosts of FinancialManager1 (10.12.12.11), FinancialManager2 (10.12.12.12), FinancialManager3 (10.12.12.13), and FinancialManager4 (10.12.12.14) are assigned to public IP addresses within the source pool's range of addresses.

Public IP address assignments are made when a packet is transmitted, and they are not static. A different address might be assigned to a private IP address when a new session is started.

This sample configuration performs the following tasks:

- Configures an interface called ge-0/0/1.1 and assigns it to the FinancialALBQ zone.
- Configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- Adds the financial managers' addresses—10.12.12.11, 10.12.12.12, 10.12.12.13, and 10.12.12.14—to the FinancialALBQ zone's address book.
- Defines a source pool called src-nat-multiple-addr on the InternetALBQ egress interface and assigns the public IP address range 11.12.1.20 to 11.12.1.30 to it.
- Defines a policy (src-nat-multiple-addr-policy) that permits traffic from financial managers 1, 2, 3, and 4 if the packets are HTTP traffic destined for the Internet. The policy directs the router to translate the private IP addresses in the packet headers of matching traffic from the managers to different public IP addresses in the specified range with PAT also performed.

Interface and zone configurations

```
user@host# set interfaces ge-0/0/1 unit 1 family inet address 10.12.12.1/24
user@host# set security zones security-zone FinancialALBQ interfaces ge-0/0/1.1
user@host# set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2
```

For details on configuring interfaces and zones, see “Configuring Interfaces and Assigning Them to Zones” on page 49.

Address book assignments

user@host# **set security zones security-zone FinancialALBQ address-book address FinancialManager1 10.12.12.11/32**

user@host# **set security zones security-zone FinancialALBQ address-book address FinancialManager2 10.12.12.12/32**

user@host# **set security zones security-zone FinancialALBQ address-book address FinancialManager3 10.12.12.13/32**

user@host# **set security zones security-zone FinancialALBQ address-book address FinancialManager4 10.12.12.14/32**

For details on configuring address books, see “Configuring Address Books for Zones” on page 68.

Source pool with PAT and address range on on egress interface

user@host# **set security nat interface ge-0/0/3.2 source-nat pool src-nat-multiple-addr address-range low 11.12.1.20 high 11.12.1.30**

For the statement syntax for configuring source pools with PAT, see Table 130 on page 210.

Policy specifying source NAT with PAT translation for multiple addresses

user@host# **set security policies from-zone FinancialALBQ to-zone InternetALBQ policy src-nat-multiple-addr-policy match source-address FinancialManager1 FinancialManager2 FinancialManager3 FinancialManager4 destination-address any application http**
user@host# **set security policies from-zone FinancialALBQ to-zone InternetALBQ policy src-nat-multiple-addr-policy then permit source-nat pool src-nat-multiple-addr**

For the syntax for configuring policies, see Table 129 on page 206.

Table 130 shows the statement syntax for mapping an address range with PAT.

Table 130: Source Pool with PAT Configuration Statement for Mapping an IP Address Range

Source Pool with PAT for Address Range	Syntax
Defines a source pool containing a range of public IP addresses.	`set security nat interface <interface-name> source-nat pool <pool-name> address-range low <ipaddr> high <ipaddr>` ■ `interface-name`: Name of the interface on which the source pool is to be configured. ■ `pool-name`: Name of the source pool. ■ (address low) `ip-addr`: For a source pool with a range of addresses, the IP address of the beginning of the public IP address range. ■ (address high) `ip-addr`: For a source pool with a range of addresses, the IP address of the end of the public IP address range.

Source NAT with PAT: Mapping Multiple Private IP Addresses to a Single Public IP Address

This sample configuration creates a source NAT with PAT address pool with a single public IP address. It specifies a policy that translates the private IP address of any packets from two of the financial managers' hosts that match the policy criteria to the single public IP address in the address pool. The router uses PAT to differentiate the sessions belonging to the hosts sharing the same public IP address.

In the following sample configuration, the private IP addresses of the hosts for financial manager 3 (10.12.12.13) and financial manager 4 (10.12.12.14) are assigned to the single public IP address (11.12.1.2) in the source pool.

The sample configuration performs the following tasks:

- It configures an interface called ge-0/0/1.1 and assigns it to the FinancialALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It creates an address set called FinTelnet including the private IP addresses 10.12.12.13/32 and 10.12.12.14/32 for the two financial manager's hosts and assigns the address set to the FinancialALBQ zone's address book.
- It defines a source pool called src-nat-1-address on the InternetALBQ egress interface ge-0/0/3.2 and adds the single public IP address 11.12.1.2 to it.
- It defines a policy (src-nat-1-addr-policy) that permits traffic from financial managers 3 and 4 if the packets are Telnet traffic destined for the Internet. The policy directs the router to translate the private IP addresses in the packet headers of matching traffic from either of the two managers to the same public IP address 11.12.1.2.

Interface and zone configurations

```
user@host# set interfaces ge-0/0/1 unit 1 family inet address 10.12.12.1/24
user@host# set security zones security-zone FinancialALBQ interfaces ge-0/0/1.1
```

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

```
user@host# set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2
```

Address book assignments

```
user@host# set security zones security-zone FinancialALBQ address-book
address-set FinTelnet 10.12.12.13/32
```

```
user@host# set security zones security-zone FinancialALBQ address-book
address-set FinTelnet 10.12.12.14/32
```

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Source pool with PAT and single address on egress interface

```
user@host# set security nat interface ge-0/0/3.2 source-nat pool src-nat-1-address
address 11.12.1.2
```

For the statement syntax for configuring source pools with PAT, see Table 131.

Policy specifying source NAT with PAT translation

user@host# **set security policies from-zone FinancialALBQ to-zone InternetALBQ policy src-nat-1-addr-policy match source-address FinTelnet destination-address any application telnet**

user@host# **set security policies from-zone FinancialALBQ to-zone InternetALBQ policy src-nat-1-addr-policy then permit source-nat pool src-nat-1-address**

For the statement syntax for configuring policies, see Table 129 on page 206.

Table 131 shows the statement syntax for configuring many private IP addresses mapped to one public IP address.

Table 131: Source Pool with PAT Configuration Statement for Mapping Many-to-One

Source Pool with PAT	Syntax
Defines a source pool containing one public IP address. It can be used in a policy to map many private IP addresses to one public IP address	`set security nat interface <interface-name> source-nat pool <pool-name> address <ip-addr>` ■ `interface-name`: Name of the interface on which the source pool is to be configured. ■ `pool-name`: Name of the source pool. ■ (address) `ip-addr`: The IP address for a source pool containing a single IP address.

Static Source NAT: Statically Mapping a Private IP Address Range to a Range of Public IP Addresses

When you configure source NAT with PAT to map a range of addresses, a private IP address is not always assigned to the same public IP address from within the range of public ones. In some cases, you might want to ensure that the same private IP addresses are mapped to the same public IP addresses in a range. For this purpose, you can use a static source pool. Because the addresses are statically mapped, PAT is not required.

NOTE: You can configure a policy that applies to a range of private source IP addresses which exceed the number of public IP addresses in the source pool range. In such cases, the router applies static source NAT to those addresses that fall within the range of the pool's public IP addresses. It transmits the other packets, whose addresses fall outside the static source pool range, but it leaves their IP addresses unchanged.

Figure 17 shows static source NAT translation for packets sent from a range of hosts in the FinancialALBQ zone to a server in the InternetALBQ zone.

Figure 17: Statically Mapping a Range of Private Source IP Addresses to a Public IP Address Range

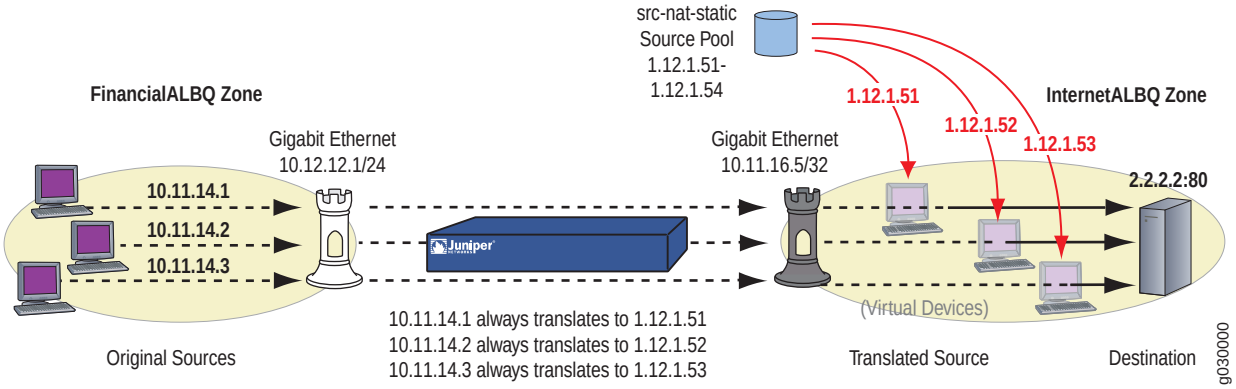

This sample configuration creates a static source NAT address pool containing a range of public IP addresses mapped statically to a range of private IP addresses that are identified by the first private IP address, specified as the low one.

The sample configuration specifies a policy that directs the router to translate the private source IP addresses of packets from the three bank managers' hosts that match the policy criteria to the public IP addresses specified by the pool range. For all of its sessions that match the policy, a bank manager's source IP address is always translated to the same public IP address to which it is mapped.

In this sample configuration:

- The private IP address assigned to BankManager1's host (10.11.14.1) is assigned to the first public IP address in the source pool address range, address 1.12.1.51.
- The private IP address assigned to BankManager2's host (10.11.14.2) is assigned to the second public IP address in the source pool address range, address 1.12.1.52.
- The private IP address assigned to BankManager3's host (10.11.14.3) is assigned to the third public IP address in the source pool address range, address 1.12.1.53.

The sample configuration performs the following tasks:

- It configure an interface called ge-0/0/1.3 and assigns it to the BankManagerALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It assigns the private IP addresses 10.11.14.1 through 10.11.14.3 to the BankManagersALBQ zone.
- It defines a source pool called src-static-nat-policy on the InternetALBQ egress interface ge-0/0/3.2 and adds the public IP address range 1.12.1.51 through 1.12.1.54 to it. It also adds the beginning of the bank manager's address range, that of BankManager1 (10.11.14.1), to the source pool, to identify the beginning of the private source IP addresses to map.

The router begins mapping in order from the first private IP address—one private IP address to one public address in the range—through the third one. You only need to specify the first host IP address in the range.

- It defines a policy (src-static-nat-policy) that permits traffic from any of the three bank managers if the packets are SSH traffic destined for the Internet.

  Regardless of how packets are sent from their hosts, the IP addresses are assigned in order of the range. For example, if BankManager2's host transmits traffic that matches the policy before BankManager1's host does, the router still translates the private IP address 10.11.14.2 in packets from BankManager2 to the public IP address 1.12.1.52. This is not the case when you use source NAT with PAT to map multiple public IP addresses to multiple public IP addresses in a range.

Interfaces and zones configuration

```
user@host# set interfaces ge-0/0/1 unit 3 family inet address 10.11.14.3/24
user@host# set security zones security-zone BankManagersALBQ interfaces
ge-0/0/1.3

user@host# set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2
```

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignments

```
user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan1 10.11.14.1/32

user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan2 10.11.14.2/32

user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan3 10.11.14.3/32
```

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Static source pool configuration

```
user@host# set security nat interface ge-0/0/1.3 source-nat pool src-nat-static
address-range low 1.12.1.51 high 1.12.1.54
user@host# set security nat interface ge-0/0/1.3 source-nat pool src-nat-static
host-address-low 10.11.14.1
```

When you configure the pool, you must specify the public IP address range as a separate statement from the private IP address of the first host in the host range, as shown in this sample configuration.

Policy specifying source static NAT

```
user@host# set security policies from-zone BankManagersALBQ to-zone
InternetALBQ policy src-static-nat-policy match source-address bankMan1
bankMan2 bankMan3 bankMan4 destination-address any application ssh

user@host# set security policies from-zone BankManagersALBQ to-zone
InternetALBQ policy src-static-nat-policy then permit source-nat pool src-nat-static
```

For the statement syntax for configuring policies, see Table 129 on page 206.

Table 132 shows the statement syntax for configuring a static NAT source pool.

Table 132: Source Pool with PAT Configuration Statement for Statically Mapping a Range of IP Addresses

Static NAT Source Pool	Syntax
Defines a source pool containing multiple public IP addresses. It can be used to statically map a range of private IP addresses to the public range	`set security nat interface <interface-name> source-nat pool <pool-name> address-range low <ip-addr> high <ip-addr> host-address-low <ip-addr>` ■ `interface-name`: Name of the interface on which the source pool is to be configured. ■ `pool-name`: Name of the source pool. ■ `(address-range low) ip-addr`: For a source pool with a range of addresses, the IP address of the beginning of the public IP address range. ■ `(high) ip-addr`: For a source pool with a range of addresses, the IP address of the end of the public IP address range. ■ `host-address-low`: For a range of addresses, the private IP address of the host at the beginning of the range.

Source NAT Without PAT

Certain configurations or situations might require you to perform source NAT without performing PAT. For example, a custom application might require a specific number for the source port address. If you want the port number to remain fixed when the router maps the private source IP address to a public one, you must disable PAT. You disable PAT on the source pool to be referred to in the policy.

For the sample configuration described in this section, Figure 18 shows how source NAT without PAT is performed to translate a single private source IP address to a single public one.

Figure 18: Source NAT Without PAT

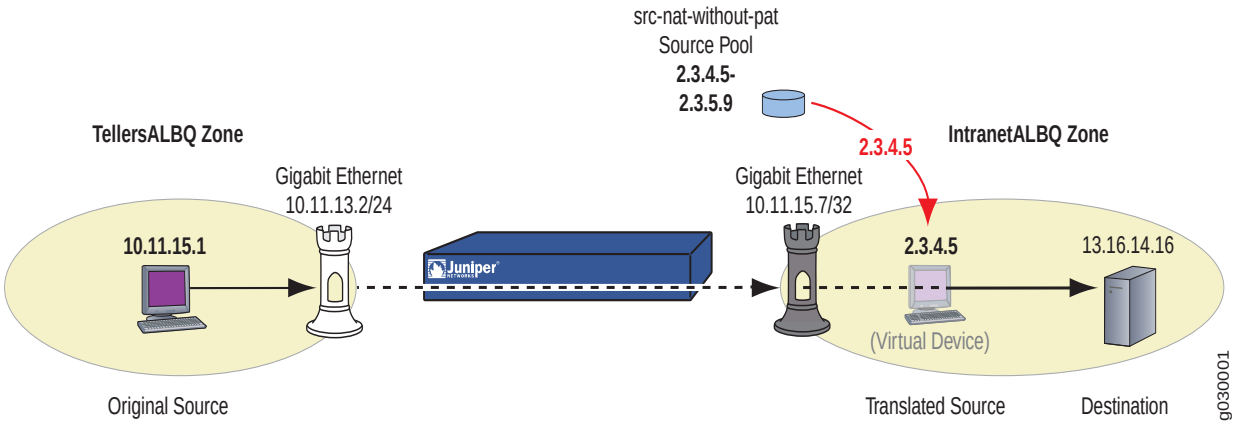

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/3.1 and assigns it to the IntranetALBQ zone.
- It configures an interface called ge-0/0/1.2 and assigns it to the TellersALBQ zone.

- It assigns the teller's subnet IP address (10.11.13.0) to the TellersALBQ zone's address book.
- It assigns the IP addresses (13.16.14.16 and 13.15.14.17) of the servers at corporate headquarters to the IntranetALBQ zone's address book.
- It defines a source pool called src-nat-without-pat on the IntranetALBQ's egress interface ge-0/0/3.1 containing the public IP address range 2.3.4.5 to 2.3.5.9.
- It specifies that the router should not do PAT.
- It defines a policy (src-nat-nopat-policy) that permits traffic from any of the hosts on the tellers subnet (TellersALBQ) if the packets are custApp (custom application) traffic destined for either of the two servers at corporate headquarters. In this case, it specifies that the private IP addresses in packets from tellers that match the policy are to be translated to the public IP addresses in the source pool's configured range.

Interfaces and zones configuration

user@host# **set interfaces ge-0/0/3 unit 1 family inet address 10.11.15.7/32**
user@host# **set security zones security-zone IntranetALBQ interfaces ge-0/0/3.1**

user@host# **set interfaces ge-0/0/1 unit 2 family inet address 10.11.13.2/24**
user@host# **set security zones security-zone TellersALBQ interfaces ge-0/0/1.2**

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignments

user@host# **set security zones security-zone TellersALBQ address-book address tellers 10.11.13.0/24**

user@host# **set security zones security-zone IntranetALBQ address-book address BankRec1 13.16.14.16/32**

user@host# **set security zones security-zone IntranetALBQ address-book address BankRec2 13.15.14.17/32**

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Source NAT without PAT pool configuration

user@host# **set security nat interface ge-0/0/3.1 source-nat pool src-nat-without-pat address-range low 2.3.4.5 high 2.3.5.9**

Disable PAT

user@host# **set security nat interface ge-0/0/3.1 source-nat pool src-nat-without-pat no-port-translation**

Policy specifying source NAT without PAT

user@host# **set security policies from-zone TellersALBQ to-zone IntranetALBQ policy src--nat-nopat-policy match source-address tellers destination-address BankRec1 BankRec2 application custApp**

user@host# **set security policies from-zone TellersALBQ to-zone IntranetALBQ policy src--nat-nopat-policy then permit source-nat pool src-nat-without-pat**

For the statement syntax for configuring policies, see Table 129 on page 206.

Table 133 shows the syntax for configuring a NAT source pool IP address range with no PAT. As shown in the preceding sample configuration, you must specify the source pool IP address contents and the no-port-translation designation separately.

Table 133: Source Pool with Address Range and Without PAT Configuration Statement

Source Pool without PAT	Syntax
Defines a source pool containing an address range and specifying that PAT is not performed.	`set security nat interface <interface-name> source-nat pool <pool-name> address-range low <ip-addr> high <ip-addr> no-port-translation` ■ `interface-name`: Name of the interface on which the source pool is to be configured. ■ `pool-name`: Name of the source pool. ■ `(address-range low) ip-addr`: For a source pool with a range of addresses, the IP address of the beginning of the public IP address range. ■ `(high) ip-addr`: For a source pool with a range of addresses, the IP address of the end of the public IP address range. ■ `no-port-translation`: Directs the router not to do PAT

Interface Source NAT

If you apply source NAT to a policy that does not specify a source pool but instead specifies that the IP address of the egress interface in the destination zone is to be used, for all traffic that matches the policy the router translates the private source IP address to the public IP address of the egress interface. The router always applies PAT for interface source pool.

TIP: This sample configuration includes no source pool and instead specifies source NAT interface.

For this sample configuration, Figure 19 shows how interface source NAT is applied to translate the private source IP address to that of the egress interface.

Figure 19: Interface Source NAT

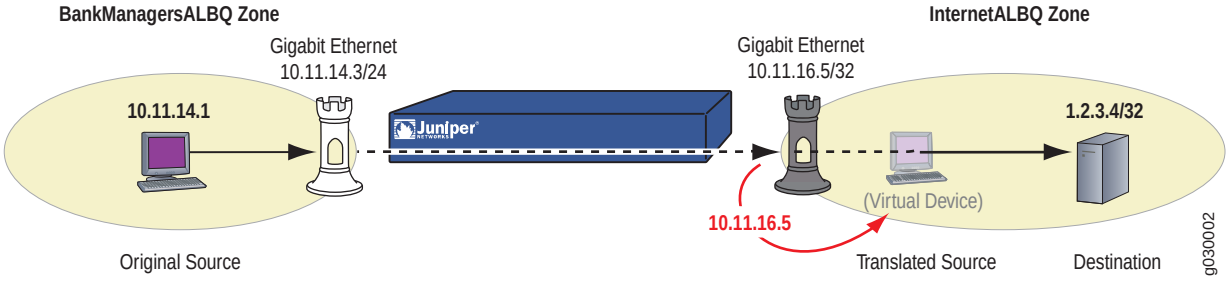

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/1.3 and assigns it to the BankManagersALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.

- It assigns the private IP addresses 10.11.14.1 through 10.11.14.4 to the BankManagersALBQ zone.
- It defines a policy called src-interface-policy that allows traffic from any of the specified bank managers in the BankManagersALBQ zone to any destination in the InternetALBQ zone if the packet specifies HTTPS as the application. In this case, it translates the private IP address in the packet's header to the IP address of the egress interface (10.11.16.5) in the InternetALBQ zone.

Interfaces and zones configuration

```
user@host# set interfaces ge-0/0/1 unit 3 family inet address 10.11.14.3/24
user@host# set security zones security-zone BankManagersALBQ interfaces
ge-0/0/1.3

user@host# set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2
```

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignments

```
user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan1 10.11.14.1/32

user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan2 10.11.14.2/32

user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan3 10.11.14.3/32

user@host# set security zones security-zone BankManagersALBQ address-book
address bankMan4 10.11.14.4/32
```

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Policy specifying source NAT without PAT

```
user@host# set security policies from-zone BankManagersALBQ to-zone
InternetALBQ policy src-interface-policy match source-address bankMan1 bankMan2
bankMan3 bankMan4 destination-address any application https

user@host# set security policies from-zone BankManagersALBQ to-zone
InternetALBQ policy src-interface-policy then permit source-nat interface
```

For the statement syntax for configuring policies, see Table 129 on page 206.

Because the security device translates all original IP addresses to the same translated IP address (that of the egress interface), the router uses the unique, translated port number to identify each session to which a packet belongs and to maintain session state information for traffic to and from the same, single IP address.

Source NAT with PAT: Configuring Address Persistence

When the router performs source NAT with PAT for a host that initiates several sessions concurrently, it assigns a different IP address for each session. Services that create multiple sessions require the same source IP address for each session. To ensure that the router assigns the same IP address from a source pool to a host for multiple concurrent sessions, you use the address persistence feature. Address persistence especially addresses the requirements of Voice-over-IP (VoIP) applications, such as the SIP ALG.

NOTE: When you enable address persistence, it applies to all source IP address mappings because it is enabled for the router.

The following statement enables source NAT address persistence for the router. Table 134 shows the syntax statement.

Address persistence — user@host# **set security nat source-nat address-persistent**

Table 134: IP Address Assignment Persistence Configuration Statement

Source	Syntax
Specifies that all public IP addresses are to be statically assigned to their mapped private source IP addresses.	`set security nat source-nat address-persistent` ■ `address-persistent`: Directs the router to maintain the IP address mappings across sessions globally, for all mappings on the router.

Configuring Destination NAT

You can use destination Network Address Translation (NAT) to map an incoming public destination IP address in a packet header to a private IP address. For policy-based destination NAT, the destination port number can also be mapped.

JUNOS software with enhanced services provides the following ways to perform destination NAT:

- Static destination NAT defined on the ingress interface.

  You can translate

  - One public destination IP address to its mapped private IP address statically so that the mapping applies across sessions.
  - A range of public destination IP addresses to a mapped private IP address range (such as one subnet to another), so that the security device consistently maps each original public destination address to a specific translated private IP address for all of its sessions. This is a one-to-one mapping by subnet. In this case, PAT is not supported. For destination static mapping, both the source and destination IP addresses are mapped.

  See “Destination NAT: Configuring Static NAT” on page 221

- Policy-based destination NAT, defined within a policy. For details, see “About Policy-based Destination NAT” on page 224.

JUNOS software with enhanced services also supports a form of NAT called allow-incoming that is used for VoIP ALGs. This NAT type is not discussed in this guide. For details, see the *JUNOS Software Security Configuration Guide*.

Figure 20 shows two of the methods you can use to perform destination NAT. In the first case, the router maps one public destination IP address to one private IP address. In the second case, the router maps two public destination IP addresses to a single private IP address.

Figure 20: Types of Destination NAT

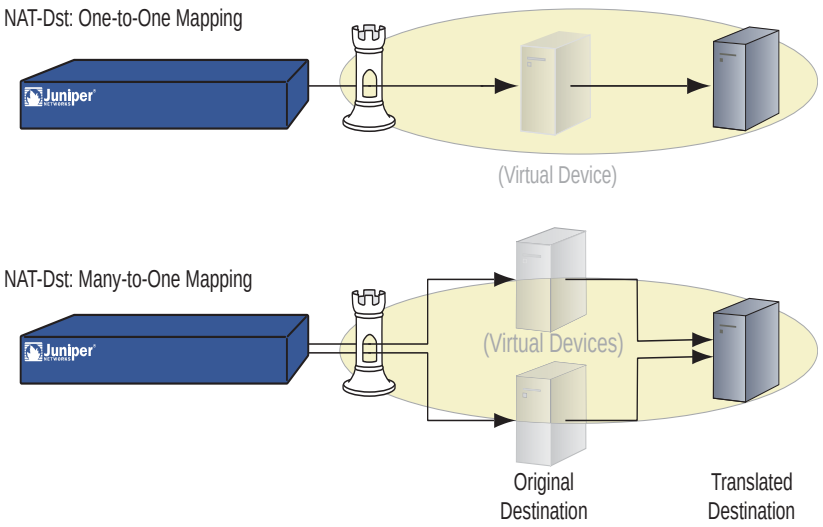

Note: The original and the translated destination IP addresses must be in the same security zone.

About Policy-based Destination NAT and Route Configuration

For policy-based destination NAT to work, the routing table must contain entries for both the public destination IP address and the translated private destination IP address so that the router can perform two route lookups:

1. First the router performs a route lookup using the public destination IP address to determine the destination zone for a subsequent policy lookup.

   This lookup allows the router to determine the egress interface. The egress interface, in turn, provides the destination zone so that the router can do a policy lookup.

   When the router finds a policy match, the router translates the public destination IP address to the translated private destination IP address.

2. The router then performs a second route lookup using the translated private IP address to determine where to send the packet.

   The second route lookup allows the router to determine the interface through which it must forward the packet to reach the new destination IP address.

To ensure that the routing decision is in accord with the policy, both the public destination IP address and the translated private IP address must be in the same security zone.

In summary, the route to the public IP address provides a means to perform the policy lookup, and the route to the translated private IP address specifies the egress interface through which the router is to forward the packet.

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

Destination NAT: Configuring Static NAT

Destination static NAT is a direct static mapping of one public IP address to a private IP address without port address translation (PAT). The mapping can be a single IP address or it can be a range of IP addresses for a subnet. You configure destination static NAT on an ingress interface.

For static NAT, for incoming traffic that matches the policy, the router translates the static NAT public IP address to the private IP address of the host to which it is mapped and transmits the traffic to it. The one-to-one mapping is constant across sessions so that a public destination IP address is always mapped to the same private IP address of a host.

NOTE: Unlike other forms of source and destination NAT, static NAT is bidirectional. When a mapped host initiates outbound traffic, the router translates the private source IP address of the host to the public IP address to which it is statically mapped through destination NAT.

To refer to static NAT in a policy, you must specify static_nat_*ipprefix* as the destination address. For the destination zone, you must specify junos-global, which is a predefined zone. The router stores all its static NAT mappings in relation to the global zone.

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/1.3 and assigns it to the BankManagersALBQ zone.
- It configures an interface called ge-0/0/3.1 and assigns it to the IntranetALBQ zone.
- It assigns the IP addresses (13.16.14.16 and 13.15.14.17) of the servers at corporate headquarters to the IntranetALBQ zone's address book
- It assigns the addresses of the bank managers (bankMan1 10.11.14.1, bankMan2 10.11.14.2, bankMan3 10.11.14.3, and bankMan4 10.11.14.4) to the BankManagersALBQ zone's address book.
- It creates the following static mappings:
  - It maps bank manager 1's (bankMan1) address 10.11.14.1 statically to the public address 5.6.7.8.
  - It maps bank manager 2's (bankMan2) address 10.11.14.2 to 5.6.7.9.
  - It maps bank manager 3's (bankMan3) address 10.11.14.3 to 5.6.8.0.

  - It maps bank manager 4's (bankMan4) address 10.11.14.4 to 5.6.8.1.

- It creates policies (static-nat-policy1 through static-nat-policy8) that refer to the appropriate static mappings for the four managers. Each policy stipulates that when incoming traffic from one of the IntranetALBQ zone's servers is addressed to the specified statically mapped IP address, the public address is to be translated to the mapped private IP address.

Interfaces and zones configuration

user@host# **set interfaces ge-0/0/1 unit 3 family inet address 10.11.14.3/24**
user@host# **set security zones security-zone BankManagersALBQ interfaces ge-0/0/1.3**

user@host# **set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32**
user@host# **set security zones security-zone IntranetALBQ interfaces ge-0/0/3.1**

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49

Address book assignments

user@host# **set security zones security-zone BankManagersALBQ address-book address bankMan1 10.11.14.1/32**

user@host# **set security zones security-zone BankManagersALBQ address-book address bankMan2 10.11.14.2/32**

user@host# **set security zones security-zone BankManagersALBQ address-book address bankMan3 10.11.14.3/32**

user@host# **set security zones security-zone BankManagersALBQ address-book address bankMan4 10.11.14.4/32**

user@host# **set security zones security-zone IntranetALBQ address-book address BankRec1 13.16.14.16/32**

user@host# **set security zones security-zone IntranetALBQ address-book address BankRec2 13.15.14.17/32**

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Static NAT mappings

user@host# **set security nat interface ge-0/0/1/3 static-nat 5.6.7.8 host 10.11.14.1**
user@host# **set security nat interface ge-0/0/1/3 static-nat 5.6.7.9 host 10.11.14.2**
user@host# **set security nat interface ge-0/0/1/3 static-nat 5.6.8.0 host 10.11.14.3**
user@host# **set security nat interface ge-0/0/1/3 static-nat 5.6.8.1 host 10.11.14.4**

Policies specifying static NAT

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy1 match source-address 13.16.14.16 destination-address user@host static_nat_5.6.7.8_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy1 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy2 match source-address 13.16.14.17 destination-address static_nat_5.6.7.8_32 application custApp2**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy2 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy3 match source-address 13.16.14.16 destination-address static_nat_5.6.7.9_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy3 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy4 match source-address 13.16.14.17 destination-address static_nat_5.6.7.9_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy4 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy5 match source-address 13.16.14.16 destination-address static_nat_5.6.8.0_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy5 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy6 match source-address 13.16.14.17 destination-address static_nat_5.6.8.0_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy6 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy7 match source-address 13.16.14.16 13.16.14.17 destination-address static_nat_5.6.8.1_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy7 then permit**

user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy8 match source-address 13.16.14.17 destination-address static_nat_5.6.8.1_32 application custApp2**
user@host# **set security policies from-zone IntranetALBQ to-zone junos-global policy static-nat-policy8 then permit**

For the statement syntax for configuring policies, see Table 129 on page 206.

Table 135 shows the statement syntax for configuring destination static NAT.

Table 135: Destination Static NAT Configuration Statement

Inbound Services Configuration	Syntax
Maps a single public destination IP address to a single private one without PAT.	`set security nat interface <interface-name> static-nat <ip-prefix> host <ip-prefix> virtual-router <vr-name>` ■ `interface-name`: Name of the interface on which static NAT is to be configured. ■ `static-nat ip-prefix`: Specifies a public IP address to be statically mapped to the private one identified as the host. ■ `host ip-prefix`: Specifies the private IP address of the host to be statically mapped to the public one. ■ `(virtual router) vr-name`: Specifies the virtual router to use to do route lookup for the host address.

About Policy-based Destination NAT

Policy-based destination Network Address Translation (NAT) entails mapping one or more public destination IP addresses to one or more private destination IP addresses with the mapping established in a policy. Not only the public IP address but also the port address can be translated, if port mapping is enabled.

TIP: Because policy-based destination NAT is more extensible than static destination NAT, which is interface-based, it is especially useful for smaller routers.

There are many ways to use policy-based destination NAT. You can translate:

- A single public destination IP address to a single private destination IP address.

  For details, see “Destination NAT: Mapping a Single Public IP Address to a Single Private IP Address” on page 226.

- A range of public destination IP addresses to a range of private IP addresses (such as one subnet to another).

  For details, see “Destination NAT: Mapping Multiple Public IP Addresses to Multiple Private IP Addresses” on page 227

- A single public destination IP address to multiple private destination IP addresses.

  For details, see “Destination NAT: Mapping a Single Public IP Address to Multiple Private IP Addresses” on page 229

- Multiple public destination IP addresses to a single mapped private IP address.

  For details, see “Destination NAT: Mapping Multiple Public IP Addresses to a Single Private IP Address” on page 231

- One public destination IP address to a private one with PAT to a different port number.

For details, see “Destination NAT with Port Mapping” on page 233..

TIP: You can use source and destination NAT in combination.

After your zones and address books are set up and your routes defined, to use destination NAT, you perform the following tasks:

- Define the destination NATs.
- Create policies that specify the destination NAT strategy to be performed on packets that match the policy criteria.

Table 136 gives an overview of the entire policy configuration statement syntax. The action that applies to destination NAT (destination-nat) and its attributes is shown in bold.

Table 136: Policy Configuration Statement

Policy Configuration	Syntax
Configure a policy.	`set security policies from-zone <zone-name> to-zone <zone-name> policy <policy-name> match source-address (<address-name> | <address-set-name>) destination-address (<address-name> | <address-set-name>) application (<application-name> | <application-set-name>)` `set security policies from-zone <zone-name> to-zone <zone-name> policy <policy-name> then (permit | firewall-authentication | tunnel ipsec-vpn <vpn-name> pair-policy <pair-policy-name> | source-nat (pool <pool-name> | pool-set <pool-set-name>) | interface | destination-nat <dest-nat-name> | deny | reject) schedule <scheduler-name> log <session-init session-close> count alarm per-second-threshold <value> per-minute-threshold <value>`  ■ `from-zone zone-name`: Source zone. Name of the zone from which traffic is sent. ■ `to-zone zone-name`: Destination zone. Name of the zone to which traffic is to be sent. ■ `policy-name`: Unique name used to refer to the policy. ■ `(source-address) address-name`: Name of an address (or address set) as entered in the source zone’s address book. ■ `(destination-address) address-name`: Name of an address (or address set) as entered in the destination zone’s address book. ■ `(application) application-name`: Name of a preconfigured or custom application (or application set). ■ action: Any one of the actions listed after the term ‘then’. For destination NAT, the actions is   ■ (`destination-nat`) `dest-nat-name`: The name of the previously configured destination NAT containing the private IP address of the host that the public destination IP address is to be translated to.   ■ `permit` | `deny` | `reject`: Whether a packet matching the policy is transmitted, denied or rejected. ■ `scheduler-name`: Optionally, the name of a scheduler whose schedule determines when the policy is active and when it can be used. ■ `(log) session-init session-close`: Log the traffic that matches the policy. ■ count `values`: Count setters.

Destination NAT: Mapping a Single Public IP Address to a Single Private IP Address

You can configure destination NAT to perform a one-to-one mapping between a public destination IP address and a private one. This kind of policy-based destination NAT is similar to that of a static destination NAT mapping, except that it is configured in a policy and the mapping pertains to a single direction. (For static destination NAT, both the source and destination IP addresses are mapped.)

This sample configuration creates a destination NAT definition containing a single address. It specifies a policy that translates a single public destination IP address to the single private one, for packets that match the policy criteria, without changing the port number.

For the sample configuration, Figure 21 shows how the router maps the original, public IP address to a single private IP address.

Figure 21: Destination NAT Mapping a Single Public IP Address to a Single Private IP Address

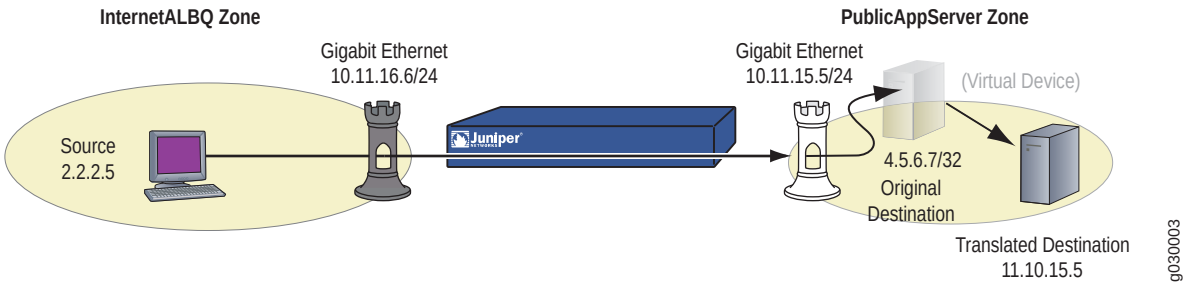

The policy permits SMTP traffic from any address in the InternetALBQ zone to be transmitted to the PubsApps1Mail server in the PublicAppServersALBQ zone with the public IP address single-dest-addr (4.5.6.7) mapped to the private PubsApps1Mail one (11.10.15.5).

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/2.1 and assigns it to the PublicAppServersALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It assigns the PubApps1Mail private IP address to the PublicAppServersALBQ zone's address book.
- It assigns the public destination NAT IP address to the PublicAppServersALBQ zone's address book.
- It creates a destination NAT definition called single-addr for the single PubApps1Mail server address.
- It configures a policy specifying that any SMTP mail addressed to the public destination IP address 4.5.6.7 is to be transmitted to the private destination IP address 11.10.15.5, which is the IP address of the PubApps1Mail server, as specified in the single-addr destination NAT.

Interface and zones configuration

user@host# **set interfaces ge-0/0/2 unit 1 family inet address 10.11.15.5/32**
user@host# set security zones security-zone PublicAppServersALBQ interfaces ge-0/0/2.1

user@host# **set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32**
user@host# **set security zones security-zone InternetALBQ interfaces ge-0/0/3.2**

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignment

user@host# **set security zones security-zone PublicAppServersALBQ address-book address PubApps1Mail 11.10.15.5/32**

user@host# **set security zones security-zone PublicAppServersALBQ address-book address single-dest-addr 4.5.6.7/32**

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Destination NAT single IP address

user@host# **set security nat destination-nat single-addr address 11.10.15.5**

Policy specifying a destination NAT single address

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy single-addr-policy match source-address any destination-address single-dest-addr application smtp**

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy single-addr-policy then permit destination-nat single-addr**

For the statement syntax for configuring policies, see Table 136 on page 225.

Table 137 shows the statement syntax for configuring destination NAT for a single IP address.

Table 137: Destination NAT Configuration Statement for Mapping a Single Private IP Address

Inbound Services Configuration	Syntax
Maps a single public destination IP address to a single private one.	`set security nat destination-nat <dest-nat-name> address <ip-addr>` ■ `dest-nat-name`: Name of the destination NAT definition, to be specified in a policy. ■ `ip-addr`: The private IP address to which the public destination IP address is to be mapped.

Destination NAT: Mapping Multiple Public IP Addresses to Multiple Private IP Addresses

You can use destination NAT to map a range of public destination IP addresses to a range of private destination IP addresses. The range can be a subnet or a range of addresses within a subnet. JUNOS software with enhanced services maintains the relationships among the public destination IP addresses after translating them to the private range.

In this sample configuration, the router maps the public destination IP address range 1.2.2.3 to 1.2.3.9 to a part of the FinancialALBQ subnet.

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/1.1 and assigns it to the FinancialALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It adds the IP address of the financial manager's subnet (10.12.12.0) to the FinancialALBQ zone's address book.
- It adds the IP address range 1.2.2.3 to 1.2.2.6 to the FinancialALBQ zone's address book.
- It configures a policy specifying that any HTTPS traffic addressed to one of six public IP addresses in a range beginning with 1.2.2.3 is to be transmitted to the corresponding private IP address in the range of 10.12.12.1 to 10.12.12.6.

Interfaces and zones configuration

```
user@host# set interfaces ge-0/0/1 unit 1 family inet address 10.12.12.1/24
user@host# set security zones security-zone FinancialALBQ interfaces ge-0/0/1.1

user@host# set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32
user@host# set security zones security-zone InternetALBQ interfaces ge-0/0/3.2
```

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignments

```
user@host# set security zones security-zone FinancialALBQ address-book
address FinancialManagers 10.12.12.0/24

user@host# set security zones security-zone FinancialALBQ address-book
Public-FinMan-range2 address 1.2.2.3/24
```

- For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Destination NAT address range

```
user@host# set security nat destination-nat dest-addr-range addr-range low
10.12.12.1 high 10.12.12.6
```

Policy specifying a destination NAT address range

```
user@host# set security policies from-zone InternetALBQ to-zone FinancialALBQ
policy addr-range-policy match source-address any destination-address
Public-FinMan-range2 application https
user@host# set security policies from-zone InternetALBQ to-zone FinancialALBQ
policy addr-range-policy then permit destination-nat dest-addr-range
```

For the statement syntax for configuring policies, see Table 136 on page 225.

Table 138 shows the statement syntax for configuring a destination IP address range.

Table 138: Destination NAT Configuration Statement for Specifying a Private IP Address Range

Inbound Services Configuration	Syntax
Defines a range of private IP addresses to be mapped to a destination public range.	`set security nat destination-nat <dest-nat-name> address-range low <ip-addr> high <ip-addr> port <port-number>` ■ `dest-nat-name`: Name of the destination NAT mapping to be referred to from within a policy. ■ (address-range) `low`: The beginning address in a range of private IP addresses to which a range of public destination IP addresses is to be mapped. ■ `high`: The ending address in a range of private IP addresses to which a range of public destination addresses is to be mapped.

Destination NAT: Mapping a Single Public IP Address to Multiple Private IP Addresses

You can configure destination NAT to translate a single public destination IP address to different private IP addresses depending on the type of traffic or the source address by creating different destination NAT definitions and policies to address each case.

For example, you might want all HTTP and SMTP traffic from the InternetALBQ zone transmitted to the public destination IP address 3.4.4.5. You might want the router to separate the traffic and redirect SNMP traffic to the PubApps1Mail server and HTTP traffic to the PubApps2Web server.

For this sample configuration, Figure 22 shows how the router maps a single public destination IP address to two private ones.

Figure 22: Mapping a Single Public IP Address to Many Private Ones

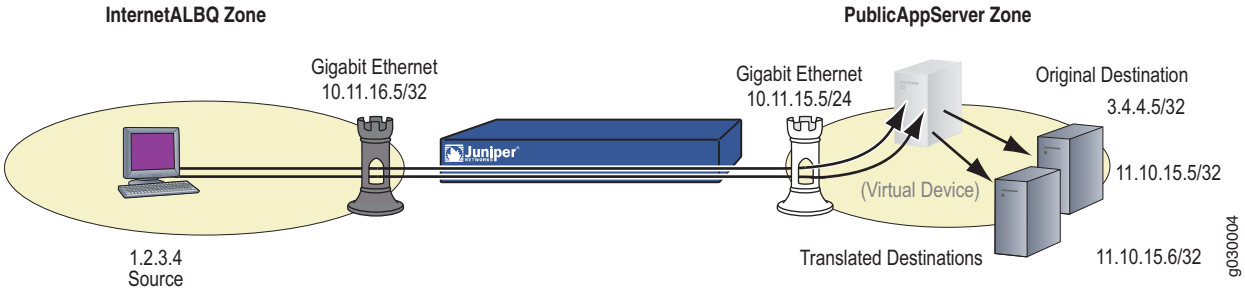

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/2.1 and assigns it to the PublicAppServersALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It adds the PubApps1Mail private IP address to the PublicAppServersALBQ zone's address book.
- It adds the PubApps2Web private IP address to the PublicAppServersALBQ zone's address book.

- It adds the public destination NAT IP address 3.4.4.5 to the PublicAppServerALBQ zone's address book.
- It creates a destination NAT definition called one-to-many-addrMail for the single PubApps1Mail server address.
- It creates a destination NAT definition called one-to-many-addrWeb for the single PubApps2Web server address.
- It configures a policy specifying that any SMTP traffic addressed to the public destination IP address 3.4.4.5 is to be transmitted to the private destination IP address 11.10.15.5, which is the IP address of the PubApps1Mail server, as specified in the one-to-many-addrMail destination NAT.
- It configures a policy specifying that any HTTP traffic addressed to the public destination IP address 3.4.4.5 is to be transmitted to the private destination IP address 11.10.15.6, which is the IP address of the PubApps2Web server, as specified in the one-to-many-addrWeb destination NAT.

Interface and zones configuration

user@host# **set interfaces ge-0/0/2 unit 1 family inet address 10.11.15.5/32**
user@host# set security zones security-zone PublicAppServersALBQ interfaces ge-0/0/2.1

user@host# **set interfaces ge-0/0/3 unit 2 family inet address 10.11.15.6/32**
user@host# **set security zones security-zone InternetALBQ interfaces ge-0/0/3.2**

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignments

user@host# **set security zones security-zone PublicAppServersALBQ address-book address PubApps1Mail 11.10.15.5/32**

user@host# **set security zones security-zone PublicAppServersALBQ address-book address PubApps2Web 11.10.15.6/32**

user@host# **set security zones security-zone PublicAppServersALBQ address-book address one--to-many-dest-addr 3.4.4.5/32**

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Destination NAT private IP addresses

user@host# **set security nat destination-nat one-to-many-addrMail address 11.10.15.5**
user@host# **set security nat destination-nat one-to-many-addrWeb address 11.10.15.6**

Policies specifying one public destination NAT IP address to different private ones

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy one-to-many-mail-policy match source-address any destination-address one--to-many-dest-addr application smtp**
user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy one-to-many-mail-policy then permit destination-nat one-to-many-addrMail**

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy one-to-many-web-policy match source-address any**

destination-address one--to-many-dest-addr application http
user@host# set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy one-to-many-web-policy then permit destination-nat one-to-many-addrWeb

For the statement syntax for configuring policies, see Table 136 on page 225.

Table 139 shows the statement syntax for configuring destination NAT to specify a single private address. To map a single public destination IP address to multiple private ones, you need to define discrete destination-nat definitions for the private IP addresses and use individual policies to direct traffic to them.

Table 139: Destination NAT Configuration Statement for Mapping a Single Public IP Address

Inbound Services Configuration	Syntax
Defines a single private destination IP address to be mapped to a public IP address	`set security nat destination-nat <dest-nat-name> address <ip-addr>` ■ `dest-nat-name`: Name of the destination NAT mapping to be referred to from within a policy. ■ `(address) ip-addr`: A single private IP address to which one or more public destination IP addresses are to be mapped.

Destination NAT: Mapping Multiple Public IP Addresses to a Single Private IP Address

You can configure destination NAT to translate multiple public IP addresses to a single private IP address. In this case, the router forwards traffic from the two or more public destination IP addresses to a single private IP address. Port mapping is optional.

The network administrator of the Albuquerque branch is migrating the network from one address range to another. During this phase, a user might have two public IP addresses—the old address and the new one. To ensure that users get all their communications, the network administrator must map the two public IP addresses to the same private one.

In this sample configuration, traffic addressed to either of the two public addresses 2.4.3.5 or 2.4.3.6 is mapped and transmitted to the single private address of bankMan1 (10.11.14.1).

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/1.3 and assigns it to the BankManagersALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It adds the single private address bankMan1 (10.11.14.1) to the BankManagersALBQ zone.
- It adds the two public destination addresses 2.4.3.5 and 2.4.3.6 to the BankManagersALBQ zone's address book.
- It creates a destination NAT definition called many-to-one-addr for the single bankMan1 (10.11.14.1) address.

- It creates a policy specifying that any HTTP traffic addressed to the public destination IP address 2.4.3.5 is to be transmitted to the private destination IP address 10.11.14.1, which is the IP address of bankMan1, as specified in the many-to-one-addr destination NAT.
- It configures a policy specifying that any HTTP traffic addressed to the public destination IP address 2.4.3.6 is to be transmitted to the private destination IP address 10.11.14.1, which is the IP address of bankMan1, as specified in the many-to-one-addr destination NAT.

Interfaces and zones configuration

user@host# **set interfaces ge-0/0/1 unit 3 family inet address 10.11.14.3/24**
user@host# set security zones security-zone BankManagersALBQ interfaces ge-0/0/1.3

user@host# **set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32**
user@host# **set security zones security-zone InternetALBQ interfaces ge-0/0/3.2**

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book assignments

user@host# **set security zones security-zone BankManagersALBQ address-book address bankMan1 10.11.14.1/32**

user@host# **set security zones security-zone BankManagersALBQ address-book address publicAddr1 2.4.3.5**
user@host# **set security zones security-zone BankManagersALBQ address-book address publicAddr2 2.4.3.6**

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Destination NAT private IP address

user@host# **set security nat destination-nat many-to-one-addr address 10.11.14.1**

Policies specifying two public destination NAT IP addresses mapped to a single private one

user@host# **set security policies from-zone InternetALBQ to-zone BankManagersALBQ policy manyto1-addr1 match source-address any destination-address publicAddr1 application http**
user@host# **set security policies from-zone InternetALBQ to-zone BankManagersALBQ policy manyto1-addr1 then permit destination-nat many-to-one-addr**

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy manyto1-addr2 match source-address any destination-address publicAddr2 application http**
user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy manyto1-addr2 then permit destination-nat many-to-one-addr**

For the statement syntax for configuring policies, see Table 136 on page 225.

Table 140 shows the statement syntax for configuring destination NAT for a single private address. Multiple public destination IP addresses can be mapped to a single one specified in a destination-nat definition.

Table 140: Destination NAT Configuration Statement for a Single Private IP Address

Inbound Services Configuration	Syntax
Defines a single private destination IP address to be mapped to one or more public ones.	`set security nat destination-nat <dest-nat-name> address <ip-addr>` ■ `dest-nat-name`: Name of the destination NAT mapping to be referred to from within a policy. ■ `(address) ip-addr`: A single private IP address to which one or more public destination IP addresses are to be mapped.

Destination NAT with Port Mapping

When you configure the router to perform destination network address translation, you can optionally enable port mapping. One reason to enable port mapping is to support multiple server processes for a single service on a single host.

In this sample configuration, the PubApps2Web server whose IP address 11.10.15.16 is configured to run two Webservers: one at port 80 and one at port 8081.

For HTTP service 1, the router translates the port number of 7 to 80 for HTTP traffic. For HTTP service 2, the router translates the port number in HTTP traffic from 8 to 8081. Because the port addresses are unique, the router can sort HTTP traffic with these port numbers and map each to the appropriate Webserver.

Port mapping translates one public destination port number to a specific private port number. (Recall that PAT translates any public source port number randomly assigned by the initiating host to another port number, randomly assigned by the router.)

This sample configuration performs the following tasks:

- It configures an interface called ge-0/0/2.1 and assigns it to the PublicAppServersALBQ zone.
- It configures an interface called ge-0/0/3.2 and assigns it to the InternetALBQ zone.
- It adds the PubApps2Web private IP address (11.10.15.16) to the PublicAppServersALBQ zone's address book.
- It adds the public destination IP address 5.7.9.1 to the PublicAppServersALBQ zone's address book.
- It creates a destination NAT definition called one-with-port1 for the single PubApps2Web server address (11.10.15.16) with a port number of 80 to receive junos-http traffic.
- It creates a destination NAT definition called one-with-port2 for the single PubApps2Web server address (11.10.15.16) with a port number of 8081 to receive http traffic other than junos-http.

- It configures a policy specifying that any HTTP traffic addressed to the public destination IP address 5.7.9.1 whose packet header specifies a port number of 7) is to be transmitted to the private destination IP address 11.10.15.16 at port 80, which is the IP address and port number of the PubApps2Web server webserver1 handling junos-http traffic, as specified in the one-with- port1 destination NAT.

- It configures a policy specifying that any HTTP traffic addressed to the public destination IP address 5.7.9.1 whose header specifies a port number of 80 is to be transmitted to the private destination IP address for the single PubApps2Web server (11.10.15.16) at port 8081, which is the port number of the PubApps2Web server webserver2, as specified in the one-with- port2 destination NAT.

Interfaces and zones configuration

user@host# **set interfaces ge-0/0/2 unit 1 family inet address 10.11.15.5/32**
user@host# **set security zones security-zone PublicAppServersALBQ interfaces ge-0/0/2.1**

user@host# **set interfaces ge-0/0/3 unit 2 family inet address 10.11.16.5/32**
user@host# **set security zones security-zone InternetALBQ interfaces ge-0/0/3.2**

For details on configuring interfaces and zones, see "Configuring Interfaces and Assigning Them to Zones" on page 49.

Address book

user@host# **set security zones security-zone PublicAppServersALBQ address-book address PubApps2Web 11.10.15.16/32**

user@host# **set security zones security-zone PublicAppServersALBQ address-book address pub--addr-with-portnumber 5.7.9.1/32**

For details on configuring address books, see "Configuring Address Books for Zones" on page 68.

Destination NAT private IP address with different port numbers

user@host# **set security nat destination-nat one-with-port1 address 11.10.15.16 port 80**
user@host# **set security nat destination-nat one-with-port2 address 11.10.15.16 port 8081**

Policies specifying two public destination NAT IP addresses mapped to a single private one

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy addr-with-port1-policy match source-address any destination-address pub--addr-with-portnumber application junos-http**
user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy addr-with-port1-policy then permit destination-nat one-with-port1**

user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy addr-with-port2-policy match source-address any destination-address pub--addr-with-portnumber application http**
user@host# **set security policies from-zone InternetALBQ to-zone PublicAppServersALBQ policy addr-with-port2-policy then permit destination-nat one-with-port2**

For the statement syntax for configuring policies, see Table 136 on page 225.

Table 141 shows the statement syntax for configuring destination NAT for a single private address and port number.

Table 141: Destination NAT Configuration Statement for a Single Private IP Address with Port Number

Inbound Services Configuration	Syntax
Defines a single private destination IP address and port number.	`set security nat destination-nat <dest-nat-name> address <ip-addr> port port-number` ■ `dest-nat-name`: Name of the destination NAT mapping to be referred to from within a policy. ■ `(address) ip-addr`: A single private IP address to which one or more public destination IP addresses are to be mapped. ■ `(port) port-number`: The private port number to which the public destination port number is to be mapped.