J-series™ Services Router

User Guide

Release 7.0

Juniper Networks, Inc.

1194 North Mathilda Avenue

Sunnyvale, California 94089

USA

408-745-2000

www.juniper.net

Part Number: 530-011657-01, Revision 1

J-series™ Services Router User Guide,
Copyright © 2004, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Michael Bushong, Taffy Everts, Walter Goralski, Joshua Kim, Jerry Isaac, Frank Reade, Swapna Steiger, and Alan Twhigg
Editing: Taffy Everts
Illustration: Faith Bradford Brown and Nathaniel Woodward
Cover Design: Edmonds Design

Revision History
4 November 2004—Revision 1.

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.

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 Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller, unless the applicable Juniper documentation expressly permits installation on non-Juniper equipment.

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.

c. Other Juniper documentation for the Software (such as product purchase documents, documents accompanying the product, the Software user manual(s), Juniper's website for the Software, or messages displayed by the Software) may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, concurrent users, sessions, subscribers, nodes, or transactions, or require the purchase of separate licenses to use particular features, functionalities, or capabilities, or provide temporal or geographical limits. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

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, 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; (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, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, 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 Software on non-Juniper equipment where the Juniper documentation does not expressly permit installation 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; or (k) 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.

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. If the Software is distributed on physical media (such as CD), Juniper warrants for 90 days from delivery that the media on which the Software is delivered will be free of defects in material and workmanship under normal use. This limited warranty extends only to the Customer. Except as may be expressly provided in separate documentation from Juniper, no other warranties apply to the Software, and the Software is otherwise provided AS IS. Customer assumes all risks arising from use of the Software. Customer's sole remedy and Juniper's entire liability under this limited warranty is that Juniper, at its option, will repair or replace the media containing the Software, or provide a refund, provided that Customer makes a proper warranty claim to Juniper, in writing, within the warranty period. Nothing in this Agreement shall give rise to any obligation to support the Software. Any such support shall be governed by a separate, written agreement. To the maximum extent permitted by law, Juniper shall not be liable for any liability for 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 HEREIN OR IN SEPARATE DOCUMENTATION PROVIDED FROM JUNIPER AND 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.

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 you may contain encryption or other capabilities restricting your 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. Miscellaneous. This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. 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.

If you have any questions about this agreement, contact Juniper Networks at the following address:

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
USA
Attn: Contracts Administrator

Abbreviated Table of Contents

Table of Contents

About This Guide

This preface provides the following guidelines for using this manual and related Juniper Networks, Inc., technical documents:

Objectives

This guide contains instructions for installing, configuring, and managing a Services Router. It explains how to prepare your site for installation, unpack and install the hardware, power on the router, configure secure routing, monitor network operations, and perform routine maintenance and troubleshooting.

NOTE: This guide documents Release 7.0 of the JUNOS software. For additional information about J-series Services Routers—either corrections to or omissions from this guide—see the J-series release notes at http://www.juniper.net.

J-series Services Routers run on the JUNOS Internet software, which you control through either a Web browser or a command-line interface (CLI) to perform the tasks shown in Table 1.

Table 1: Capabilities of J-series Interfaces

J-series Interface	Capabilities
J-Web graphical browser interface	■ Quick (basic) configuration ■ Monitoring, configuration, diagnosis, and management
JUNOS CLI	Monitoring, configuration, diagnosis, and management

This guide provides complete instructions for using the J-Web interface, but it is not a comprehensive resource for using the JUNOS CLI. For CLI information, see "Related Juniper Networks Documentation" on page xxviii.

Audience

This guide is designed for anyone who installs, configures, and maintains a J-series Services Router or prepares a site for Services Router installation. The guide is intended for the following audiences:

- Customers with technical knowledge of and experience with networks and the Internet
- Network administrators who install, configure, and manage Internet routers but are unfamiliar with the JUNOS software
- Network administrators who install, configure, and manage products of Juniper Networks

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

How to Use This Guide

Because you can configure and manage a J-series Services Router in several ways, most chapters in this guide contain multiple sets of instructions:

- Configuration—For many Services Router features, you can use J-Web Quick Configuration for basic setup. For more extensive configuration of all Services Router features, use the J-Web configuration editor or the JUNOS CLI configuration editor.
- Maintenance—To monitor, diagnose, and manage a Services Router, use the J-Web interface for common tasks, or use CLI operational mode commands.

Table 2 shows where Quick Configuration, J-Web, and CLI instructions are located.

Table 2: Location of Instructions in a Chapter

Configuration or Management Method	Location of Instructions in a Chapter
J-Web Quick Configuration pages (where applicable)	In a table, before configuration editor instructions.
■ J-Web configuration editor pages ■ JUNOS CLI configuration editor (configuration mode statements)	Together in a task table—after Quick Configuration instructions (where applicable).
■ J-Web monitor, diagnose, and manage pages	In a verification section at the end of a configuration chapter.
■ JUNOS CLI operational mode commands	Information about common monitoring and diagnostic tasks is located in "Monitoring and Diagnosing a Services Router" on page 197.

Document Conventions

Table 3 defines notice icons used in this guide.

Table 3: Notice Icons

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

Table 4 defines the text and syntax conventions used in this guide.

Table 4: Text and Syntax Conventions

Convention	Description	Examples
Bold sans serif typeface	Represents text that you type.	To enter configuration mode, type the configure command: user@host> **configure**
Italic typeface	■ 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*

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

Related Juniper Networks Documentation

Although this guide provides instructions for configuring and managing a J-series Services Router with the JUNOS CLI, it is not a comprehensive JUNOS software resource. For complete documentation of the statements and commands described in this guide, see the JUNOS software manuals listed in Table 5.

Table 5: Related JUNOS Software Publications

Chapter in This Guide	Corresponding JUNOS Software Manual
Part 2, "Installing the J-series Services Router"	
"Configuring Network Interfaces" on page 79	■ *JUNOS Network Interfaces and Class of Service Configuration Guide* ■ *JUNOS Network and Services Interfaces Command Reference*
Part 3, "Using the J-series User Interfaces"	
"J-series User Interface Overview" on page 109	*JUNOS System Basics Configuration Guide*
"Using J-series Configuration Tools" on page 127	*JUNOS System Basics Configuration Guide*
Part 4, "Managing the Services Router"	
"Managing Users and Operations" on page 163	*JUNOS System Basics Configuration Guide*
"Monitoring and Diagnosing a Services Router" on page 197	■ *JUNOS Protocols, Class of Service, and System Basics Command Reference* ■ *JUNOS Network and Services Interfaces Command Reference*
"Configuring SNMP for Network Management" on page 241	*JUNOS Network Management Configuration Guide*
Part 5, "Configuring Routing Protocols"	
"Routing Overview" on page 255 "Configuring Static Routes" on page 285 "Configuring a RIP Network" on page 297 "Configuring an OSPF Network" on page 309 "Configuring BGP Sessions" on page 331	*JUNOS Routing Protocols Configuration Guide*
Part 6, "Configuring Routing Policy, Firewall Filters, and Class of Service"	
"Policy, Firewall Filter, and Class-of-Service Overview" on page 351 "Configuring Routing Policies" on page 375	*JUNOS Policy Framework Configuration Guide*
"Configuring Firewall Filters and NAT" on page 389	■ *JUNOS Network Interfaces and Class of Service Configuration Guide* ■ *JUNOS Policy Framework Configuration Guide* ■ *JUNOS Services Interfaces Configuration Guide*
"Configuring Class of Service with DiffServ" on page 427	*JUNOS Network Interfaces and Class of Service Configuration Guide*
Part 7, "Managing Multicast Transmissions"	
"Multicast Overview" on page 461 "Configuring a Multicast Network" on page 471	*JUNOS Multicast Protocols Configuration Guide*

Chapter in This Guide	Corresponding JUNOS Software Manual
Part 8, "Managing Packet Security"	
"Configuring IPSec for Secure Packet Exchange" on page 483	■ *JUNOS System Basics Configuration Guide* ■ *JUNOS Services Interfaces Configuration Guide*

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at 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 Support

For technical support, open a support case using the Case Manager link at http://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (outside the United States).

Part 1
J-series Overview

Chapter 1

Introducing the J-series Services Router

J-series Services Routers provide stable, reliable, efficient IP routing, WAN and LAN connectivity, and management services for small to medium-sized enterprise networks. Services Routers typically connect small, branch, or regional offices to a central site router, and link Internet service provider (ISP) networks.

This chapter contains the following topics:

J-series Services Router Overview

J-series Services Routers are available in three models of increasing bandwidth, described in Table 6.

All J-series Services Routers run on the JUNOS software and are reachable through the J-Web browser interface on the JUNOS command-line interface. For details, see "J-series User Interface Overview" on page 109.

Table 6: J-series Models

Model	Description	Bandwidth
J2300 Services Router	Remote or branch office customer premises equipment (CPE). Smaller chassis (1 U) with a nonredundant AC power supply, 256 MB to 512 MB of memory, and a Universal Serial Bus (USB) port for external storage. Three available versions have two Fast Ethernet LAN interfaces plus one of the following sets of fixed WAN interfaces: ■ Dual T1 interfaces ■ Dual E1 interfaces ■ Two synchronous serial ports	Up to 4 Mbps 20,000 to 50,000 packets per second (pps)

Model	Description	Bandwidth
J4300 Services Router	Regional office CPE. Larger chassis (2 U) with a nonredundant AC power supply, 256 MB to 512 MB of memory, and a Universal Serial Bus (USB) port for external storage. In addition to two Fast Ethernet LAN interfaces, this model has six open slots for the following WAN Physical Interface Modules (PIMs): ■ 2–port Fast Ethernet PIM ■ 2–port T1 or E1 PIM ■ 2–port Serial PIM	Up to 16 Mbps 50,000 to 80,000 pps
J6300 Services Router	Corporate CPE. Larger chassis (2 U) with a redundant AC power supply, 256 MB to 1 GB of memory, and a Universal Serial Bus (USB) port for external storage. In addition to two Fast Ethernet LAN interfaces, this model has six open slots for the following WAN Physical Interface Modules (PIMs): ■ 2–port Fast Ethernet PIM ■ 2–port T1 or E1 PIM ■ 2–port Serial PIM ■ 1–port DS3 PIM	Up to 90 Mbps 100,000 to 150,000 pps

J-series Software Features and Licenses

J-series Services Routers provide the software features listed in Table 7. You must purchase a separate software license to obtain some software features.

Table 7: Summary of J-series Features and License Requirements

Feature Category	J-series Feature	Separate License
Internet Protocols	IPv4 only	
Routing and Multicast	Open Shortest Path First (OSPF)	
	Border Gateway Protocol (BGP)	License required for advanced BGP
	Routing Information Protocol version 1 (RIPv1) and RIPv2	
	Static routes	

Feature Category	J-series Feature	Separate License
	Intermediate System-to-Intermediate System (IS-IS)	
	Multicast: ■ Internet Group Management Protocol version 3 (IGMPv3) ■ Protocol Independent Multicast (PIM) ■ Distance Vector Multicast Routing Protocol (DVMRP) ■ Single-source multicast	
IP Address Management	Static addresses	
Encapsulation	Ethernet: ■ Media access control (MAC) encapsulation ■ 802.1p tagging	
	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 EIA-530 connections	
	802.1Q filtering and forwarding	
	Multilink Frame Relay	
	Multilink PPP	
Traffic Management	Policing and shaping	
	Class-based queuing with prioritization	
	Weighted random early detection (WRED)	
	Queuing by virtual LAN (VLAN), data link connection identifier (DLCI), interface, or bundle	
Security	Network attack detection	
	Denial-of-service (DoS) and distributed DoS protection	
	Generic routing encapsulation (GRE), IP-in-IP, and IP Security (IPSec) tunnels	License required for IPSec
	56-bit Data Encryption Standard (DES) and 168-bit 3DES encryption	
	MD5 and Secure Hash Algorigthm (SHA-1) authentication	
	Replay attack prevention	
	Stateful firewall packet filters	License required
Voice Support	Compressed Real-time Transport Protocol (CRTP)	
High Availability	Virtual Router Redundancy Protocol (VRRP)	
	Graceful restart according to IETF standards	
	Redundant interfaces	
System Management	JUNOScope network manager	

Feature Category	J-series Feature	Separate License
	J-Web browser interface—for Services Router configuration and management	
	JUNOScript XML application programming interface (API)	
	JUNOS command-line interface (CLI)—for Services Router configuration and management through the console, telnet, or SSH	
	Simple Network Management Protocol version 1 (SNMPv1) and SNMPv2	
Traffic Analysis	J-Flow flow monitoring and accounting	License required for J-Flow
Activity Logging and Monitoring	System log	
	Traceroute	
Administration	Supports the following external administrator databases: ■ RADIUS ■ Lightweight Directory Access Protocol (LDAP) ■ SecurID	
	Autoinstallation	
	Configuration rollback	
	Button-operated configuration rescue (CONFIG)	
	Confirmation of configuration changes	
	Software upgrades	

Chapter 2
System Overview

J-series Services Routers are available in three models.

This chapter contains the following topics:

J2300 Services Router Hardware Features

This section contains the following topics:

J2300 Chassis

The J2300 Services Router chassis is a rigid sheet metal structure that houses all the other router components (see Figure 1, Figure 2, and Figure 3). The chassis can be installed in many types of racks or cabinets, on a wall, or on a desk. For information about acceptable rack types, see “Rack Requirements” on page 542.

In addition to the features described in subsequent sections, the chassis includes the following features (see Figure 1 and Figure 2):

- One pair of metal brackets that can be attached to the side of the chassis. You can use the brackets for mounting the chassis in a rack or cabinet or on a wall.
- One electrostatic discharge (ESD) point, a PEM nut at the rear of the chassis.

WARNING: Before removing or installing components of a functioning router, attach an ESD strap to an ESD point and place the other end of the strap around your bare wrist. Failure to use an ESD strap could result in damage to the router.

The router is connected to earth ground through the AC power cord. The router must be connected to earth ground during normal operation.

For additional safety information, see "Safety and Regulatory Compliance Information" on page 563.

Figure 1: Front of J2300 Chassis

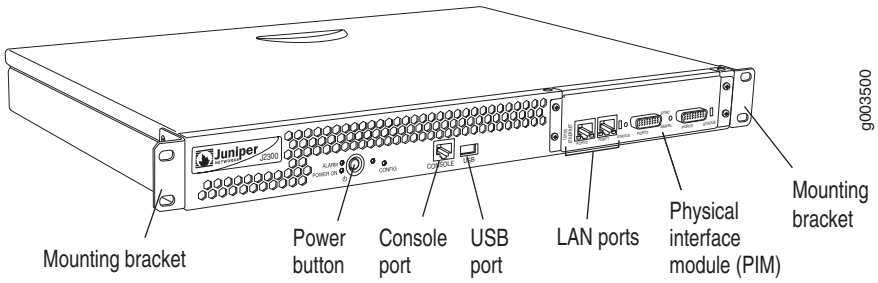

Figure 2: Rear of J2300 Chassis

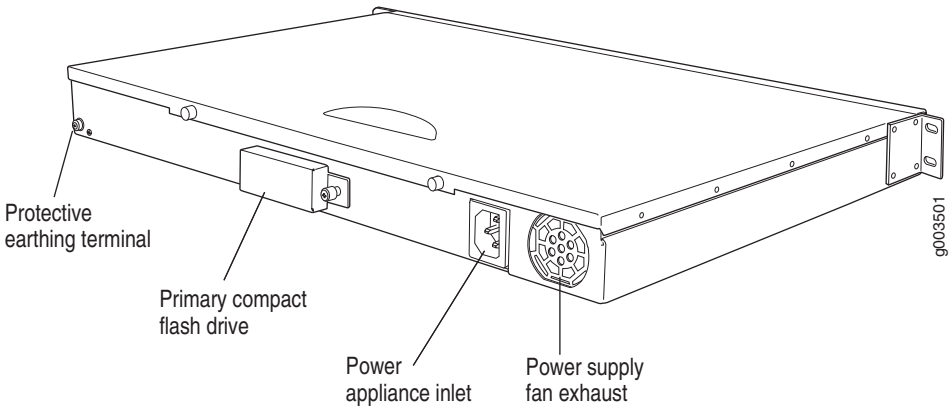

Figure 3: J2300 Hardware Components

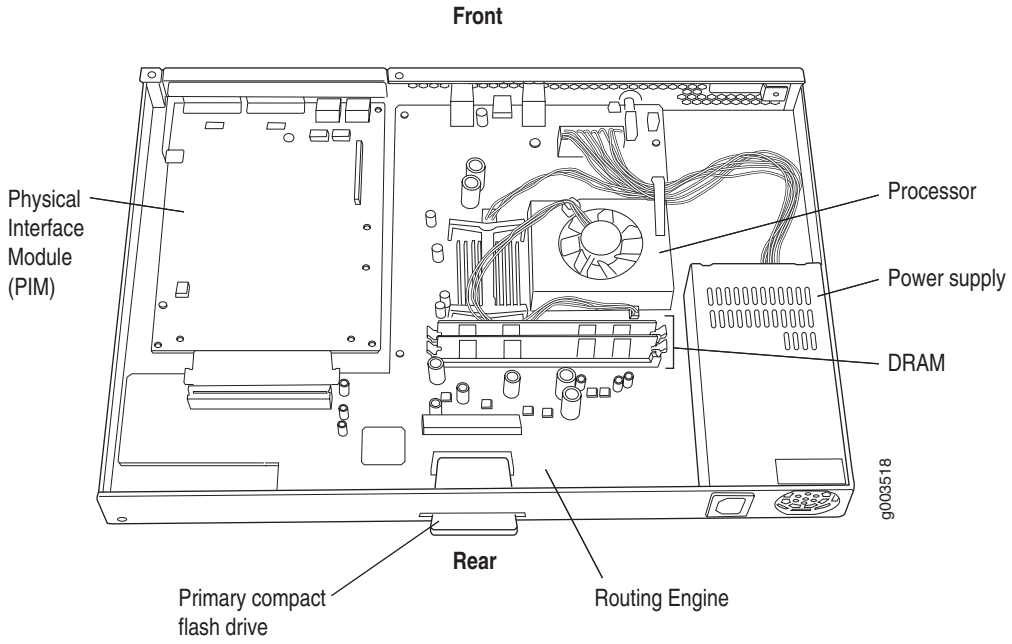

Table 8 summarizes the physical specifications for the router chassis.

Table 8: J2300 Physical Specifications

Description	Value
Chassis dimensions	■ 1.75 in. (4.4 cm) high ■ 17.25 in. (43.8 cm) wide—19 in. (48.3 cm) wide with mounting brackets attached ■ 12.37 in. (31.4 cm) deep—plus 0.5 in. (1.27 cm) of hardware that protrudes from the chassis front
Router weight	12 lb (5.4 kg)

J2300 Routing Engine

The Routing Engine provides three main functions:

- Creates the packet forwarding switch fabric for the Services Router, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network.
- Maintains the routing tables used by the router and controls the routing protocols that run on the router.
- Provides control and monitoring functions for the router, including controlling power and monitoring system status.

The Routing Engine consists of the following components:

- Processor—Creates the packet forwarding switch fabric for the router and runs JUNOS Internet software to maintain the router's routing tables and routing protocols. The Routing Engine has a Pentium-class processor.
- DRAM—Buffers incoming packets and provides storage for the routing and forwarding tables and for other Routing Engine processes.
- Compact flash drive—Provides primary storage for software images, configuration files, and microcode. The compact flash drive is accessible from the rear of the router, and is field-replaceable. For information about replacing the compact flash drive, see "Removing and Installing the Primary Compact Flash Disk" on page 523.
- PCI bus—Provides the interface to the PIMs.
- EPROM—Stores the serial number of the Routing Engine.

NOTE: For specific information about Routing Engine components (for example, the amount of DRAM installed), issue the show chassis routing-engine command.

J2300 Boot Devices

The J2300 Services Router can boot from two devices:

- Primary compact flash disk
- USB drive

J2300 Boot Sequence

Normally, the Services Router boots from the primary compact flash disk. If the compact flash disk fails, the router attempts to boot from the removable USB drive, if present, which is the alternate boot device.

J2300 Front Panel

The front panel of the Services Router (Figure 4) allows you to view router status LEDs, access the console port, and perform simple control functions.

Figure 4: Front Panel of J2300 Services Router

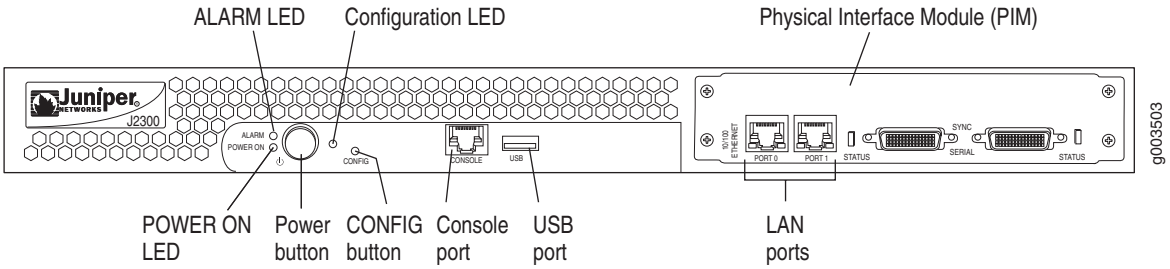

For information about the components of the front panel, see the following sections:

- ALARM LED on page 11
- Power Button and POWER ON LED on page 12
- CONFIG Button and LED on page 12
- Console Port on page 13
- J2300 USB Port on page 13

ALARM LED

The ALARM LED is located to the left of the power button on the front panel (see Figure 4). The yellow (amber) LED lights to indicate a critical condition that can result in a system shutdown or a less severe condition that requires monitoring or maintenance.

NOTE: The ALARM LED on the Services Router is a single-color alarm regardless of the severity of the alarm condition (critical, major, or minor). When an alarm condition triggers the LED, the yellow light turns on.

To deactivate alarms, you must clear the condition that caused the alarm. For a list of alarms that can occur on the router, see “Chassis Alarm Conditions” on page 536.

Power Button and POWER ON LED

The power button is located on the left side of the front panel (see Figure 4). You can use the power button to power the Services Router on and off. When you power on the router, the Routing Engine boots as the power supply completes its startup sequence.

The POWER ON LED is located to the left of the power button on the front panel. Table 9 describes the POWER ON LED.

Table 9: POWER ON LED

Color	State	Description
Green	Off	Router is unplugged, or is powered off and in standby mode.
	On steadily	Router is powered on and is either booting or functioning normally.
	Blinking	Power button has been pressed and quickly released, and the router is gracefully shutting down.

After the router is powered on, status indicators—such as LEDs on the front panel and show chassis command output—can take up to 60 seconds to indicate that the power supply is functioning normally. Ignore error indicators that appear during the first 60 seconds.

If you need to power off the router after the Routing Engine finishes booting, use the J-Web interface or the CLI to halt the Services Router first. For instructions, see “Rebooting or Halting a Services Router with the J-Web Interface” on page 512.

CONFIG Button and LED

You can use the CONFIG button to return the router to a configuration that you have determined is a stable, known configuration. The CONFIG button is recessed to prevent it from being pressed accidentally.

- When you press and release the CONFIG button, the rescue configuration is loaded and committed.
- When you press and hold the CONFIG button for more than 15 seconds, all configurations on the router (including the rescue configuration and backup configurations) are deleted, and the factory configuration is loaded and committed.

Table 10 describes the configuration LED.

Table 10: Configuration LED

Color	State	Description
Green	Blinking	Rescue configuration is being loaded.
	On steadily	Rescue or factory configuration is loaded and committed.
Red	Blinking	■ Current committed configuration and all previous versions are being deleted. ■ Factory configuration is being loaded.
	On steadily	Operation to return the router to the rescue or factory configuration failed.

Console Port

You can use the console port to connect to the Routing Engine through an RJ-45 serial cable. From the console port, you can use the CLI to configure the router. The console port is configured as data terminal equipment (DTE) and supports the RS-232 (EIA-232) standard.

J2300 USB Port

The slot labeled USB on the front panel of the router (see Figure 4) accepts a USB drive or USB drive adapter with a compact flash disk installed, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the USB drive is installed and configured, it automatically acts as a secondary boot device, if the primary compact flash disk fails on startup. Depending on the size of the USB drive, you can also configure it to receive any core files generated during a failure. For information about configuring a USB drive, see "Configuring Boot Devices with the CLI" on page 508.

NOTE: For a list of supported USB drives, see the J-series release notes at http://www.juniper.net.

J2300 Physical Interface Module (PIM)

The fixed Physical Interface Modules (PIM) in a J2300 Services Router provide the physical connection to various network media types, receiving incoming packets from the network and transmitting outgoing packets to the network. The PIM is equipped with a dedicated network processor that forwards incoming data packets to the Routing Engine, and receives outgoing data packets from the Routing Engine. During this process, the PIM performs framing and line-speed signaling for its medium type.

Each PIM supported on the router has the following components:

- One or more cable connector ports—Accept a network media connector.
- Status LED—Indicates port status. Table 11 describes the meaning of the LED states.

For pinouts of PIM cable connectors, see "Network Cable Specifications and Connector Pinouts" on page 551. For PIM replacement instructions, see "Replacing a PIM" on page 518.

Table 11: PIM Status LED

Color	State	Description
Green	On steadily	Online with no alarms or failures.
Red	On steadily	Active with a local alarm; router has detected a failure.

J2300 LAN Ports

All J-series Services Routers include two fixed 10/100Base-TX Fast Ethernet ports. The LAN ports receive incoming packets from the network and transmit outgoing packets to the network. Each port is equipped with a dedicated network processor that forwards incoming data packets to the Routing Engine, and receives outgoing data packets from the Routing Engine.

The LAN ports are located on the front panel of the router (see Figure 4) and are configured like the ports on a Physical Interface Module (PIM). The LAN ports are not field-replaceable. The ports, labeled PORT 0 and PORT 1, correspond to fe-0/0/0 and fe-0/0/1 respectively, for configuration.

For pinouts of Fast Ethernet cable connectors, see "Network Cable Specifications and Connector Pinouts" on page 551.

Each port has two LEDs located on each side of the bottom of the port. Table 12 describes the LAN port LEDs.

Table 12: LAN Port LEDs

Function	Color	State	Description
Link	Green	On steadily	Port is online.
Activity	Green	Blinking	Port is receiving data.
		Off	Port might be on, but is not receiving data.

J2300 Power System

The J2300 Services Router uses AC power. The autosensing power supply (see Figure 2) distributes the different output voltages to the router components according to their voltage requirements.

The power supply is fixed in the chassis, and is not field-replaceable. It has a single AC appliance inlet that requires a dedicated AC power feed.

For information about site power preparations, see "Power Guidelines, Requirements, and Specifications" on page 545. For information about connecting the router to power and ground, see "Connecting Power to the Services Router" on page 43.

J2300 Cooling System

The cooling system consists of the following components:

- A fan on the Routing Engine's processor
- A fan on the power supply

The airflow produced by these fans keeps router components within the acceptable temperature range (see Figure 5).

Figure 5: Airflow Through the J2300 Chassis

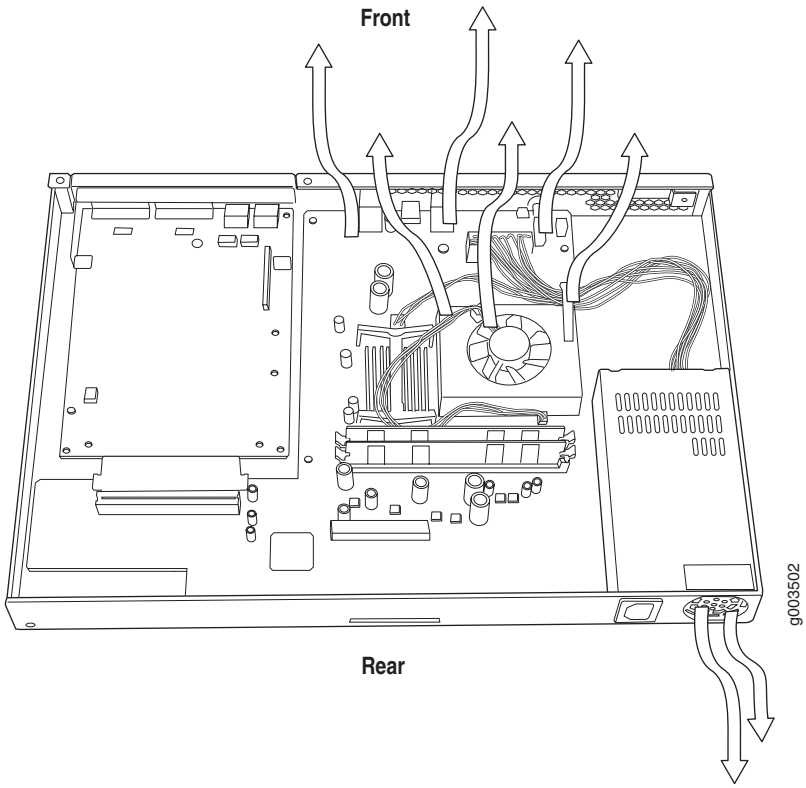

The Routing Engine monitors the temperature of the router components. If the ambient maximum temperature specification is exceeded and the router cannot be adequately cooled, the Routing Engine shuts down the hardware components.

J4300 and J6300 Services Router Hardware Features

This section contains the following topics:

- J4300 and J6300 Cooling System on page 27

J4300 and J6300 Chassis

The J4300 and J6300 Services Router chassis is a rigid sheet metal structure that houses all the other router components (see Figure 6, Figure 7, Figure 8, and Figure 9). The chassis can be installed in many types of racks or cabinets. For information about acceptable rack types, see "Rack Requirements" on page 542.

In addition to the features described in subsequent sections, the chassis includes the following features (see Figure 6, Figure 7, and Figure 8:

- One pair of metal brackets attached to the side of the chassis. You can use the brackets for mounting the chassis in a rack or cabinet.
- One electrostatic discharge (ESD) point, a banana plug receptacle at the front of the chassis.

WARNING: Before removing or installing components of a functioning router, attach an ESD strap to the ESD point and place the other end of the strap around your bare wrist. Failure to use an ESD strap could result in damage to the router.

The router is connected to earth ground through the AC power cord. The router must be connected to earth ground during normal operation.

For additional safety information, see "Safety and Regulatory Compliance Information" on page 563.

- One protective earthing terminal, a PEM nut at the rear of the chassis.

Figure 6: Front of J4300 and J6300 Chassis

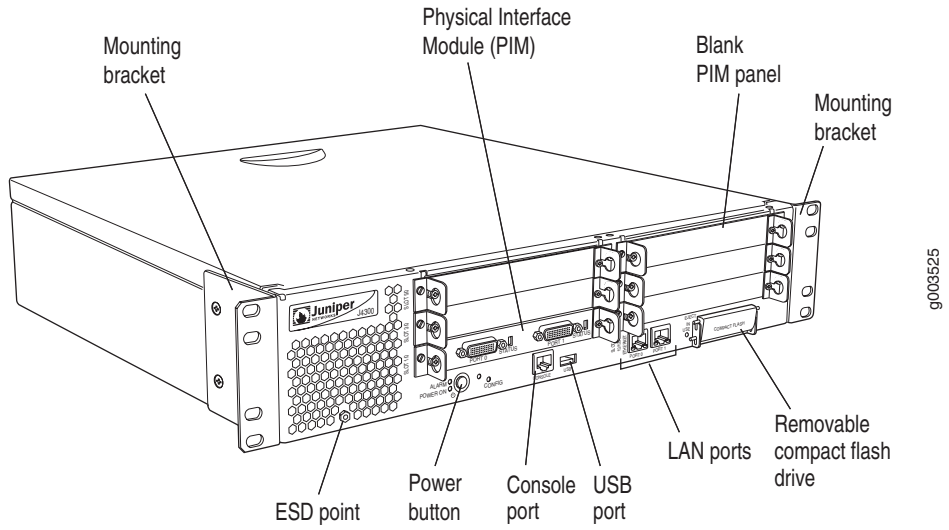

Figure 7: Rear of J4300 Chassis

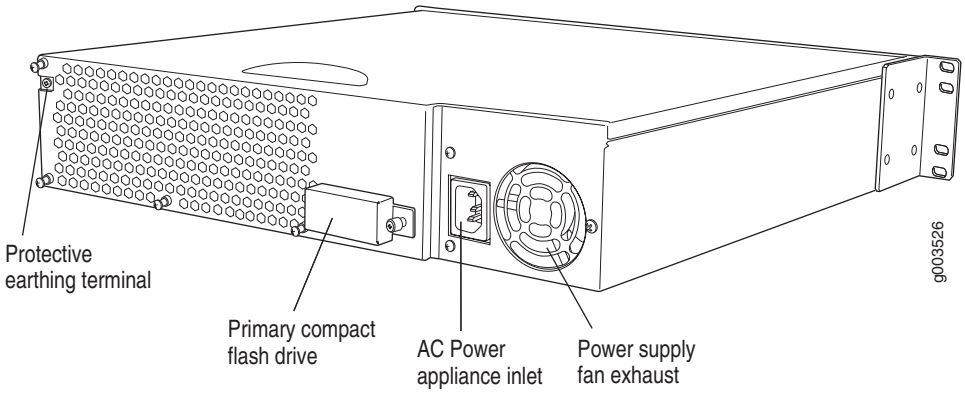

Figure 8: Rear of J6300 Chassis

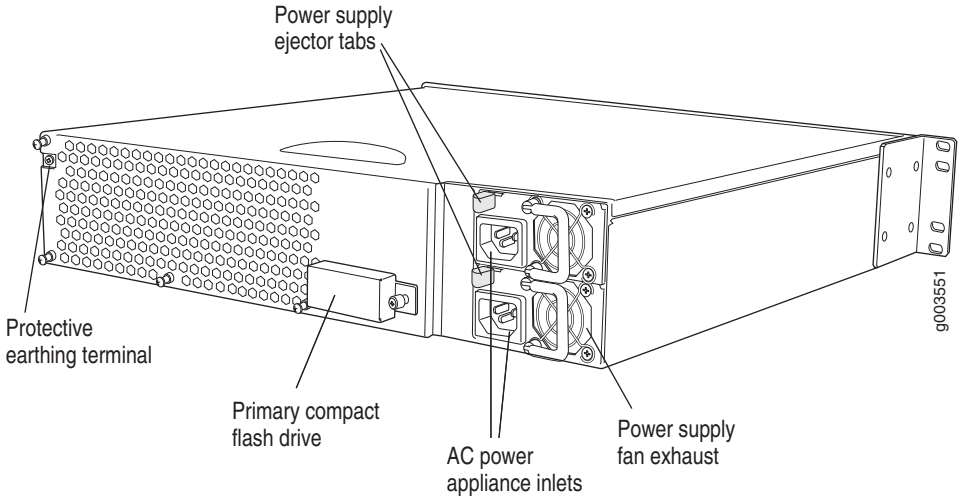

Figure 9: J4300 and J6300 Hardware Components

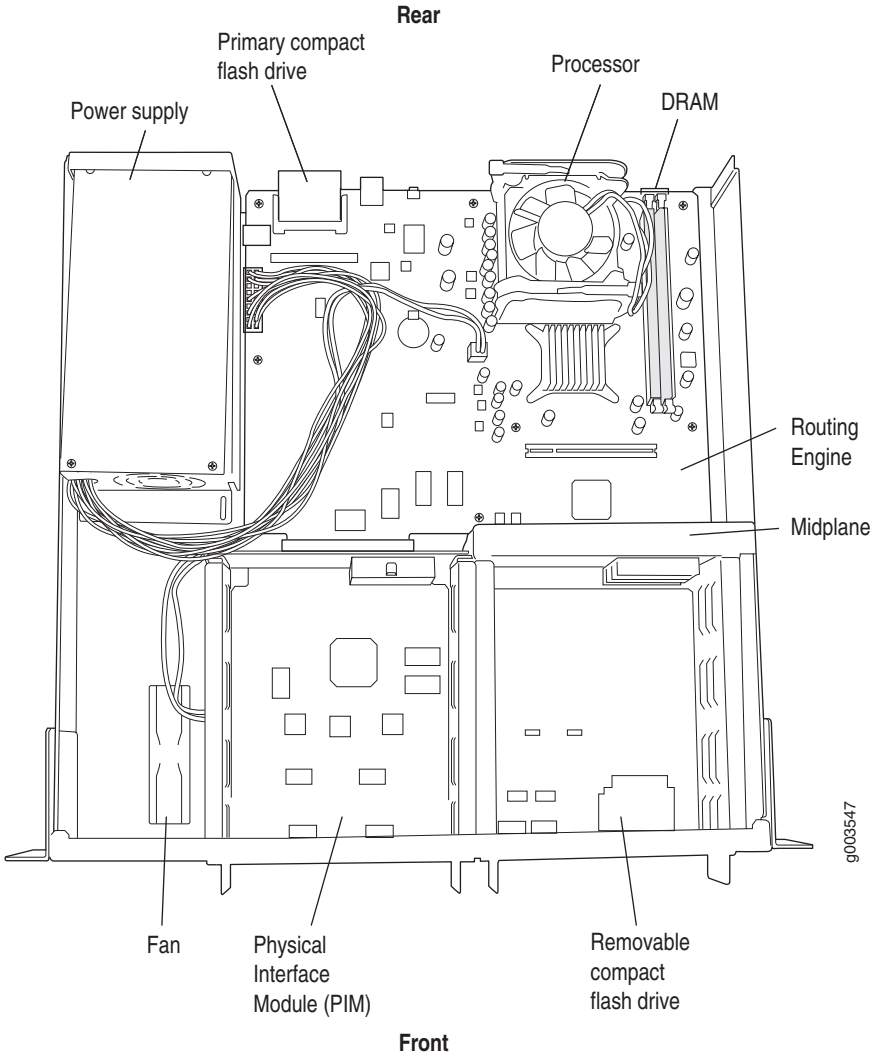

Table 13 summarizes the physical specifications for the router chassis.

Table 13: J4300 and J6300 Physical Specifications

Description	Value
Chassis dimensions	■ 3.50 in. (8.9 cm) high ■ 17.00 in. (43.2 cm) wide—19 in. (48.3 cm) wide with mounting brackets attached ■ 19.00 in. (48.3 cm) deep—plus 0.5 in. (1.27 cm) of hardware that protrudes from the chassis front
Router weight	■ J4300 router minimum configuration (no PIMs): 18 lb (8.2 kg) ■ J4300 router maximum configuration (six PIMs): 21 lb (9.5 kg) ■ J6300 router minimum configuration (no PIMs and one power supply): 18.5 lb (8.4 kg) ■ J6300 router maximum configuration (six PIMs and two power supplies): 24 lb (10.9 kg)

Midplane

The midplane is located in the center of the chassis and forms the rear of the PIM card cage (see Figure 9). You install the PIMs into the midplane from the front of the chassis. Data packets are transferred across the midplane from the PIM to the Routing Engine, and from the Routing Engine across the midplane to the destination PIM.

J4300 and J6300 Routing Engine

The Routing Engine provides three main functions:

- Creates the packet forwarding switch fabric for the Services Router, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network.
- Maintains the routing tables used by the router and controls the routing protocols that run on the router.
- Provides control and monitoring functions for the router, including controlling power and monitoring system status.

The Routing Engine consists of the following components:

- Processor—Creates the packet forwarding switch fabric for the router and runs JUNOS Internet software to maintain the router's routing tables and routing protocols. The Routing Engine has a Pentium-class processor.
- DRAM—Buffers incoming packets and provides storage for the routing and forwarding tables and for other Routing Engine processes.
- Compact flash drive—Provides primary storage for software images, configuration files, and microcode. The compact flash drive is accessible from the rear of the router, and is field-replaceable. For information about

replacing the compact flash drive, see “Removing and Installing the Primary Compact Flash Disk” on page 523.

- PCI bus—Provides the interface to the PIMs.
- EPROM—Stores the serial number of the Routing Engine.

NOTE: For specific information about Routing Engine components (for example, the amount of DRAM installed), issue the show chassis routing-engine command.

J4300 and J6300 Boot Devices

The J4300 and J6300 Services Routers can boot from three devices:

- Primary compact flash disk
- Removable compact flash disk
- USB drive

J4300 and J6300 Boot Sequence

Normally, the Services Router boots from the primary compact flash disk. If the compact flash disk fails, the router attempts to boot from the removable compact flash disk, which is the alternate boot device. If the removable compact flash disk is not present or fails, the router attempts to boot from the USB drive.

J4300 and J6300 Front Panel

The front panel of a J4300 or J6300 Services Router (Figure 10) allows you to view router status LEDs, access the console port, connect to LAN ports, and perform simple control functions.

Figure 10: Front Panel of J4300 and J6300

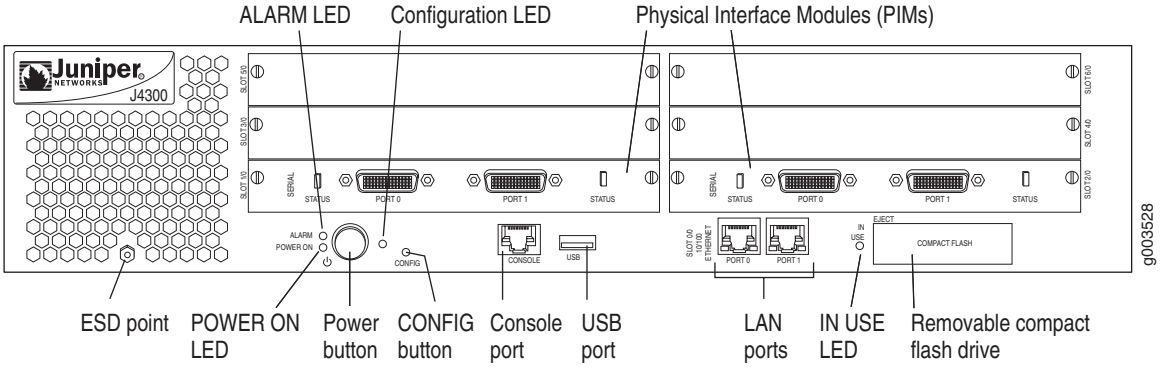

The components of the front panel, from left to right, are described in the following sections:

- ALARM LED on page 23
- Power Button and POWER ON LED on page 23
- CONFIG Button and Configuration LED on page 23
- Console Port on page 24
- J4300 and J6300 USB Port on page 24
- J4300 and J6300 LAN Ports on page 24
- J4300 and J6300 Removable Compact Flash Drive on page 25

ALARM LED

The ALARM LED on J4300 and J6300 Services Routers functions identically to the ALARM LED on the J2300 Services Router. See "ALARM LED" on page 11.

Power Button and POWER ON LED

The power button and POWER ON LED on J4300 and J6300 Services Routers function identically to the power button and POWER ON LED on the J2300 Services Router. See "Power Button and POWER ON LED" on page 12.

CONFIG Button and Configuration LED

The CONFIG button and LED on J4300 and J6300 Services Routers function identically to the CONFIG button and configuration LED on the J2300 Services Router. See "CONFIG Button and LED" on page 12.

Console Port

The console port on J4300 and J6300 Services Routers functions identically to the console port on the J2300 Services Router. See “Console Port” on page 13.

J4300 and J6300 USB Port

The slot labeled USB on the front panel of the router (see Figure 10) accepts a USB drive or USB drive adapter with a compact flash disk installed, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the USB drive is installed and configured, it automatically acts as a secondary boot device, if the primary or removable compact flash disk fails on startup. Depending on the size of the USB drive, you can also configure it to receive any core files generated during a failure. For information about configuring a USB drive, see “Configuring Boot Devices with the CLI” on page 508.

NOTE: For a list of supported USB drives, see the J-series release notes at http://www.juniper.net.

J4300 and J6300 LAN Ports

All J-series Services Routers include two fixed 10/100Base-TX Fast Ethernet ports. The LAN ports receive incoming packets from the network and transmit outgoing packets to the network. Each port is equipped with a dedicated network processor that forwards incoming data packets to the Routing Engine, and receives outgoing data packets from the Routing Engine.

The LAN ports are located on the front panel of the router (see Figure 10) and are configured like the ports on a Physical Interface Module (PIM). The LAN ports are not field-replaceable. The ports, labeled PORT 0 and PORT 1, correspond to fe-0/0/0 and fe-0/0/1 respectively, for configuration.

For pinouts of Fast Ethernet cable connectors, see “Network Cable Specifications and Connector Pinouts” on page 551.

Each port has two LEDs located on each side of the bottom of the port. Table 14 describes the LAN port LEDs.

Table 14: J4300 and J6300 LAN Port LEDs

Function	Color	State	Description
Link	Green	On steadily	Port is online.
Activity	Green	Blinking	Port is receiving data.
		Off	Port might be on, but is not receiving data.

J4300 and J6300 Removable Compact Flash Drive

The slot labeled COMPACT FLASH on the front panel of the Services Router (see Figure 10) is a removable compact flash drive that accepts a type I or II compact flash disk, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the removable compact flash disk is installed and configured, it automatically acts as the secondary boot device if the primary compact flash drive fails on startup.

Depending on the capacity of the removable compact flash disk, you can also configure it to receive any core files generated during a failure. For information about configuring a removable compact flash disk, see "Configuring Boot Devices with the CLI" on page 508.

The IN USE LED indicates that the removable compact flash is being accessed. Table 15 describes the meaning of the LED states.

Table 15: IN USE LED

Color	State	Description
Red	On steadily	■ Router has booted from the removable compact flash drive. ■ request system snapshot operation has been executed, and files are being copied to or from the removable compact flash drive. ■ Core dump of the kernel is being written to the removable compact flash drive. ■ savecore process is retrieving core dump information.

J4300 and J6300 Physical Interface Modules (PIMs)

Physical Interface Modules (PIMs) provide the physical connection to various network media types, receiving incoming packets from the network and transmitting outgoing packets to the network (see Figure 11). Each PIM is equipped with a dedicated network processor that forwards incoming data packets to the Routing Engine, and receives outgoing data packets from the Routing Engine. During this process, the PIM performs framing and line-speed signaling for its medium type.

Figure 11: PIM

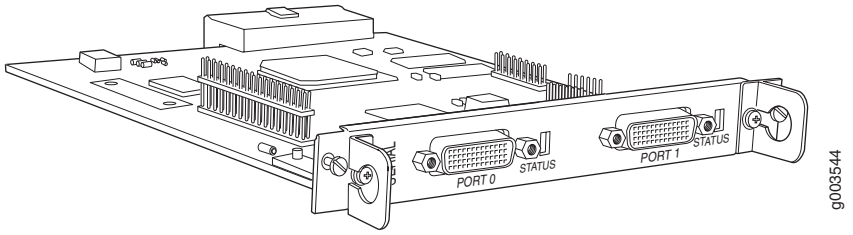

PIMs are removable and insertable when the router is powered off. You can install a PIM into one of the six slots in the router chassis. If a slot is not occupied by a PIM, a PIM blank panel must be installed to shield the empty slot and to allow cooling air to circulate properly through the router.

- One or more cable connector ports—Accept a network media connector.
- LED—Indicates port status. Table 11 describes the meaning of the LED states.

For pinouts of PIM cable connectors, see "Network Cable Specifications and Connector Pinouts" on page 551. For PIM replacement instructions, see "Replacing a PIM" on page 518.

J4300 Power System

The J4300 Services Router uses AC power. The autosensing power supply (see Figure 7) distributes the different output voltages to the router components according to their voltage requirements.

The power supply is fixed in the chassis, and is not field-replaceable. It has a single AC appliance inlet that requires a dedicated AC power feed.

For information about site power preparations, see "Power Guidelines, Requirements, and Specifications" on page 545. For information about connecting the router to power and ground, see "Connecting Power to the Services Router" on page 43.

J6300 Power System

The J6300 Services Router uses AC power. You can install one or two autosensing, load-sharing power supplies at the bottom rear of the chassis, as shown in Figure 8. The power supplies distribute the different output voltages to the router components, depending on their voltage requirements. When the power supplies are installed and operational, they automatically share the electrical load.

For full redundancy, two power supplies are required. If a power supply stops functioning for any reason, the second power supply instantly begins providing all the power the router needs for normal functioning. It can provide full power indefinitely.

Each power supply has an LED located on the power supply faceplate. Table 16 describes the J6300 power supply LED.

Table 16: J6300 Power Supply LED

State	Description
Off	No power flowing to the power supply.
Green	Power supply is working correctly.
Red	Power supply is starting up, or has failed.

For information about site power preparations, see "Power Guidelines, Requirements, and Specifications" on page 545. For information about connecting the router to power and ground, see "Connecting Power to the Services Router" on page 43.

Power supplies are hot-removable and hot-insertable. You can remove and replace a redundant power supply without powering down the router or disrupting the routing functions. To avoid electrical injury, carefully follow the instructions in "Replacing Power System Components in a J6300 Router" on page 533.

J4300 and J6300 Cooling System

The cooling system consists of the following components:

- A fan on the midplane
- A fan on the Routing Engine's processor
- An internal fan on the power supply

The airflow produced by these fans keeps router components within the acceptable temperature range (see Figure 12).

Figure 12: Airflow Through the J4300 and J6300 Chassis

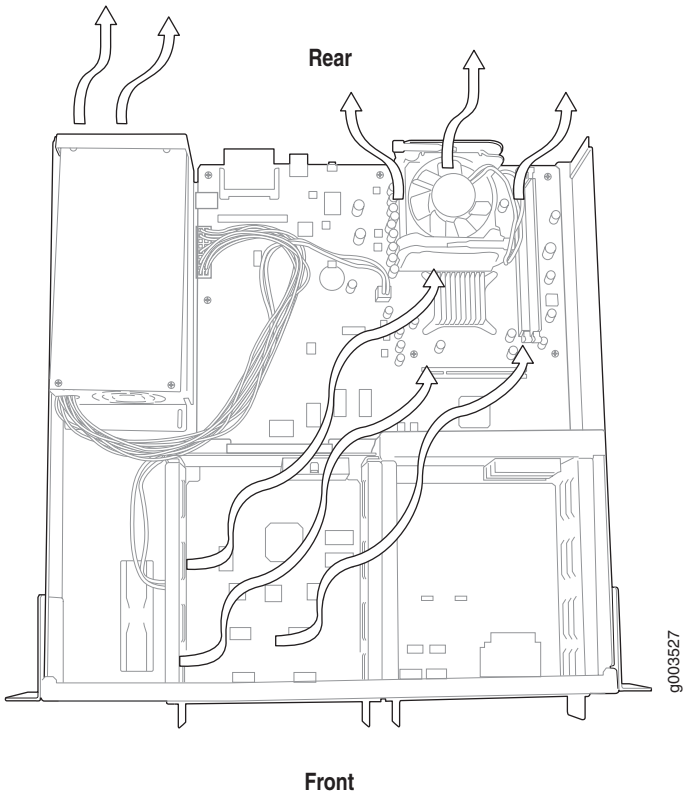

The Routing Engine monitors the temperature of the router components. If the ambient maximum temperature specification is exceeded and the router cannot be adequately cooled, the Routing Engine shuts down the hardware components.

Software Overview

Each J-series Services Router runs the JUNOS Internet software on its general-purpose processors. Designed for the large production networks typically supported by Internet service providers (ISPs), the JUNOS software includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the router chassis.

The JUNOS Internet software runs on the Routing Engine. The Routing Engine kernel coordinates communication among the JUNOS software processes and provides a link to the Packet Forwarding Engine.

With the J-Web interface and the command-line interface (CLI) to the JUNOS software, you configure the routing protocols that run on the Services Router and set the properties of its network interfaces. After activating a software configuration, use either user interface to monitor the protocol traffic passing through the router, manage operations, and diagnose protocol and network connectivity problems.

This section contains the following topics:

Routing Engine and Packet Forwarding Engine

A Services Router has two primary software processing components:

- Routing Engine—Creates and maintains the routing tables that determine how packets are routed through the network.
- Packet Forwarding Engine—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.

For information about Routing Engine hardware, see "J2300 Routing Engine" on page 10 and "J4300 and J6300 Routing Engine" on page 21.

Kernel and Microkernel

The Routing Engine kernel provides the underlying infrastructure for all JUNOS software processes by doing the following:

- Linking the routing tables maintained by the routing protocol process with the forwarding table maintained by the Routing Engine.
- Coordinating communication with the Packet Forwarding Engine, primarily by synchronizing the Packet Forwarding Engine's forwarding table with the master forwarding table maintained by the Routing Engine.

The microkernel contains device drivers and processes that the Packet Forwarding Engine uses to govern the flow of packets through the Services Router.

Processes

The JUNOS software running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual Services Router functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of the JUNOS software, for added flexibility.

The following processes are primary:

Management Process

The JUNOS management process (mgd) manages the Services Router system as follows:

- Provides communication between the other processes and an interface to the configuration database
- Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured
- Interacts with the other processes when commands are issued through one of the user interfaces on the router

Chassis Process

The JUNOS chassis process (chassisd) controls a Services Router chassis and its components as follows:

- Detects hardware on the system that is used to configure network interfaces with the J-Web user interface
- Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered
- Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully

Routing Protocols Process

The Services Router forwards packets through a network by means of the routing protocols it uses and the routing and forwarding tables it maintains. By selecting routes and maintaining forwarding tables, the JUNOS routing protocols process (rpd) defines how routing protocols such as RIP, OSPF, and BGP operate on the router.

Interface Process

The JUNOS interface process (ifd) supplies the programs that configure and monitor network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.

Forwarding Process

The JUNOS forwarding process (fwdd) is responsible for most of the packet transmission through a Services Router. The overall performance of the router is largely determined by the effectiveness of the forwarding process.

User Interfaces

The user interfaces on a Services Router interact with the management process to execute commands and store and retrieve information from the configuration database. The user interfaces operate as clients that communicate with the JUNOS Internet software through an application programming interface (API).

The following primary user interfaces are shipped with the router:

- J-Web graphical user interface—Includes quick configuration capabilities for performing the minimum required steps to enable a feature, plus a built-in configuration editor with access to the entire configuration hierarchy to fully configure the router. The J-Web interface also provides tools for monitoring, managing, and diagnosing router operation.

- Command-line interface (CLI)—Grants access to the complete JUNOS command and configuration hierarchies, to monitor and diagnose the router and configure it completely.

For more information, see “J-series User Interface Overview” on page 109.

Other user interfaces for the Services Router interact with the management process through the common API interface. These interfaces are designed to facilitate the configuration of one or, in some cases, many routers on the network. Among the supported interfaces are the JUNOScope and Service Deployment System (SDX) applications. For more information about these products, see the *JUNOScope Software User Guide* and the *SDX Software Basics Guide*.

Part 2

Installing the J-series Services Router

Chapter 3

Installing and Connecting a Services Router

Make the appropriate preparations and verify the J-series equipment before installing a J-series Services Router and connecting it to a power source and the network.

This chapter contains the following topics:

Before You Begin

Before you begin installation, complete the following tasks:

- Read the information in "Maintenance and Operational Safety Guidelines and Warnings" on page 588, with particular attention to "Chassis Lifting Guidelines" on page 577.
- Determine where to install the Services Router, and verify that the rack or installation site meets the requirements described in "Preparing for Router Installation" on page 541.
- For installation, gather the equipment and tools listed in Table 17.

Table 17: Equipment and Tools Required for Services Router Installation

Desk Installation—J2300 Services Router Only	Wall Installation—J2300 Services Router Only	Rack Installation
Rubber feet (provided)	■ Rubber feet (provided) ■ Mounting brackets and screws (provided) ■ Number 2 Phillips screwdriver ■ Four wall screws or four mounting screws and anchors capable of supporting the full weight of the chassis, up to 12 lb (5.4 kg)	■ Mounting brackets and screws (provided) ■ Number 2 Phillips screwdriver ■ Four (J2300) or eight (J4300 and J6300) mounting screws appropriate for your rack

- To connect the router to power and ground, have ready a 14 AWG grounding cable and lug, as specified in “Chassis Grounding” on page 42, and the power cord or cords shipped with the router.
- To connect network interfaces, have ready a length of cable used by the interface, as specified in “Network Cable Specifications and Connector Pinouts” on page 551.

Unpacking the J-series Services Router

The Services Router is shipped in a cardboard carton and secured with foam packing material. The carton also contains an accessory box and quick start instructions.

NOTE: The router is maximally protected inside the shipping carton. Do not unpack it until you are ready to begin installation.

To unpack the router:

1. Move the shipping carton to a staging area as close to the installation site as possible, but where you have enough room to remove the router.
2. Position the carton so that the arrows are pointing up.
3. Open the top flaps on the shipping carton.
4. Remove the accessory box, and verify the contents against the parts inventory on the label attached to the carton.
5. Pull out the packing material holding the router in place.
6. Verify the contents of the carton against the packing list included with the router. A generic parts inventory appears in Table 18.

7. Save the shipping carton and packing materials in case you later need to move or ship the router.

Table 18: Generic Inventory of Services Router Shipping Carton

Component	J2300 Services Router	J4300 Services Router	J6300 Services Router
Chassis	1	1	1
Physical Interface Module (PIM)	2 Fast Ethernet ports and 1 of the following interfaces: ■ 2-port E1 ■ 2-port Serial ■ 2-port T1 **NOTE:** The interfaces installed in the J2300 Services Router are not field-replaceable. For more information, see "J2300 Physical Interface Module (PIM)" on page 13.	Between 0 and 6 of the following in any combination: ■ 1-port DS3 (T3) PIM ■ 2-port E1 PIM ■ 2-port Fast Ethernet PIM ■ 2-port Serial PIM ■ 2-port T1 PIM	
Power supply	1 (fixed)	1 (fixed)	1 or 2
Mounting brackets	2	2 (fixed)	2 (fixed)
Blank panels for slots without components	0	Depends on router configuration	Depends on router configuration

Installing the J2300 Services Router

You can install the J2300 Services Router on a desk, on a wall, or in a rack. The J2300 Services Router includes mounting brackets that support either wall or rack mounting, and rubber feet for desk and wall mounting.

Install the J2300 Services Router as appropriate for your site, with one of the following procedures:

- Installing the J2300 Services Router on a Desk on page 37
- Installing the J2300 Services Router on a Wall on page 38
- Installing the J2300 Services Router into a Rack on page 39

Installing the J2300 Services Router on a Desk

You can install the J2300 Services Router on a desk, table, or other level surface. The router is shipped with rubber feet in the accessory box. The rubber feet are necessary to stabilize the router on the desk.

To install the J2300 router on a desk:

1. Turn the chassis upside-down on the desk or work surface where you intend to operate the router.
2. Attach the provided rubber feet to the bottom of the chassis, as shown in Figure 13.
3. Turn the chassis right-side up on the desk or work surface.

Figure 13: Attaching Rubber Feet to the J2300 Services Router

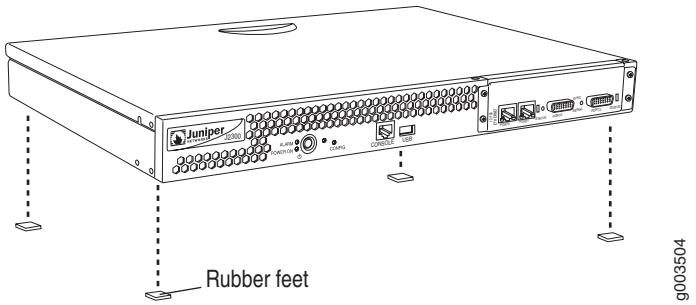

Installing the J2300 Services Router on a Wall

You can install the J2300 Services Router on a wall. The router is shipped with mounting brackets and rubber feet in the accessory box. The rubber feet help stabilize the router on the wall and enhance airflow.

To install the J2300 router on a wall:

1. Turn the chassis upside-down on a desk or work surface near where you intend to operate the router.
2. Attach the provided rubber feet to the bottom of the chassis, as shown in Figure 13.
3. Turn the chassis right-side up, and place it on a flat surface.
4. Position a mounting bracket on each side of the chassis as shown in Figure 14.
5. Use a number 2 Phillips screwdriver to install the screws that secure the mounting brackets to the chassis.
6. If you are using wall anchors to support the chassis, install two pairs of anchors on the wall, spaced 19 in. (48.3 cm) apart.

CAUTION: Mounting screws and wall anchors must be strong enough to support the full weight of the chassis, up to 12 lb (5.4 kg). Attaching the router to wall studs or using wall anchors provides extra support for the chassis.

7. Have one person grasp the sides of the router, lift the router, and position it on the wall.
8. Have a second person install two pairs of mounting screws through the bracket holes on either side of the router, to secure the router to the wall.
9. Verify that the mounting screws on one side are aligned with the mounting screws on the opposite side and that the router is level.

Figure 14: Attaching Mounting Brackets to Install a J2300 Services Router on a Wall

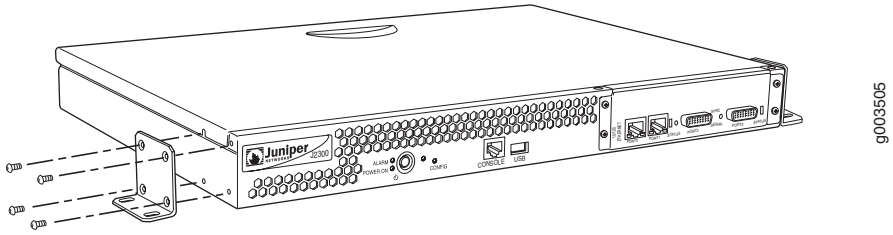

Installing the J2300 Services Router into a Rack

You can front-mount the J2300 Services Router in a rack. The router is shipped with mounting brackets in the accessory box. Many types of racks are acceptable, including four-post (telco) racks, enclosed cabinets, and open-frame racks. For more information about the type of rack or cabinet the J-series router can be installed into, see "Rack Requirements" on page 542.

NOTE: If you are installing multiple routers in one rack, install the lowest one first and proceed upward in the rack.

CAUTION: The chassis weighs up to 12 lb (5.4 kg). Installing it into the rack requires one person to lift and a second person to secure the mounting screws.

To install the J2300 router into a rack:

1. Position a mounting bracket on each side of the chassis as shown in Figure 15.

2. Use a number 2 Phillips screwdriver to install the screws that secure the mounting brackets to the chassis.
3. Have one person grasp the sides of the router, lift the router, and position it in the rack.
4. Align the bottom hole in each mounting bracket with a hole in each rack rail, making sure the chassis is level.
5. Have a second person install a mounting screw into each of the two aligned holes. Use a number 2 Phillips screwdriver to tighten the screws.
6. Install the second screw in each mounting bracket.
7. Verify that the mounting screws on one side of the rack are aligned with the mounting screws on the opposite side and that the router is level.

Figure 15: Attaching Mounting Brackets to Install a J2300 Services Router in a Rack

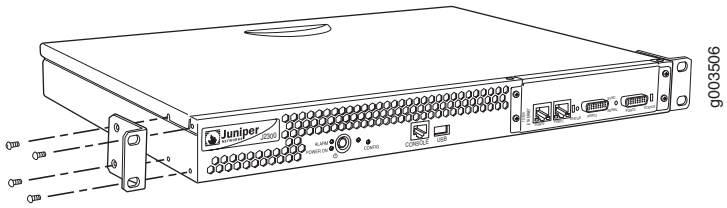

Installing the J4300 or J6300 Services Router

You can front-mount the J4300 Services Router or J6300 Services Router in a rack. The router is shipped with mounting brackets installed. Many types of racks are acceptable, including four-post (telco) racks, enclosed cabinets, and open-frame racks. For more information about the type of rack or cabinet the J-series router can be installed into, see “Rack Requirements” on page 542.

NOTE: If you are installing multiple routers in one rack, install the lowest one first and proceed upward in the rack.

CAUTION: The chassis weighs between 18 lb (8.2 kg) and 24 lb (10.9 kg). Installing it into the rack requires one person to lift and a second person to secure the mounting screws.

To install the J4300 router or J6300 router into a rack:

1. Have one person grasp the sides of the router, lift the router, and position it in the rack.
2. Align the bottom hole in each mounting bracket with a hole in each rack rail as shown in Figure 16, making sure the chassis is level.
3. Have a second person install a mounting screw into each of the two aligned holes. Use a number 2 Phillips screwdriver to tighten the screws.
4. Install the remaining screws in each mounting bracket.
5. Verify that the mounting screws on one side of the rack are aligned with the mounting screws on the opposite side and that the router is level.

Figure 16: Installing the J4300 or J6300 Services Router

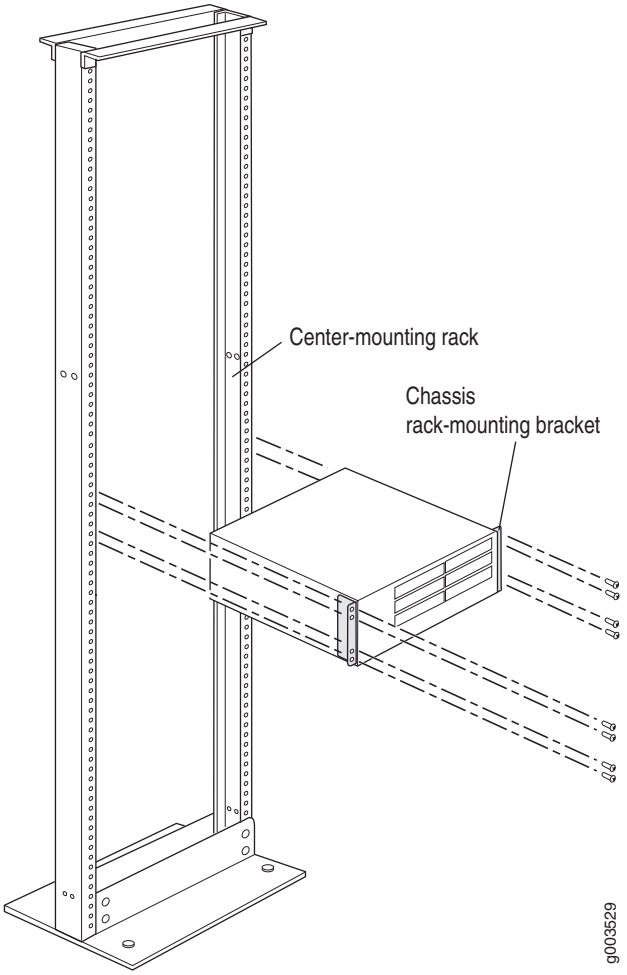

Connecting Interface Cables to the Services Router

You connect the interfaces installed in the Services Router to various network media. For more information about the network interfaces supported on the router, see “Configuring Network Interfaces” on page 79.

1. Have ready a length of the type of cable used by the interface, as specified in “Network Cable Specifications and Connector Pinouts” on page 551.
2. Insert the cable connector into the cable connector port on the interface faceplate.
3. Arrange the cable as follows to prevent it from dislodging or developing stress points:
   a. Secure the cable so that it is not supporting its own weight as it hangs to the floor.
   b. Place excess cable out of the way in a neatly coiled loop.
   c. Place fasteners on the loop to help maintain its shape.

Chassis Grounding

To meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, the Services Router must be adequately grounded before power is connected. In addition to the grounding pin on the power plug cord, a threaded insert (PEM nut), screw, and washer are provided on the rear of the chassis to connect the router to earth ground.

CAUTION: Before router installation begins, a licensed electrician must attach a cable lug to the grounding cable that you supply. A cable with an incorrectly attached lug can damage the router (for example, by causing a short circuit).

The grounding cable must be 14 AWG single-strand wire cable, and must be able to handle the following amperage:

- J2300 router—up to 4 A
- J4300 router and J6300 router—up to 6 A

The grounding lug must be a ring-type, vinyl-insulated TV14–10R lug, or equivalent, to accommodate the 14 AWG cable.

To ground the router before connecting power, you connect the grounding cable to earth ground and then attach the lug on the cable to the chassis grounding point, with the screw. (See “Connecting Power to the Services Router” on page 43.)

Connecting Power to the Services Router

J2300 and J4300 Services Routers have a single fixed power supply. The J6300 Services Router has one or two field-replaceable power supplies. For more information about the J-series power specifications, see “Power Guidelines, Requirements, and Specifications” on page 545.

The AC power cord shipped with the router connects the router to earth ground when plugged into an AC grounding-type power outlet. The router must be connected to earth ground during normal operation.

To connect power to the router:

1. Locate the power cord or cords shipped with the router, which has a plug appropriate for your geographical location. For power cord specifications, see “Power Guidelines, Requirements, and Specifications” on page 545.

2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist, and connect the strip to the ESD point on the chassis. For more information about ESD, see “Preventing Electrostatic Discharge Damage” on page 567.

3. Use a grounding cable to connect the router to earth ground: (For cable requirements, see “Chassis Grounding” on page 42.)

   a. Verify that a licensed electrician has attached an appropriate grounding cable lug to the grounding cable.

   b. Connect one end of the grounding cable to a proper earth ground, such as the rack in which the router is installed.

   c. With a Phillips screwdriver, remove the screw and washer from the PEM nut at the grounding point on the rear of the chassis.

   d. Place the grounding lug at the other end of the cable over the grounding point, as shown in Figure 17, Figure 18, and Figure 19.

   e. Secure the cable lug to the grounding point, first with the washer, then with the screw.

4. For each power supply:

   a. Insert the appliance coupler end of a power cord into the appliance inlet on the power supply faceplate, as shown in Figure 17, Figure 18, and Figure 19.

   b. Insert the plug into an AC power source receptacle.

5. Verify that the power cord does not block access to router components or drape where people can trip on it.

Figure 17: Connecting Power to the J2300 Services Router

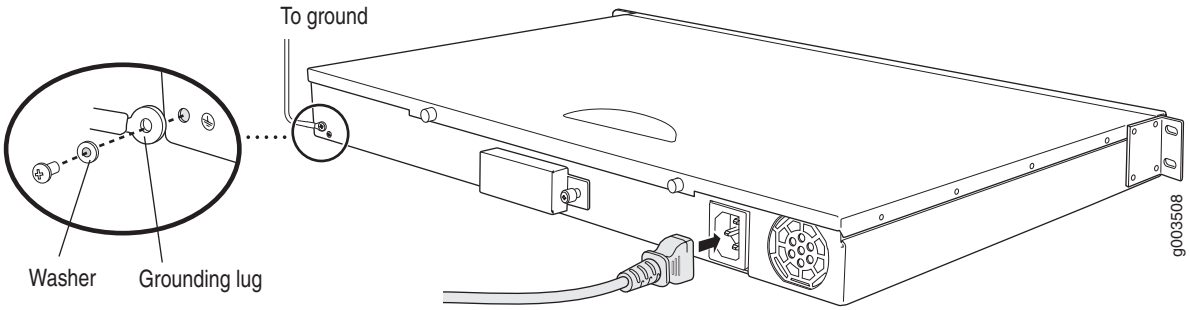

Figure 18: Connecting Power to the J4300 Services Router

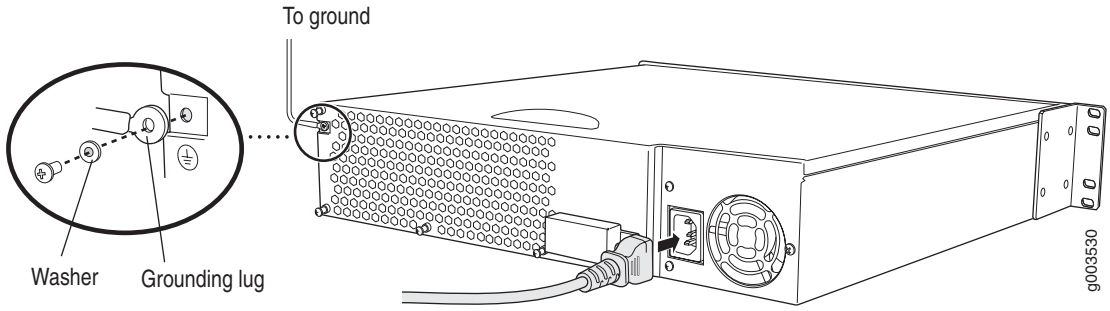

Figure 19: Connecting Power to the J6300 Services Router

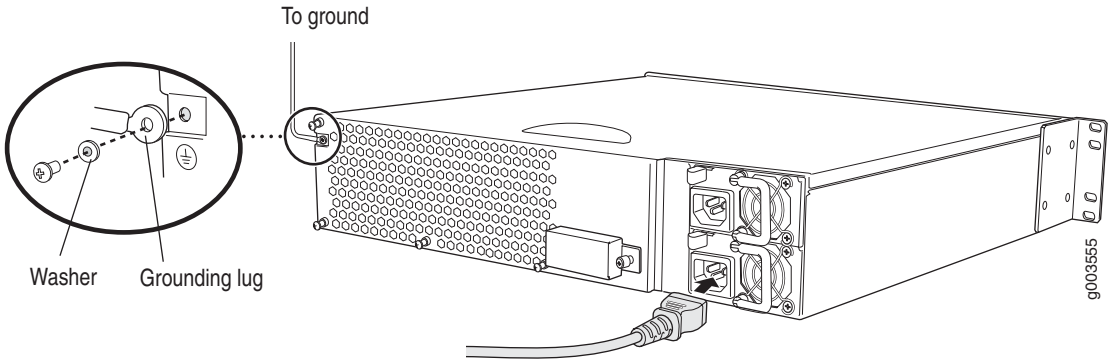

Powering a Services Router On and Off

To power on a Services Router, press the power button. The Routing Engine boots as the power supply completes its startup sequence. The POWER ON LED lights during startup and remains on steadily when the router is operating normally.

To power off a Services Router, do one of the following:

- Press and release the power button. The router begins gracefully shutting down the operating system and then powers itself off.
- Press the power button and hold it for more than 5 seconds. The router immediately powers itself off without shutting down the operating system.

To remove power completely from the router, unplug the power cord. The power button on the Services Router is a standby power switch. If the router is connected to an AC power source receptacle when you press the power button to power the router off, the router remains in standby mode and a small amount (5 V and 3.3 V) of standby voltage is still available in the chassis.

Chapter 4

Establishing Basic Connectivity

The JUNOS software is preinstalled on the Services Router. When the router is powered on, it is ready to be configured. If the router does not have a configuration from the factory or your service provider, you must configure the software to establish basic connectivity.

If you are setting up a Services Router for the first time, you can use either J-Web Quick Configuration or the JUNOS CLI configuration editor to configure basic connectivity.

If you are setting up many Services Routers, autoinstallation can help automate the installation process.

This chapter contains the following topics. For more information about basic connectivity, see the *JUNOS System Basics Configuration Guide*.

Basic Connectivity Terms

Before configuring basic connectivity, become familiar with the terms defined in Table 19.

Table 19: Basic Connectivity Terms

Term	Definition
domain name	Name that identifies the network or subnetwork a router belongs to.
Dynamic Host Configuration Protocol (DHCP)	Protocol for assigning dynamic IP addresses to devices on a network.
gateway	Packets destined for IP addresses not identified in the routing table are sent to the default gateway.
hostname	Unique name that identifies a router on the network.
loopback address	IP address of a Services Router on logical interface lo0.0 that is always active and available to external hosts and as the source address for outgoing packets.
Network Time Protocol (NTP)	Protocol that provides a reliable way of synchronizing the system time of a router.
root user	A superuser or system administrator who can perform any task in the file system.
secure shell (SSH)	Software that provides a secured method of logging in to a remote network system.
telnet	Software that allows a computer to act as a remote terminal on a network system.

Basic Connectivity Overview

To connect your Services Router to the network and establish basic connectivity, you enter information about your network. This overview includes the following topics:

Router Identification

The domain name defines the network or subnetwork that the Services Router belongs to. The hostname refers to the specific machine, while the domain name is shared among all the devices in a given network. Together the hostname and domain name identify the router in the network.

Root Password

The root user has complete privileges to configure the Services Router, and manage files in the router's file system. Initially, the root password is not defined on the router. To ensure basic security, you must define the root password during initial configuration. If you use a plain-text password, the router displays the password as a encrypted string so that users viewing the configuration cannot easily see the password.

The root password must meet the following conditions:

- The password must be at least 6 characters long.
- You can include most character classes in a password (alphabetic, numeric, and special characters), except control characters.
- Valid passwords must contain at least one change of case or character class.

Time Zone and System Time

You define the time zone for the location where you plan to operate the Services Router by using a designation that consists of the following information for the location:

- Name of the continent or ocean—For example, America or Atlantic
- Name of the major city or other geographic feature in the time zone—For example, Detroit or Azores

A Network Time Protocol (NTP) server provides accurate time across a network. The router synchronizes the system time with the NTP server, and periodically accesses the NTP server to maintain the correct time.

The time zone and system time must be accurate so that the router schedules events and operations as expected.

Network Settings

A Domain Name System (DNS) server on the network maintains a database for resolving hostnames and IP addresses. Network devices can query the DNS server by hostnames rather than IP addresses. The router accesses the DNS servers that are added to the configuration to resolve hostnames in the order in which you list them.

If you plan to include your router in several domains, you can add these domains to the configuration so that they are included in a DNS search. When DNS searches are requested, the domain suffixes are appended to the hostnames.

Default Gateway

A default gateway is a static route that is used to direct packets addressed to networks not explicitly listed in the router's routing table. If a packet arrives at the Services Router with an address that the router does not have routing information for, the router sends the packet to the default gateway. The default gateway entry is always present in the routing and forwarding tables.

Backup Router

You can specify a backup router to take over when the Services Router's routing protocol process is not running, usually when the Services Router is booting, or if its routing protocol process has failed. Packets arriving at a Services Router in this situation are routed to the backup router. When the routing protocol process starts up again, the address of the backup router is removed from the Services Router's routing and forwarding tables. The backup router must be located on the same subnet.

NOTE: To configure a backup router, you must use the CLI or J-Web configuration editor. You cannot configure a backup router with J-Web Quick Configuration.

Loopback Address

The loopback address is the IP address of the Services Router itself. The loopback address ensures that the Services Router provides an IP address to management applications. Because it must always be available to hosts attempting to route packets to the Services Router, the loopback address resides on an interface that is always active, known as the loopback interface (lo0.0). Setting a loopback address ensures that the Services Router can receive packets addressed to the loopback address as long as the router is reachable though any ingress interface. In addition, applications such as NTP, RADIUS, and TACACS+ can use the loopback address as the source address for outgoing packets.

If you use the J-Web Set Up Quick Configuration page, you can either set a loopback address of your choice or have the loopback address automatically set to 127.0.0.1 when you click **Apply** or **OK** to commit the configuration.

Management Interface Address

The Fast Ethernet interface fe-0/0/0, labeled PORT 0 on the front panel of the Services Router, is the network interface through which you perform initial router setup. After the router is initially configured, you can attach fe-0/0/0 to the management network for use as a management interface.

Before Initial Configuration

Before initial configuration, when the factory default configuration is active:

1. The Services Router attempts to perform autoinstallation by obtaining a router configuration through all its connected interfaces, including fe-0/0/0. The Services Router acts as a DHCP client out the fe-0/0/0 interface.

2. If the Services Router does not find a DHCP server within a few seconds, it sets the address of fe-0/0/0 to 192.168.1.1/24 and becomes a DHCP server out the fe-0/0/0 interface.

   With the router temporarily acting as a DHCP server, you can manually configure it with the J-Web interface. Any DHCP client host (a PC or laptop computer, for example) directly connected to fe-0/0/0 receives an address on the 192.168.1.1/24 network.

During Initial Configuration

Once you connect your laptop or PC to fe-0/0/0, you can use a Web browser to visit the address 192.168.1.1/24, access the J-Web Set Up Quick Configuration page, and initially configure the router.

After Initial Configuration

After you perform the initial configuration and commit it by clicking **Apply** or **OK** on the Set Up page, the configured router can no longer act as a DHCP server. You can do either of the following:

- Continue to use the J-Web Quick Configuration and leave the IP address and prefix length as 192.168.1.1/24. You can continue configuring the router until the DHCP lease expires, or the physical connection is lost because the cable is disconnected or the router is rebooted.

- Change the IP address and prefix length. You lose access to the router until you either adjust the IP address of the management device to be on the same subnetwork as the router, or connect to the router through the console port.

Management Access

Telnet allows you to connect to the Services Router and access the CLI to execute commands from a remote system. Telnet connections are not encrypted and therefore can be intercepted.

Telnet access to the root user is prohibited. You must use more secure methods, such as SSH, to log in as root.

If you are using a JUNOScript server to configure and monitor routers, you can activate clear-text access on the router to allow unencrypted text to be sent directly over a TCP connection without using any additional protocol (such as

SSH, SSL, or telnet). Information sent in clear-text is not encrypted and therefore can be intercepted. For more information about the JUNOScript application programming interface (API), see the *JUNOScript API Guide*

SSH also allows you to connect to the router and access the CLI to execute commands from a remote system. However, unlike telnet, SSH encrypts the password so that it cannot be intercepted.

SSH connections are authenticated by a digital certificate. SSH uses public-private key technology for both connection and authentication. The SSH client software must be installed on the machine where the client application runs. If the SSH private key is encrypted (for greater security), the SSH client must be able to access the passphrase used to decrypt the key. For information about obtaining SSH software, see http://www.ssh.com and http://www.openssh.com.

Before You Begin

Before you begin initial configuration, complete the following tasks:

- Install the Services Router in its permanent location, as described in “Installing and Connecting a Services Router” on page 35.
- Gather the following information:
  - Hostname for the router on the network
  - Domain that the router belongs to on the network
  - Password for the root user
  - Time zone where the router is located
  - IP address of an NTP server (if NTP is used to set the time on the router)
  - IP address of a DNS server
  - List of domains that can be appended to hostnames for DNS resolution
  - IP address of the default gateway
  - IP address to be used for the loopback interface
  - IP address of the fe-0/0/0 interface
- If you are performing the initial configuration with the J-Web interface, collect the following equipment:
  - A management device, such as a laptop, with an Ethernet port
  - An Ethernet cable
- If you are performing the initial configuration with the CLI, collect the following equipment:
  - A management device, such as a PC or laptop, with a serial port and an asynchronous terminal application (such as Microsoft Windows Hyperterminal)
  - An RJ-45 to DB-9 serial port adapter (provided)
  - A rollover Ethernet cable (provided)

Configuring the Services Router with J-Web Quick Configuration

If you plan to use the J-Web interface to configure the Services Router, you must connect through LAN PORT 0, as shown in Figure 20 and Figure 21.

Before you configure the router, gather the information described in “Before You Begin” on page 53.

To configure the router with J-Web Quick Configuration, perform the following procedures:

- Connecting to the J-Web Interface on page 54
- Configuring Basic Settings with Quick Configuration on page 55

Connecting to the J-Web Interface

When the Services Router is powered on for the first time, if no configuration is present, the fe-0/0/0 interface on LAN PORT 0 acts as a DHCP server and assigns an IP address within the 192.168.1/24 subnetwork to any devices connected to it.

To connect to the J-Web interface using LAN PORT 0 on the router (see Figure 20 and Figure 21):

1. On the management device, such as a PC or laptop, that you will use to access the J-Web interface, verify that the address of the port that you connect to the router is set to one of the following:
   - An Ethernet address other than 192.168.1.1 on the 192.168.1/24 subnetwork
   - An Ethernet address from a DHCP server
2. Turn off the power to the management device.
3. Plug one end of the Ethernet cable into the Ethernet port on the management device.
4. Connect the other end of the Ethernet cable to LAN PORT 0 on the router.
5. Power on the router by pressing the power button on the front panel. Verify that the POWER ON LED on the front panel turns green.
6. Turn on the power to the management device. The router assigns an IP address to the management device within the 192.168.1/24 subnetwork if the device is configured to use DHCP.
7. From the management device, open a Web browser and enter the IP address 192.168.1.1 in the address field. The Set Up Quick Configuration page appears.

Figure 20: Connecting to the Fast Ethernet Port on the J2300 Services Router

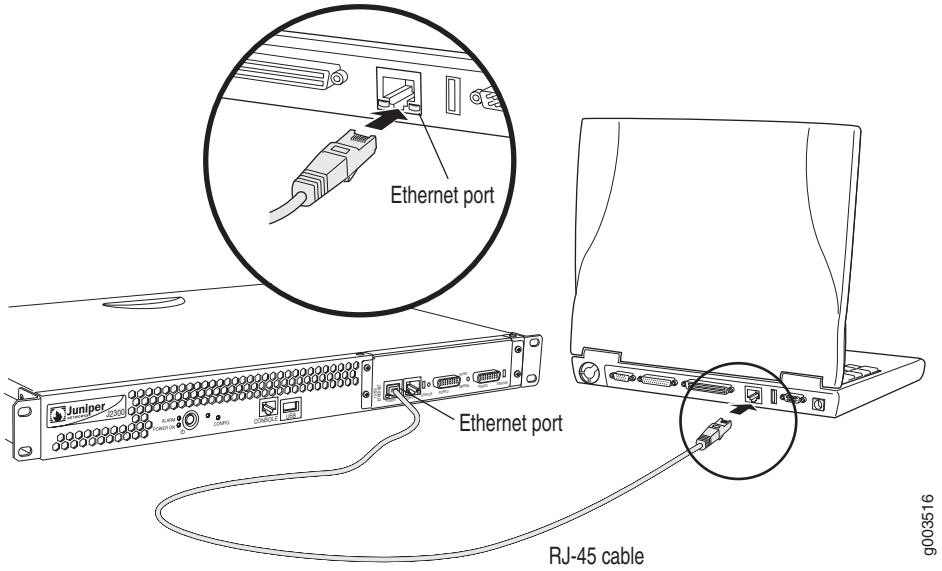

Figure 21: Connecting to the Fast Ethernet Port on the J4300 or J6300 Services Router

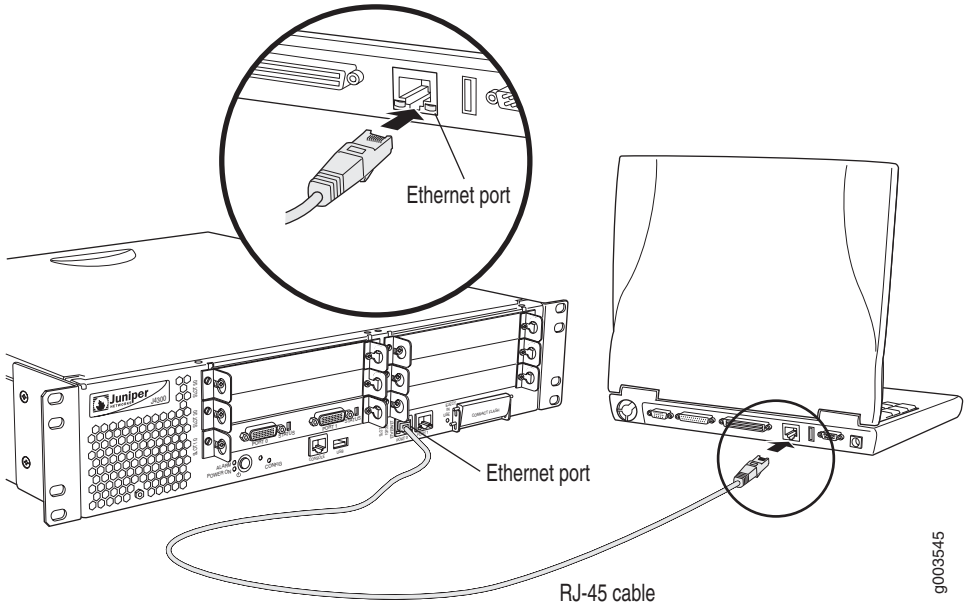

Configuring Basic Settings with Quick Configuration

To configure basic settings in the J-Web interface:

1. Enter information into the fields described in Table 20 on the Set Up Quick Configuration page.
2. Click one of the following buttons:
   - To apply the configuration and stay in the Set Up Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the Quick Configuration page, click **OK**.
   - To cancel your entries and return to the Quick Configuration page, click **Cancel**.

NOTE: Once initial configuration is complete, the Services Router stops functioning as a DHCP server. If you change the IP address of fe-0/0/0 and have the management device configured to use DHCP, you lose your DHCP lease and your connection to the router through the J-Web interface. To reestablish a connection, either set the IP address on the management device manually, or connect fe-0/0/0 to the management network and access the router another way—for example, through the console port.

3. To check the configuration, see "Displaying Basic Connectivity Configurations" on page 67.

Table 20: Set Up Quick Configuration Field Descriptions

Field	Function	Your Action
Identification		
Host Name (required)	Defines the hostname of the router.	Type the hostname.
Domain Name	Defines the network or subnetwork that the machine belongs to.	Type the domain name.
Root Password (required)	Sets the root password that user "root" can use to log in to the router.	Type a plain-text password that the system encrypts. **NOTE:** After a root password has been defined, it is required when you log in to the J-Web user interface or the CLI.
Verify Root Password (required)	Verifies the root password has been typed correctly.	Retype the password.
Time		
Time Zone	Identifies the time zone that the router is located in.	From the drop-down list, select the appropriate time zone.

Field	Function	Your Action
NTP Servers	Specify an NTP server that the router can reach to synchronize the system time.	To add an IP address, type it in the box to the left of the Add button, then click **Add**. To delete an IP address, click on it in the box above the Add button, then click **Delete**.
Current System Time	Synchronizes the system time with the NTP server, or manually set the system time and date.	■ To immediately set the time using the NTP server, click **Set Time via NTP**. The router sends a request to the NTP server and synchronizes the system time. **NOTE:** If you are configuring other settings on this page, the router also synchronizes the system time using the NTP server when you click **Apply** or **OK**. ■ To set the time manually, click **Set Time Manually**. A pop-up window allows you to select the current date and time from drop-down lists.
Network		
DNS Name Servers	Specify a DNS server that the router can use to resolve hostnames into addresses.	To add an IP address, type it in the box to the left of the Add button, then click **Add**. To delete an IP address, click on it in the box above the Add button, then click **Delete**.
Domain Search	Adds each domain name that the router is included in to the configuration so that they are included in a DNS search.	To add a domain name, type it in the box to the left of the Add button, then click **Add**. To delete a domain name, click on it in the box above the Add button, then click **Delete**.
Default Gateway	Defines a default gateway through which to direct packets addressed to networks not explicitly listed in the routing table.	Type a 32-bit IP address, in dotted decimal notation.
Loopback Address	Defines a reserved IP address that is always available on the router. If no address is entered, this address is set to 127.0.0.1/32.	Type a 32-bit IP address and prefix length, in dotted decimal notation.

Field	Function	Your Action
fe-0/0/0 Address	Defines the IP address and prefix length of fe-0/0/0. The interface fe-0/0/0 is typically used as the management interface for accessing the router. By default this address is set to 192.168.1.1/24.	Type a 32-bit IP address and prefix length, in dotted decimal notation. **NOTE:** If you change the fe-0/0/0 address, you will lose your connection to the J-Web interface when you click **Apply** or **OK**. If you need to change this address but want to continue using the J-Web interface after applying the initial configuration, set the IP address on the management device manually.
Management Access		
Allow Telnet Access	Allows remote access to the router using telnet.	To enable telnet access, select the check box.
Allow JUNOScript over Clear-Text Access	Allows JUNOScript to access the router using a protocol for sending unencrypted text over a TCP connection.	To enable JUNOScript access over clear-text, select the check box.
Allow SSH Access	Allows remote access to the router using SSH.	To enable SSH access, select the check box.

Configuring the Services Router with a Configuration Editor

If you plan to use the CLI to configure the router, you must connect through the CONSOLE port, as shown in Figure 22 and Figure 23.

You can configure basic settings in the J-Web interface from a device attached to the fe-0/0/0 on LAN PORT 0. For instructions, see "Connecting to the J-Web Interface" on page 54.

Before you configure the router, gather the information described in "Before You Begin" on page 53

This section contains the following topics:

Connecting to the CLI

To connect to the CLI using the console port on the router:

1. Turn off the power to the management device, such as a PC or laptop computer, that you are using to access the CLI.
2. Plug one end of an Ethernet rollover cable into the RJ-45 to DB-9 serial port adapter (see Figure 22 and Figure 23).

NOTE: The Ethernet rollover cable and RJ-45 to DB-9 serial port adapter are provided in the router's accessory box.

3. Plug the RJ-45 to DB-9 serial port adapter into the serial port on the management device (see Figure 22 and Figure 23).
4. Connect the other end of the Ethernet rollover cable to the console port on the router (see Figure 22 and Figure 23).
5. Turn on the power to the management device.
6. Start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate COM port to use (for example, COM1).
7. Configure the port settings as follows:
   - Bits per second: 9600
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
8. Power on the router by pressing the power button on the front panel. Verify that the POWER ON LED on the front panel turns green.

   The terminal emulation screen on your management device displays the boot sequence. When the router has finished booting, a login prompt appears.
9. Log in as the user "root". There is no password.

Figure 22: Connecting to the Console Port on the J2300 Services Router

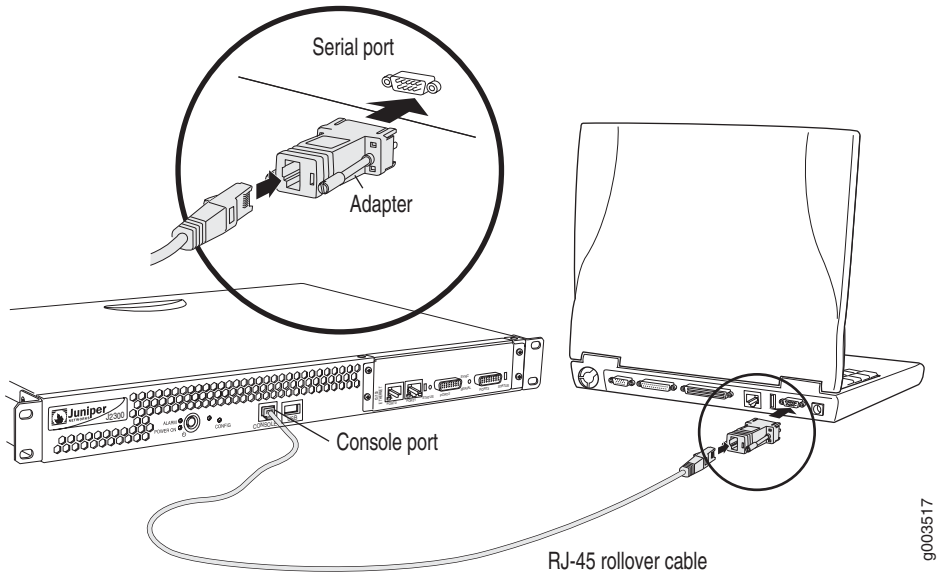

Figure 23: Connecting to the Console Port on the J4300 or J6300 Services Router

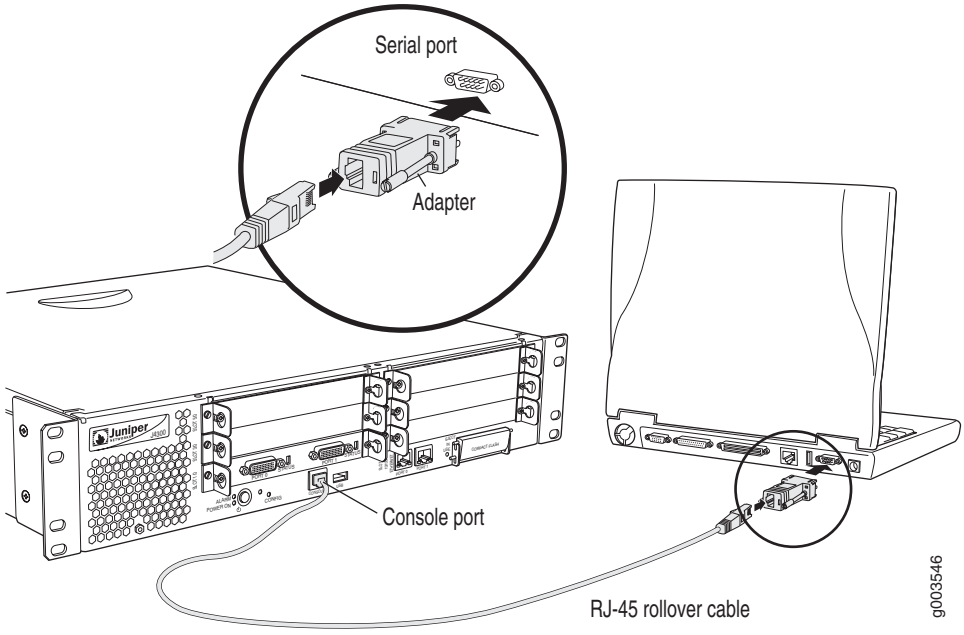

Configuring Basic Settings with a Configuration Editor

To establish basic connectivity on a Services Router, you identify the router, connect the router to the network, and specify basic network settings.

In a typical network, the Services Router has the basic settings listed in Table 21. Determine the values to set on the Services Router in your network.

Table 21: Sample Settings on a Services Router

Services Router Property	Value
Services Router hostname	routera
Access for user "root"	SSH RSA public key
IP address of the NTP server used to synchronize system time on the Services Router	10.148.2.21
Services Router location	Sunnyvale, California, USA, which is in the America/Los_Angeles time zone
IP address of the DNS server to which DNS requests are sent	10.148.2.32
Domains to which the Services Router belongs	lab.router.net and router.net
IP address of a backup router to use while the Services Router is booting or if the routing protocol processes fail to start	192.168.2.44
Loopback IP address and prefix length for the Services Router lo0 interface	172.16.1.24/32
IP address and prefix length for the Services Router fe-0/0/0 interface	192.168.2.24/24

To use a configuration editor to configure basic settings:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure basic settings, perform the configuration tasks described in Table 22.
3. If you are using the J-Web interface, click **Commit** to view a summary of your changes, then click **OK** to commit the configuration. If you are using the CLI, commit the configuration by entering the commit command.
4. To check the configuration, see "Displaying Basic Connectivity Configurations" on page 67.

Table 22: Configuring Basic Settings

Task	J-Web Interface Configuration Editor	CLI Configuration Editor
Navigate to the **System** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Configuration > View and Edit > Edit Configuration**. 2. Next to System, click **Configure** or **Edit**.	From the top of the configuration hierarchy, enter edit system.
Define the hostname of the router.	In the Host name box, type the router's hostname—for example, routera.	Set the hostname. For example: set host-name routera
Name the domain in which the router is located.	In the Domain name box, type the router's domain name—for example, lab.router.net.	Set the domain name. For example: set domain-name lab.router.net
Allow SSH remote access.	1. In the Nested configuration section, next to Services, click **Configure** or **Edit**. 2. Next to Ssh, click **Configure** or **Edit**. 3. Click **OK**. 4. Click **OK** a second time to return to the System level in the configuration editor hierarchy.	Set remote access for SSH: set services ssh
Define root authentication for access to the router. **NOTE:** For readability, the entire key is not shown.	1. In the Nested configuration section, next to Root authentication, click **Configure** or **Edit**. 2. Next to Ssh rsa, click **Add New Entry**. 3. In the Authorized key box, type the RSA password—for example, ssh-rsa AAAAB3Nza...D9Y2gXF9ac== root@routera.lab.router.net 4. Click **OK**. 5. Click **OK** a second time to return to the System level in the configuration editor hierarchy.	Set the root password. For example: set root-authentication ssh-rsa "ssh-rsa AAAAB3Nza...D9Y2gXF9ac== root@routera.lab.router.net"
Define the time zone the router is located in.	In the Time zone drop-down list, select the time zone for your router—for example, **America/Los_Angeles**.	Set the time zone. For example: set time-zone America/Los_Angeles

Task	J-Web Interface Configuration Editor	CLI Configuration Editor
Define the NTP server that NTP requests can be sent to.	1. In the Nested configuration section, next to Ntp, click **Configure** or **Edit**. 2. Next to Server, click **Add New Entry**. 3. In the Address box, type the NTP server's IP address—for example, 10.148.2.21 4. Click **OK**. 5. Click **OK** a second time to return to the System level in the configuration editor hierarchy.	Set the address of the NTP server. For example: set ntp server address 10.148.2.21
Define the DNS server that DNS requests can be sent to.	1. Next to Name server, click **Add New Entry**. 2. In the Address box, type the address of the DNS server—for example, 10.148.2.32. 3. Click **OK**.	Set the address of the DNS server. For example: set name-server 10.148.2.32
Add each domain that the router belongs to.	1. Next to Domain search, click **Add New Entry**. 2. In the Value box, type the name of the domain in which the rouer is located—for example, lab.router.net. 3. Click **OK**. 4. Next to Domain search, click **Add New Entry**. 5. In the Value box, type the name of another domain that the router belongs to—for example, router.net. 6. Click **OK**.	Set the domains to be searched. For example: set domain-search lab.router.net set domain-search router.net
Define the backup router to be used when the router is booting or the routing protocol processes are not running.	In the Backup router section, next to Address, type the IP address of the backup router—for example, 192.168.2.44.	Set the backup router. For example: set backup router address 192.168.2.44

Task	J-Web Interface Configuration Editor	CLI Configuration Editor
Define the IP address for lo0.0.	1. In the configuration editor hierarchy, next to Interfaces, click **Configure** or **Edit**. 2. In the Interface table, locate the lo0 row and click **Unit**. 3. In the Unit table, click **0**, and in the Family section next to Inet, click **Configure** or **Edit**. 4. To delete the existing IP address, click the **Discard** button. Select the **Delete Configuration Below This Point** radio button from the next display. 5. Next to Address, click **Add new entry**. 6. In the Source box, type the address and prefix length for the loopback interface—for example, 172.16.1.24/32. 7. Click **OK**.	1. Exit the system level of the hierarchy: exit 2. From the top of the configuration hierarchy, enter edit interfaces. 3. Delete the existing IP address: delete lo0 unit 0 family inet address. 4. Set the IP address and prefix length of lo0.0. For example: set lo0 unit 0 family inet address 172.16.1.24/32.
Define the IP address for fe-0/0/0.	1. In the configuration editor hierarchy, next to Interfaces, click **Configure** or **Edit**. 2. In the Interface table, locate the fe–0/0/0 row and click **Unit**. 3. In the Unit table, click **0**, and in the Family section next to Inet, click **Configure** or **Edit**. 4. To delete the existing IP address, click the **Discard** button. Select the **Delete Configuration Below This Point** radio button from the next display. 5. Next to Address, click **Add new entry**. 6. In the Source box, type the address and prefix length for the management interface—for example, 192.168.1.1/24. 7. Click **OK**.	1. Delete the existing IP address: delete fe-0/0/0 unit 0 family inet address. 2. Set the IP address and prefix length of fe-0/0/0. For example: set fe-0/0/0 unit 0 family inet address 192.168.1.1/24

Configuring Autoinstallation

This section contains the following topics:

Autoinstallation Overview

You can download a configuration file automatically from an FTP, Hypertext Transfer Protocol (HTTP), or Trivial File Transfer Protocol (TFTP) server. When you power on a Services Router configured for autoinstallation, it requests an IP address from a DHCP server. After the Services Router has an address, it sends a request to a configuration server and downloads and installs a configuration.

Autoinstallation is enabled when you turn on power to a Services Router that does not have a valid configuration file stored in nonvolatile RAM (NVRAM). When autoinstallation is enabled, a network manager loads an existing configuration file automatically to the Services Router that you are deploying. Autoinstallation is helpful for deploying many Services Router, because it centralizes and automates the installation process.

For autoinstallation to work, the Services Router must have or must acquire a unique IP address. Address resolution takes place in one of the following ways:

- For an attached LAN interface with High-Level Data Link Control (HDLC) encapsulation, autoinstallation issues DHCP requests or bootstrap protocol (BOOTP) requests, or uses the Reverse Address Resolution Protocol (RARP). For more information about DHCP, see RFC 2131, *Dynamic Host Configuration Protocol*.
- For a serial interface with Frame Relay encapsulation, autoinstallation uses BOOTP.
- For a serial interface without Frame Relay encapsulation, autoinstallation uses the Serial Line Address Resolution Protocol (SLARP).

Autoinstallation Requirements for End Users

When enabling autoinstallation as an end user, ensure that you have performed the following tasks:

- Installed the Services Router
- Powered on the Services Router
- Connected an interface on the Services Router to a network that has access to a DHCP server, a DHCP Relay Agent, and a TFTP server (if you want to use TFTP)

Both DHCP and TFTP can reside on the same server. As the DHCP client, the Services Router sends a request through the DHCP Relay Agent to the DHCP server to receive an IP address. When the server assigns the IP address, the DHCP Relay Agent sends the IP address to the Services Router. The router now has a temporary address that is taken from a DHCP pool of addresses.

Based on the IP address and based on the configuration file that it is looking up, the Services Router initiates a TFTP request. The request advertises the IP address of the Services Router and requests the configuration file. The TFTP server processes the request and sends the configuration file back to the Services Router via FTP. The Services Router then loads the configuration file.

Autoinstallation Requirements for Service Providers

As is the normal practice, ensure that you (the service provider) have installed a bootstrap configuration with the following characteristics on several Services Routers:

- Configuration under autoinstallation, and specification that autoinstallation must be enabled

  If the destination URL is not specified, but DHCP is configured correctly, autoinstallation still works properly.

- Configuration that specifies the Services Router interface on which to search for this configuration file

Enabling Autoinstallation with the CLI

To enable autoinstallation on a Services Router from the CLI:

1. Enter configuration mode, and issue the edit system autoinstallation command:

```
user@host> configure

Entering configuration mode
[edit]

user@host# edit system autoinstallation

[edit system autoinstallation]
```

2. Specify the URL or FTP site of the configuration file that you want to automatically install on the Services Router:

```
user@host# set configuration-servers url
```

   If the destination URL is not specified, but DHCP is configured correctly, autoinstallation still works properly.

3. Specify the interface that the Services Router will use to send out and receive requests, and specify the IP address procurement protocol—bootp, rarp, or slarp. Typically, the interface is fe-0/0/0. This interface needs to be connected to a network that provides access to a DHCP server for IP address resolution.

   For example, the following command configures the fe-0/0/0 interface for autoinstallation by means of BOOTP:

```
user@host# set interfaces fe-0/0/0 bootp
```

Verifying Basic Connectivity

To verify that the Services Router has the settings you configured, perform the following task.

Displaying Basic Connectivity Configurations

Purpose Verify the configuration of basic connectivity. Because the basic connectivity settings appear in different places in the configuration hierarchy, displaying the entire configuration at once makes viewing the settings easier.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the show command. The following sample output displays the sample values configured in Table 22. Your output displays the values you set.

Sample Output

```
system {
            host-name routera;
            domain-name lab.router.net;
            domain-search [ lab.router.net router.net ];
            backup-router 192.168.2.44;
            time-zone America/Los_Angeles;
  root-authentication {
            ssh-rsa "ssh-rsa AAAAB3Nza...D9Y2gXF9ac==root@routera.lab.router.net";
  }
  name-server {
            10.148.2.32;
  }
  services {
```

```
    }
    ntp {
                server 10.148.2.21;
    }
}
interfaces {
    fe-0/0/0 {
        unit 0 {
            family inet {
                address 192.168.1.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.1.24/32;
            }
        }
    }
}
```

What It Means The output shows the configuration of basic connectivity. Verify that the values displayed are correct for your Services Router. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 136.

Chapter 5
Managing J-series Licenses

To enable some JUNOS software features and use additional ports on a J-series Services Router, you must purchase, install, and manage separate software licenses. The presence on the router of the appropriate software license keys (passwords) determines the features and ports you can configure and use.

For information about how to purchase J-series software licenses, contact your Juniper Networks sales representative.

This chapter contains the following topics:

J-series License Overview

The J-series set of licenses is composed of two primary types: feature licenses and port licenses. Each type of license is valid for only a single Services Router. To manage the licenses, you must understand the components of a license key.

This section contains the following topics:

Software Feature Licenses

Each feature license is tied to exactly one software feature, and that license is valid for exactly one Services Router. Table 23 lists the Services Router software features that require licenses.

Table 23: J-series Services Router Software Feature Licenses

Licensed Software Feature	License Name
Stateful Firewall Filters and NAT	
Stateful firewall and Network Address Translation (NAT) on the J2300 platform—all configuration statements within the [edit services stateful-firewall] hierarchy.	J2300 Services Router Software License for Stateful Firewall
Stateful firewall and NAT on the J4300 platform—all configuration statements within the [edit services stateful-firewall] hierarchy.	J4300 Services Router Software License for Stateful Firewall
Stateful firewall and NAT on the J6300 platform—all configuration statements within the [edit services stateful-firewall] hierarchy.	J6300 Services Router Software License for Stateful Firewall
IPSec VPN Tunneling	
IPSec VPN tunneling on the J2300 platform—all configuration statements within the [edit services ipsec-vpn] hierarchy.	J2300 Services Router Software License for IPSec Tunneling
IPSec VPN tunneling on the J4300 platform—all configuration statements within the [edit services ipsec-vpn] hierarchy.	J4300 Services Router Software License for IPSec Tunneling
IPSec VPN tunneling on the J6300 platform—all configuration statements within the [edit services ipsec-vpn] hierarchy.	J6300 Services Router Software License for IPSec Tunneling
Traffic Analysis	
J-Flow traffic analysis—all configuration statements within the [edit forwarding-options sampling] and [edit forwarding-options accounting] hierarchies.	J-series Services Router Software License for J-Flow Traffic Analysis
BGP Route Reflectors	
Advanced Border Gateway Protocol (BGP) features that enable route reflectors—all configuration statements within the [edit protocols bgp cluster] hierarchy. BGP clusters allow routers to act as route reflectors by enabling the readvertising of BGP routes to internal peers.	J-series Services Router Software License for Advanced Border Router Protocol Support

Port Licenses

Each port license is tied to exactly one licensed port, and that license is valid for exactly one Services Router. To enable multiple ports, you must have a license for each licensed port.Table 24 lists the additional Services Router port licenses.

Table 24: J-series Services Router Port Licenses

Licensed Port	License Name
T1	
Additional port on a T1 Physical Interface Module (PIM).	J-series Services Router Software License for One Additional T1 Port
E1	

Licensed Port	License Name
Additional port on a E1 PIM.	J-series Services Router Software License for One Additional E1 Port
Serial	
Additional port on a serial PIM.	J-series Services Router Software License for One Additional Serial Port
Fast Ethernet	
Additional port on a Fast Ethernet PIM.	J-series Services Router Software License for One Additional Fast Ethernet Port

The LAN ports (fe-0/0/0 and fe-0/0/1) do not require port licenses.

Additionally, one port per PIM can be configured without a port license. A port license is required only if you configure more than one port on a particular PIM.

License Key Components

A license key consists of two parts:

- License ID—Alphanumeric string that uniquely identifies the license key. When a license is generated, it is given a license ID.
- License data—Block of binary data that defines and stores all license key objects.

For example, in the following typical license key, the string li29183743 is the license ID, and the trailing block of data is the license data:

```
li29183743 4ky27y acasck 82fsj6 jzsn4q ix8i8d adj7kr
                  8uq38t ix8i8d jzsn4q ix8i8d 4ky27y acasck
                  82fsj6 ii8i7e adj7kr 8uq38t ks2923 a9382e
```

The license data defines the device ID for which the license is valid and the version of the license.

Before You Begin

Before you begin managing the J-series licenses, complete the following tasks:

- Purchase the licenses you require.
- Establish basic connectivity. See “Establishing Basic Connectivity” on page 47.

Managing J-series Licenses with the J-Web Interface

To manage licenses with the J-Web interface, you perform the following tasks:

- Adding New Licenses with the J-Web Interface on page 73

- Deleting Licenses with the J-Web User Interface on page 74
- Displaying License Keys with the J-Web Interface on page 74
- Downloading Licenses with the J-Web Interface on page 74

Figure 24 shows the J-Web Licenses page.

Figure 24: Licenses Page

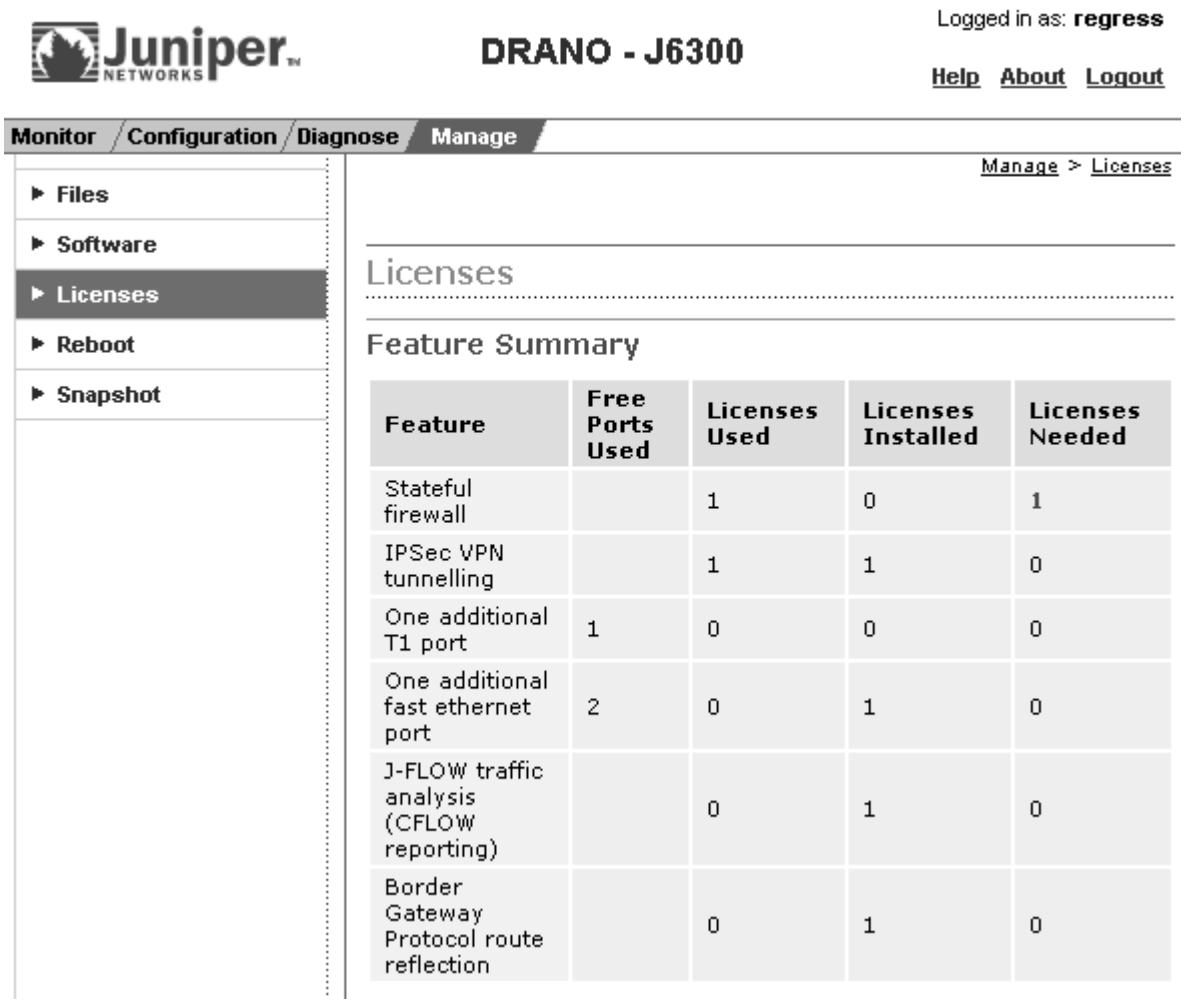

The Licenses page displays a summary of licensed features that are configured on the Services Router and a list of the licenses that are installed on the router. The information on the license management page is summarized in Table 25.

Table 25: Summary of License Management Fields

Field Name	Definition
Feature Summary	
Feature	Name of the licensed feature or port: ■ J-series licenses listed in Table 23 and Table 24 ■ All features—All-inclusive licenses
Licenses Used	Number of licenses currently being used on the router. Usage is determined by the configuration on the router. If a port license exists and that port is configured, the license is considered used.
Licenses Installed	Number of licenses installed on the router for the particular feature or port.
Licenses Needed	Number of licenses required for legal of use the feature or port. Usage is determined by the configuration on the router: ■ If a feature is configured and the license for that feature is not installed, a single license is needed. ■ If one or more ports are configured beyond the number of licenses installed on the router, a single license is needed for each additional configured port.
Installed Licenses	
ID	Unique alphanumeric ID of the license.
State	Valid—The installed license key is valid. Invalid—The installed license key is not valid.
Version	Numeric version number of the license key.
Enabled Features	Name of the feature that is enabled with the particular license.

Adding New Licenses with the J-Web Interface

To add a new license key on a Services Router with the J-Web license manager:

1. In the J-Web interface, select **Manage > Licenses**.
2. Under Installed Licenses, click **Add** to add a new license key.
3. Do *one* of the following, using a blank line to separate multiple license keys:
   - In the License File URL box, type the full URL to the destination file containing the license key to be added.
   - In the License Key Text box, paste the license key text, in plain-text format, for the license to be added.
4. Click **OK** to add the license key.

5. Go on to “Verifying J-series License Management” on page 76.

Deleting Licenses with the J-Web User Interface

To delete one or more license keys from a Services Router with the J-Web license manager:

1. In the J-Web interface, select **Manage > Licenses**.
2. Select the check box of the license or licenses you want to delete.
3. Click **Delete**.
4. Go on to “Verifying J-series License Management” on page 76.

Displaying License Keys with the J-Web Interface

To display the license keys installed on a Services Router with the J-Web license manager:

1. In the J-Web interface, select **Manage > Licenses**.
2. Under Installed Licenses, click **Display Keys** to display all the license keys installed on the router.

   A screen displaying the license keys in text format appears. Multiple licenses are separated by a blank line.
3. Go on to “Verifying J-series License Management” on page 76.

Downloading Licenses with the J-Web Interface

To download the license keys installed on the Services Router with the J-Web license manager:

1. In the J-Web interface, select **Manage > Licenses**.
2. Under Installed Licenses, click **Download Keys** to download all the license keys installed on the router to a single file.
3. Select **Save it to disk** and specify the file to which the license keys are to be written.
4. Go on to “Verifying J-series License Management” on page 76.

Managing J-series Licenses with the CLI

To manage the J-series licenses with the CLI, perform the following tasks.

- Adding New Licenses with the CLI on page 75
- Deleting a License with the CLI on page 75
- Saving License Keys with the CLI on page 76

Adding New Licenses with the CLI

To add a new license key to the Services Router with the CLI:

1. Enter operational mode in the CLI.
2. Enter one of the following CLI commands:
   - To add a license key from a file or URL, enter the following command, specifying the filename or the URL where the key is located:

     request system license add *filename* | *url*

   - To add a license key from the terminal, enter the following command:

     request system license add terminal
3. When prompted, enter the license key, separating multiple license keys with a blank line.

   If the license key you enter is invalid, an error is generated when you press Ctrl-D to exit license entry mode.
4. Go on to “Verifying J-series License Management” on page 76.

Deleting a License with the CLI

To delete a license key from the Services Router with the CLI:

1. Enter operational mode in the CLI.
2. Enter the following command for each license, specifying the license ID. You can delete only one license at a time.

   request system license delete *license-id*
3. Go on to “Verifying J-series License Management” on page 76.

Saving License Keys with the CLI

To save the licenses installed on the Services Router to a file with the CLI:

1. Enter operational mode in the CLI.
2. To save the installed license keys to a file or URL, enter the following command:

   request system license save *filename* | *url*

   For example, the following command saves the installed license keys to a file named license.config:

   request system license save ftp://user@host/license.conf

3. Go on to "Verifying J-series License Management" on page 76.

Verifying J-series License Management

To verify J-series license management, perform these tasks:

- Displaying Installed Licenses on page 76
- Displaying License Usage on page 77
- Displaying Installed License Keys on page 78

Displaying Installed Licenses

Purpose Verify that the expected licenses are installed and active on the Services Router.

Action From the CLI, enter the show system license command.

Sample Output

```
user@router> show system license

License identifier: li29183743
  State: valid
  License version: 2
  Valid for device: jp47859620

License identifier: li48293123
  State: valid
  License version: 2
  Valid for device: jp47859620
  Features:
    firewall         - Stateful firewall

License identifier: li72194673
  State: valid
  License version: 2
  Valid for device: jp47859620
```

```
  Features:
    if-t1-4          - Four additional T1 ports

License identifier: li41597793
  State: valid
  License version: 2
  Valid for device: jp47859620
  Features:
    ipsec-vpn        - IPSec VPN tunnelling
```

What It Means The output shows a list of the licenses installed on the Services Router. Verify the following information:

- Each license is present. Licenses are listed in ascending alphanumeric order by license ID.
- The state of each license is valid.

  A state of invalid indicates that the license key is not a valid license key. Either it was entered incorrectly or it is not valid for the specific device.
- The feature for each license is the expected feature. The features enabled are listed by license. An all-inclusive license has All features listed.

Displaying License Usage

Purpose Verify that the licenses fully cover the feature configuration on the Services Router.

Action From the CLI, enter the show system license usage command.

Sample Output

```
user@router> show system license usage

  Feature name     Licensed    Used      Needed
  firewall                1       1           0
  ipsec-vpn               1       0           0
  if-t1                   4       3(+1)       0
  if-se                   4       6(+2)       2
  j-flow                  1       1           0
  bgp-reflection          0       1           1
```

What It Means The output shows a list of the licenses installed on the Services Router and how they are used. Verify the following information:

- Each licensed feature and port is present. Features and ports are listed in ascending alphabetical order by license name. The number of licenses is shown in the first column. Check that the appropriate number of licenses is installed.
- The number of used licenses matches what is configured. If a licensed feature or port is configured, the feature or port is considered used. The sample output

shows that stateful firewall, J-Flow, and BGP route reflection are configured. Additionally, three T1 interfaces and six serial interfaces are configured.

If free port licenses are being used, the number of free licenses being used is listed in parentheses next to the number of used licenses. The sample output shows that the user has configured four T1 interfaces (three licensed interfaces and one free interface).

- A license is installed on the Services Router for each configured feature and port. For every feature or port configured that does not have a license, one license is needed.

  The sample output shows that the user has configured eight serial interfaces (six licensed interfaces and two free interfaces). This configuration requires six purchased licenses, but only four have been purchased. An additional two licenses are required to be in compliance with license agreements.

Displaying Installed License Keys

Purpose Verify the license keys installed on the Services Router.

Action From the CLI, enter the show system license keys command.

Sample Output

```
user@router> show system license keys


li29183743 jzsn4q ix8i8d 4ky27y jzsn4q ix8i8d adj7kr
           8uq38t 82fsj6 ii8i7e adj7kr 82fsj6 acasck
           ix8i8d 4ky27y acasck 8uq38t ks2923 a938

li48293123 4ky27y acasck 82fsj6 jzsn4q ix8i8d eksi2r
           8uq38t ix8i8d jzsn4q ix8i8d 4ky27y acasck
           82fsj6 ii8i7e adj7kr 8uq38t ks2923 a9382e

li83474929 dkdis8 adj7kr 4ky27y aclsck 82fsj6 jzsn4q
           8uq38t jzsn4q 9dk2i2 ii3i8d akd239 ks2923
           492idf oo8i7e adj7kr 8u3892 3ksio
```

What It Means The output shows a list of the license keys installed on the Services Router. Verify that each expected license key is present.

Chapter 6
Configuring Network Interfaces

Each Services Router can support types of interfaces suited to different functions. The router uses network interfaces to transmit and receive network traffic. For network interfaces to operate, you must configure properties such as logical interfaces, the encapsulation type, and certain settings specific to the interface type.

In addition to network interfaces, Services Routers uses permanent interfaces for internal communication, such as the services interfaces that provide additional features for regulating and manipulating traffic. For information about one of these interfaces, see "Loopback Address" on page 50.

This chapter includes the following topics. For more information about interfaces, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Network Interfaces Terms

To understand Services Router network interfaces, become familiar with the terms defined in Table 26.

Table 26: Network Interfaces Terms

Term	Definition
alternate mark inversion (AMI)	Original method of formatting T1 and E1 data streams.
binary 8-zero substitution (B8ZS)	Improved method of formatting T1 and E1 data streams, in which a special code is substituted whenever 8 consecutive zeros are sent over the link.
Challenge Handshake Authentication Protocol (CHAP)	Protocol that authenticates remote users. CHAP is a server-driven, three-step authentication method that depends on a shared secret password residing on both the server and the client.

Term	Definition
checksum	See *frame checksum sequence*.
channel service unit (CSU)	Unit that connects a digital telephone line to a multiplexer or other signal service.
Cisco HDLC	Cisco High-level Data Link Control protocol. Proprietary Cisco encapsulation for transmitting LAN protocols over a WAN. HDLC specifies a data encapsulation method on synchronous serial links by means of frame characters and checksums. Cisco HDLC enables the transmission of multiple protocols.
clock source	Source of the consistent, periodic signal used by a router to synchronize data communication and processing tasks.
CSU compatibility mode	Subrate on a T3 interface that allows a Services Router to connect with a channel service unit (CSU) with proprietary multiplexing at the remote end of the line. Subrating a T3 interface reduces the maximum allowable peak rate by limiting the payload encapsulated by the High-level Data Link Control protocol (HDLC).
data-link connection identifier (DLCI)	Identifier for a Frame Relay virtual connection, also called a logical interface.
data service unit (DSU)	Unit that connects a data terminal equipment (DTE) device—in this case, a Services Router—to a digital telephone line.
data terminal equipment (DTE)	RS-232 interface that a Services Router uses to exchange information with a serial device.
DS3 interface	Another name for a T3 interface.
data inversion	Transmission of all data bits in the data stream so that zeros are transmitted as ones and ones are transmitted as zeros. Data inversion is normally used only in alternate mark inversion (AMI) mode to guarantee ones density in the transmitted stream.
E1 interface	Physical WAN interface for transmitting signals in European digital transmission (E1) format. The E1 signal format carries information at a rate of 2.048 Mbps and can carry 32 channels of 64 Kbps each.
encapsulation type	Type of protocol header in which data is wrapped for transmission.
Fast Ethernet interface	Physical LAN interface for transmitting data at 100 Mbps. Fast Ethernet, also known as 100Base-T, additionally supports standard 10Base-T Ethernet transmission.
Flexible PIM Concentrator (FPC)	Logical identifier for a Physical Interface Module (PIM) installed on a Services Router.
frame check sequence (FCS)	Calculation that is added to a frame to control errors in High-level Data Link Control (HDLC), Frame Relay, and other data link layer protocols.
Frame Relay	Efficient WAN protocol that does not require explicit acknowledgement of each frame of data. Frame Relay allows private networks to reduce costs by sharing facilities between the end-point switches of a network managed by a Frame Relay service provider. Individual data link connection identifiers (DLCIs) are assigned to ensure that customers receive only their own traffic.
fractional E1	Service also known as channelized E1, in which a 2.048-Mbps E1 link is subdivided into 32 DS0 time slots (channels) in which time slot 0 is reserved. The individual channels or groups of channels connect to different destinations, and customers pay for only the channels used and not for the entire line.
fractional T1	Service also known as channelized T1, in which a 1.544-Mbps T1 link is subdivided into 24 DS0 time slots (channels) in which time slot 0 is reserved. The individual channels or groups of channels connect to different destinations, and customers pay for only the channels used and not for the entire line.

Term	Definition
High-level Data Link Control	International Telecommunication Union (ITU) standard for a bit-oriented data link layer protocol on which most other bit-oriented protocols are based.
hostname	Name assigned to the Services Router during initial configuration.
logical interface	Virtual interface that you create on a physical interface to identify its connection. Creating multiple logical interfaces allows you to associate multiple virtual circuits, data line connections, or virtual LANs (VLANs) with a single interface device.
maximum transmission unit (MTU)	Limit on the segment size that a network can transmit.
Physical Interface Module (PIM)	Network interface card that is fixed or can be interchangeably installed on a Services Router to provide the physical connections to a LAN or WAN, receiving incoming packets and transmitting outgoing packets. A PIM contains *one* of the following interfaces or sets of interfaces: ■ Two Fast Ethernet LAN interfaces ■ Two T1 or two E1 WAN interfaces ■ Single T3 (DS3) WAN interface (J6300 model only) ■ Two serial interfaces
Point-to-Point Protocol (PPP)	Link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration.
serial interface	Physical LAN interface for transmitting data between computing devices. A Services Router has two types of serial interfaces: ■ Asynchronous serial interface—Console port, with speeds up to 110.5 Kbps. The console port supports an RS-232 (EIA-232) standard serial cable with a 25-pin (DB-25) connector. ■ Synchronous serial interface—Port that transmits packets to and from, for example, a T1 device or microwave link, at speeds up to 8 Mbps. You cannot use this serial interface to connect a console. Services Router synchronous serial interfaces support the following cable types: ▪ V.35—Serial cable with a 34-pin connector for speeds up to 8 Mbps ▪ RS-232—(EIA-232) Standard serial cable with a 25-pin (DB-25) connector for speeds up to 110.5 Kbps ▪ RS-422/449—(EIA-449) Serial cable with a 37-pin (DB-37) connector, for RS-422 and RS-423 interfaces ▪ X.21—Standard serial cable, popular in Europe, with a 15-pin (DB-15) connector ▪ RS-530—(EIA-530) Serial cable with a 25-pin connector for higher speeds than RS-232 For cable details, see “Network Cable Specifications and Connector Pinouts” on page 551.
T1 interface	Physical WAN interface for transmitting digital signals in the T-carrier system used in the United States, Japan, and Canada. The T1 signal format carries 24 pulse code modulation (PCM) signals using time-division multiplexing (TDM) at an overall rate of 1.544 Mbps.
T3 interface	Physical WAN interface for transmitting digital signals in the T-carrier system used in the United States, Japan, and Canada. T3 signals are formatted like T1 signals, but carry information at the higher rate of 44.736 Mbps. T3 is also known as DS3.

Interfaces Overview

This section contains the following topics:

Network Interface Types

J-series Services Routers support the following network interface types: E1, Fast Ethernet, serial, T1, and T3.

T3 interfaces, which are also known as DS3 interfaces, are supported on J6300 Services Routers only.

Interfaces and Interface Naming

The interfaces on a Services Router are used for networking and services. Most interfaces are configurable, but some internally generated interfaces are not configurable. Each interface has a unique name that identifies its type and location and indicates whether it is a physical interface or an optional logical unit created on a physical interface:

- The name of each interface on the router has the following format, to identify the physical device that corresponds to a single physical network connector:

  type-*FPC*/*PIM*/*port*

- Network interfaces that are fractionalized into time slots include a virtual DS0 channel number in the name, preceded by a colon (:):

  type-*FPC*/*PIM*/*port*:*channel*

- Each logical interface has an additional logical unit identifier, preceded by a period (.):

  type-*FPC*/*PIM*/*port*:<*channel*>.*logical unit*

For example, **e1-5/0/0** is the E1 interface on port 0 of FPC 5, **e1-5/0/0:15** is channel 15 on that interface, and **e1-5/0/0:15.0** is logical unit 0 on that channel.

The parts of an interface name are explained in Table 27.

Table 27: Interface Name Information

Interface Name Part	Meaning	Possible Values
type	Type of network medium that can connect to this interface.	■ dsc—Virtual interface that discards packets. ■ e1—E1 WAN interface. ■ fe—Fast Ethernet LAN interface. ■ gr, gre—Generic routing encapsulation (GRE) interface for tunnel services. This interface is internally generated and not configurable. ■ ip, ipip—IP-over-IP interface. This interface is internally generated and not configurable. ■ ls, lsi—Link services interface. This interface is internally generated and not configurable. ■ lo—Loopback interface. This interface is internally generated and also configurable. ■ mtun—Multicast GRE interface. This interface is internally generated and not configurable. ■ pd, pimd—Protocol Independent Multicast (PIM) decapsulator interface. This interface is internally generated and not configurable. ■ pe, pime—PIM encapsulator interface. This interface is internally generated and not configurable. ■ se—Serial interface (including RS-232, RS-422/449, RS-530, V.35, and X.21 interfaces). ■ t1—T1 WAN interface. ■ t3—T3 (also known as DS3) WAN interface. ■ tap—This interface is internally generated and not configurable.
FPC	Number of the Flexible PIM concentrator (FPC) on which the physical interface is located.	■ On a J2300 router, always 0. ■ On a J4300 or J6300 router, a value from 0 through 6.
PIM	Number of the PIM on which the physical interface is located. For Services Router interfaces, the FPC and PIM are the same physical unit, so the PIM has no number of its own.	Always 0.
port	Number of the port on a PIM on which the physical interface is located.	Either 0 or 1.

Interface Name Part	Meaning	Possible Values
channel	Number of the channel (time slot) on a fractional T1 or E1 interface.	■ On an E1 interface, a value from 0 through 32. The 0 and 1 time slots are reserved. ■ On a T1 interface, a value from 0 through 24. The 0 time slot is reserved.
logical unit	Number of the logical unit created on the physical interface.	A value from 0 through 16384.

Before You Begin

Before you configure network interfaces, you need to perform the following tasks:

- Install Services Router hardware. For more information, see "Installing and Connecting a Services Router" on page 35.
- Establish basic connectivity. For more information, see "Establishing Basic Connectivity" on page 47.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read "Interfaces Overview" on page 82.

Although not a requirement, you might also want to plan how you are going to use the various network interfaces before you start configuring them. You can see a list of the physical interfaces installed on the J-series Services Router by displaying the Quick Configuration page, as shown in Figure 25.

Configuring Network Interfaces with Quick Configuration

The Quick Configuration page allows you to configure network interfaces on a Services Router, as shown in Figure 25.

Figure 25: Quick Configuration Interfaces Page

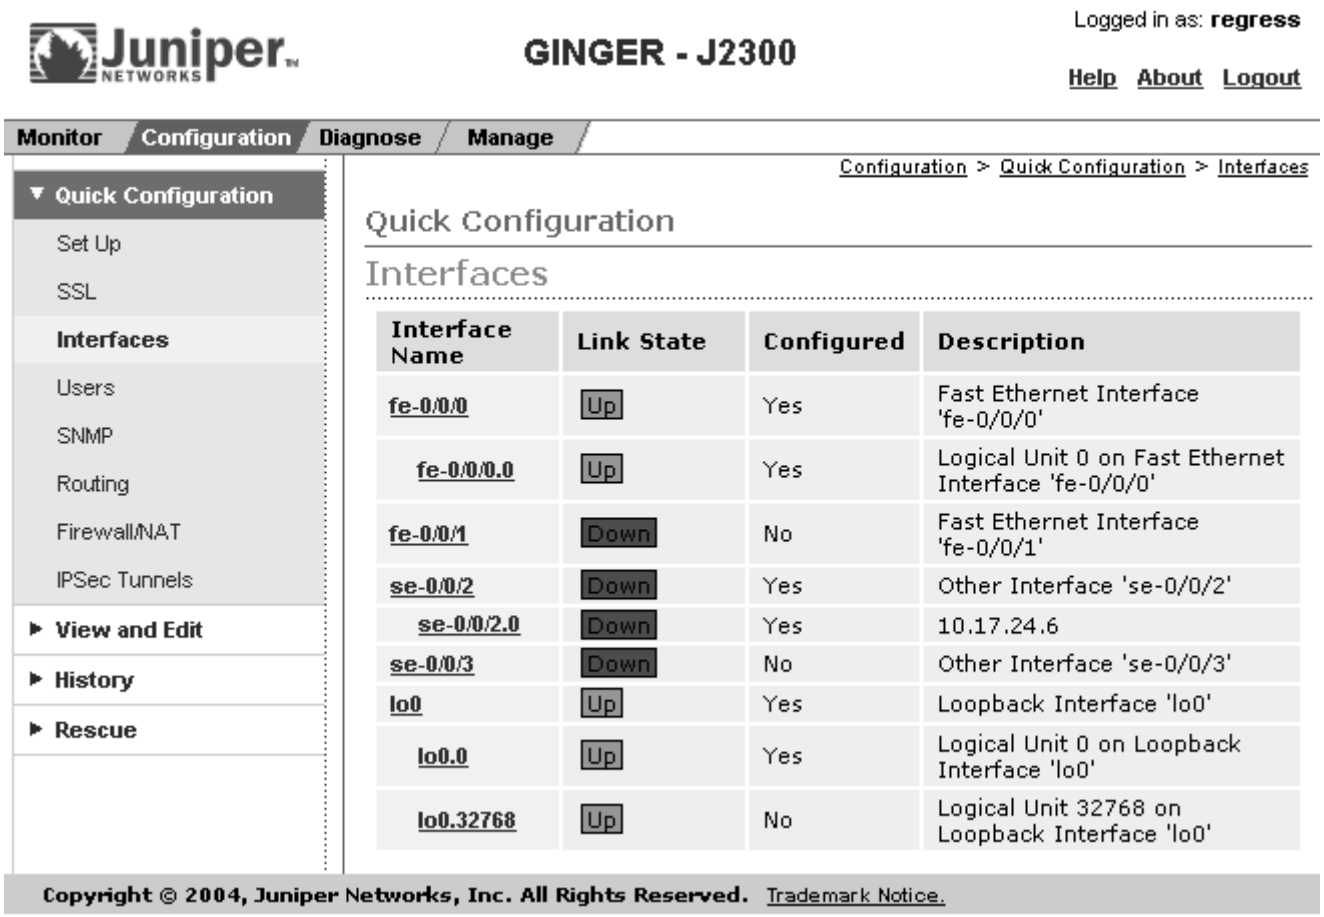

To configure a network interface with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Interfaces**. You can select **Interfaces** in the list under Router Configuration or from the left pane.

   A list of the network interfaces present on the Services Router is displayed, as shown in Figure 25. The third column indicates whether the interface has been configured.

2. To configure properties for a network interface, select the interface name and proceed with configuration as described in one of the following topics:

- Configuring an E1 Interface with Quick Configuration on page 86
- Configuring a Fast Ethernet Interface with Quick Configuration on page 89
- Configuring a T1 Interface with Quick Configuration on page 91

- Configuring a T3 Interface with Quick Configuration on page 95
- Configuring a Serial Interface with Quick Configuration on page 98

Configuring an E1 Interface with Quick Configuration

To configure properties on an E1 interface:

1. From the Quick Configuration page, as shown in Figure 25, select the E1 interface you want to configure.

   The properties you can configure on an E1 interface are displayed, as shown in Figure 26.

Figure 26: E1 Interfaces Quick Configuration Page

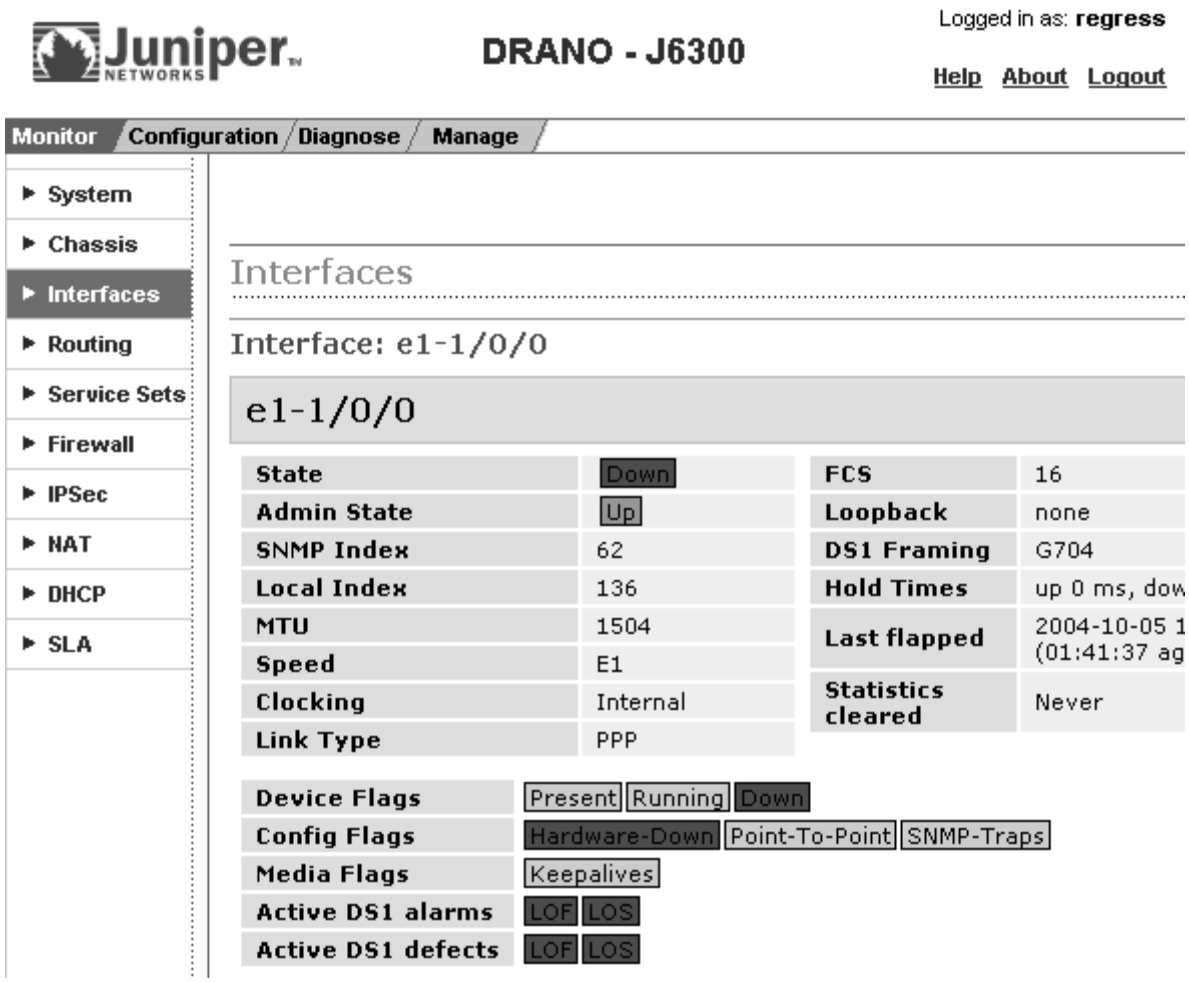

2. Enter information into the Quick Configuration page, as described in Table 28.
3. Click one of the following buttons:
   - To apply the configuration and stay in the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the E1 interface is configured correctly, see "Verifying Interface Configuration" on page 104.

Table 28: E1 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical E1 interface. You must define at least one logical unit for an E1 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click **Add**.
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	1. Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 2. Click **Add**.
Physical Interface Description	(Optional) Adds supplementary information about the physical E1 interface.	Type a text description of the E1 interface to more clearly identify it in monitoring displays.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the drop-down list, select the encapsulation for this E1 interface: ■ **PPP** ■ **Frame Relay** ■ **Cisco HDLC**

Field	Function	Your Action
Enable CHAP	Enables or disables CHAP authentication on an E1 interface with PPP encapsulation only.	■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the E1 interface use the Services Router's system hostname in CHAP challenge and response packets.	■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this E1 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this E1 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
E1 Options		
MTU	Specifies the maximum transmission unit size for the E1 interface.	Type a value between 256 and 9192 bytes. The default MTU for E1 interfaces is 1504.
Clocking	Specifies the transmit clock source for the E1 line.	From the drop-down list, select one of the following: ■ **internal**—Services Router's own system clock (the default) ■ **external**—Clock received from the E1 interface
Framing Mode	Specifies the framing mode for the E1 line.	From the drop-down list, select one of the following: ■ **g704**—The default ■ **g704–no-crc4**—G704 without cyclic redundancy check 4 (CRC4) ■ **unframed**—Unframed transmission format
Invert Data	Enables or disables data inversion. Data inversion is normally used only in alternate mark inversion (AMI) mode.	■ To enable, select the check box. ■ To disable, clear the check box.

Field	Function	Your Action
Timeslots	Specifies the number of time slots allocated to a fractional E1 interface. By default, an E1 interface uses all the time slots.	Type numeric values from 2 through 32. Separate discontinuous entries with commas, and use hyphens to indicate ranges. For example: 2,4,7–9
Frame Checksum	Specifies the number of bits in the frame checksum. A 32–bit checksum provides more reliable packet verification, but is not supported by some older equipment.	Select **16** or **32**. The default checksum is **16**.

Configuring a Fast Ethernet Interface with Quick Configuration

To configure properties on a Fast Ethernet interface:

1. From the Quick Configuration page, as shown in Figure 25, select the interface you want to configure.

   The properties you can configure on a Fast Ethernet interface are displayed, as shown in Figure 27.

Figure 27: Fast Ethernet Interfaces Quick Configuration Page

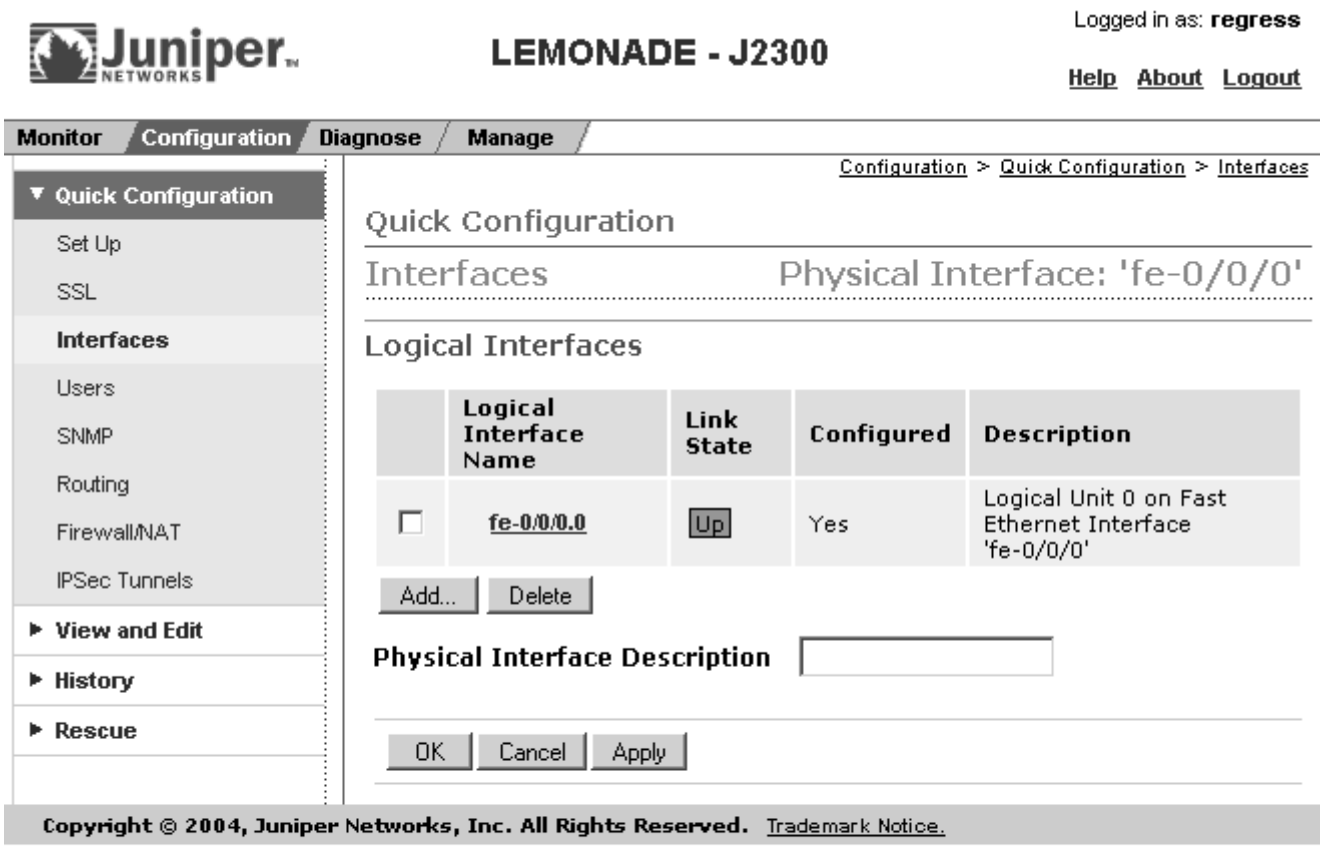

2. Enter information into the Quick Configuration page, as described in Table 29.
3. Click one of the following buttons:
   - To apply the configuration and stay in the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the Fast Ethernet interface is configured correctly, see “Verifying Interface Configuration” on page 104.

Table 29: Fast Ethernet Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical Fast Ethernet interface. You must define at least one logical unit for a Fast Ethernet interface. You can define multiple units if the encapsulation type is Frame Relay.	Click **Add**.
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	1. Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 2. Click **Add**.
Physical Interface Description	(Optional) Adds supplementary information about the physical Fast Ethernet interface.	Type a text description of the Fast Ethernet interface to more clearly identify it in monitoring displays.

Configuring a T1 Interface with Quick Configuration

To configure properties on a T1 interface:

1. From the Quick Configuration page, as shown in Figure 25, select the interface you want to configure.

   The properties you can configure on a T1 interface are displayed, as shown in Figure 28.

Figure 28: T1 Interfaces Quick Configuration Page

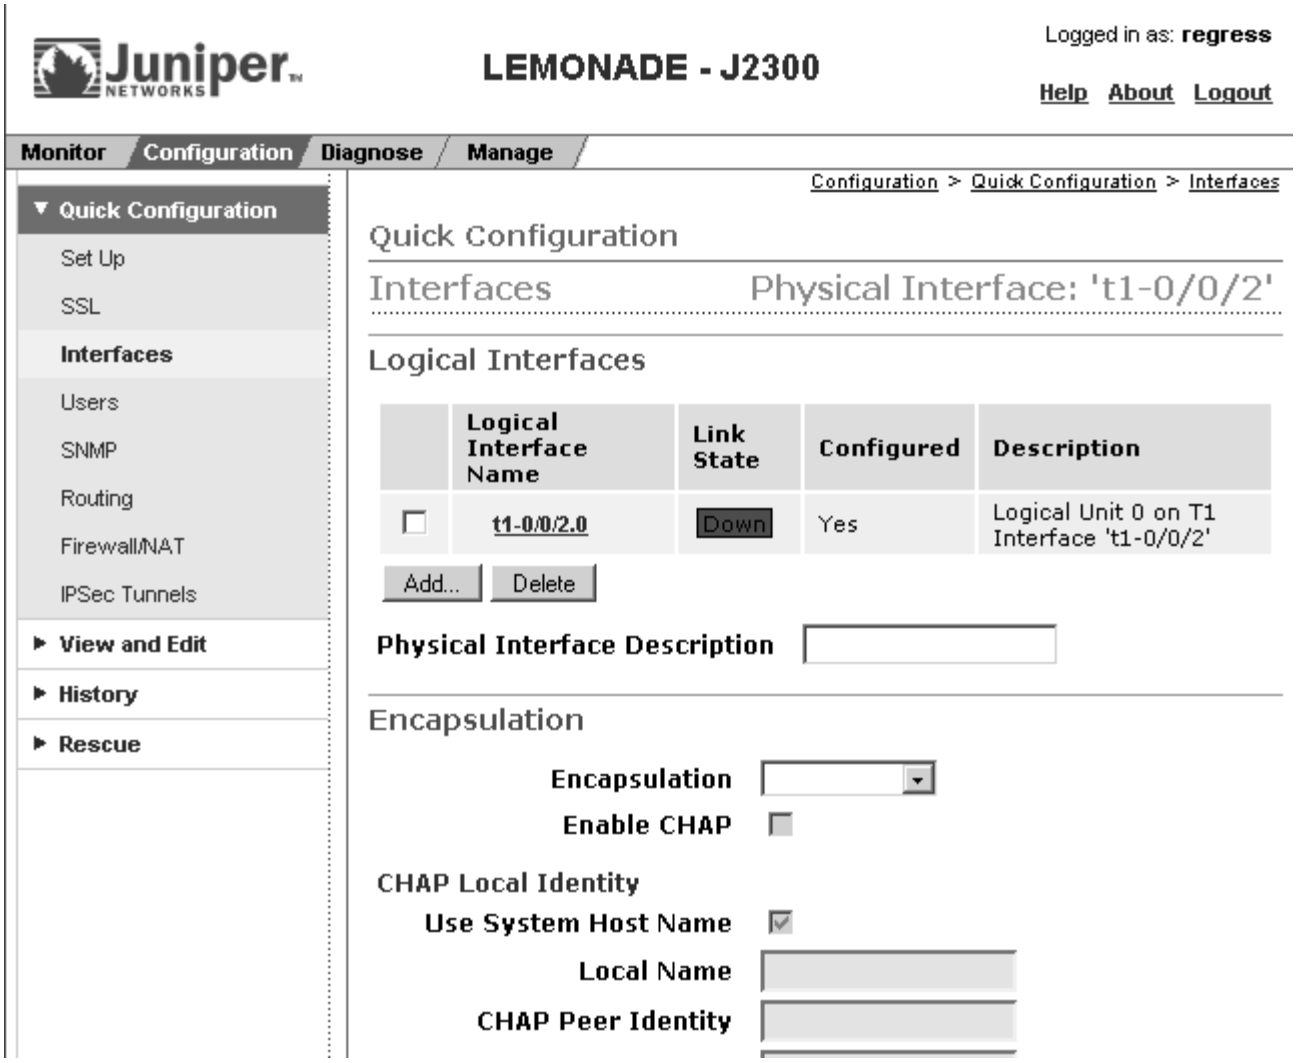

2. Enter information into the Quick Configuration page, as described in Table 30.

3. Click one of the following buttons:

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

4. To verify that the T1 interface is configured correctly, see “Verifying Interface Configuration” on page 104.

Table 30: T1 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical T1 interface. You must define at least one logical unit for a T1 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click **Add**.
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	1. Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 2. Click **Add**.
Physical Interface Description	(Optional) Adds supplementary information about the physical T1 interface.	Type a text description of the T1 interface to more clearly identify it in monitoring displays.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the drop-down list, select the encapsulation for this T1 interface: ■ **PPP** ■ **Frame Relay** ■ **Cisco HDLC**
Enable CHAP	Enables or disables CHAP authentication on a T1 interface with PPP encapsulation only.	■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the T1 interface use the Services Router's system hostname in CHAP challenge and response packets.	■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this T1 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this T1 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.

Field	Function	Your Action
T1 Options		
MTU	Specifies the maximum transmission unit size for the T1 interface.	Type a value between 256 and 9192 bytes. The default MTU for T1 interfaces is 1504.
Clocking	Specifies the transmit clock source for the T1 line.	From the drop-down list, select one of the following: ■ **internal**—Services Router's own system clock (the default) ■ **external**—Clock received from the T1 interface
Framing Mode	Specifies the framing mode for the T1 line.	From the drop-down list, select one of the following: ■ **esf**—Extended superframe (the default) ■ **sf**—Superframe
Line Encoding	Specifies the line encoding method.	From the drop-down list, select one of the following: ■ **ami**—Alternate mark inversion ■ **b8zs**—Binary 8 zero substitution (the default)
Byte Encoding	Specifies the byte encoding method.	From the drop-down list, select one of the following: ■ **nx56**—7 bits per byte ■ **nx64**—8 bits per byte (the default)
Invert Data	Enables or disables data inversion. Data inversion is normally used only in alternate mark inversion (AMI) mode.	■ To enable, select the check box. ■ To disable, clear the check box.
Timeslots	Specifies the number of time slots allocated to a fractional T1 interface. By default, a T1 interface uses all the time slots.	Type numeric values from 1 through 24. You can use any combination of time slots. To configure ranges, use hyphens. To configure discontinuous slots, use commas. For example: 1–5,10,24

Field	Function	Your Action
Frame Checksum	Specifies the number of bits in the frame checksum. A 32–bit checksum provides more reliable packet verification, but is not supported by some older equipment.	Select **16** or **32**. The default value is 16.
Line Buildout	Specifies the T1 cable length, in feet.	From the drop-down list, select one of the following cable lengths: ■ **0–132** (0 m–40 m) (the default) ■ **133–265** (40 m–81 m) ■ **266–398** (81 m–121 m) ■ **399–531** (121 m–162 m) ■ **532–655** (162m–200m)

Configuring a T3 Interface with Quick Configuration

To configure properties on a T3 (DS3) interface:

1. From the Quick Configuration page, as shown in Figure 25, select the interface you want to configure.

   The properties you can configure on a T3 interface are displayed, as shown in Figure 29.

Figure 29: T3 Interfaces Quick Configuration Page

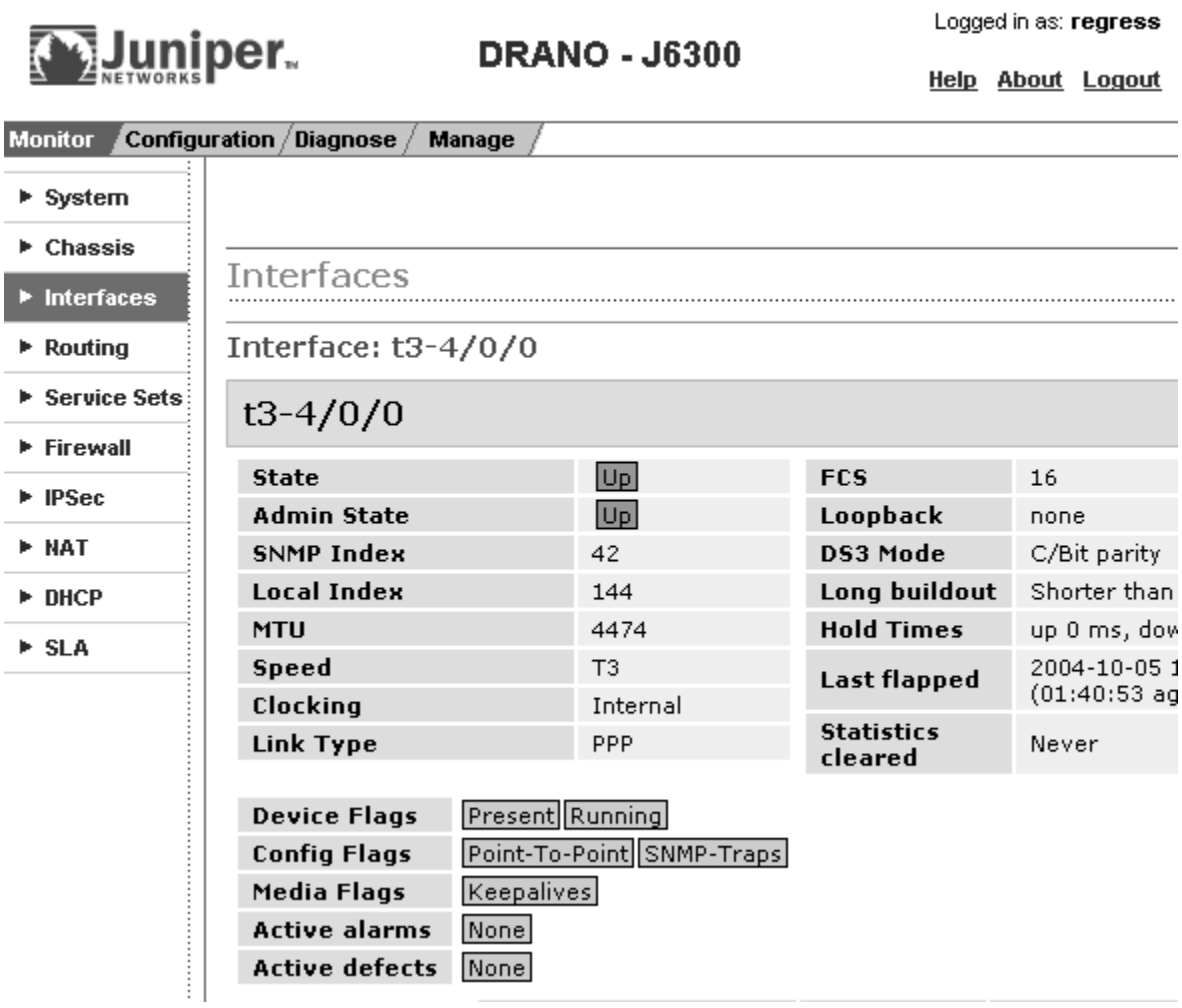

2. Enter information into the Quick Configuration page, as described in Table 31.

3. Click one of the following buttons:

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

4. To verify that the T3 interface is configured correctly, see “Verifying Interface Configuration” on page 104.

Table 31: T3 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical T3 interface. You must define at least one logical unit for a T3 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click **Add**.
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	1. Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 2. Click **Add**.
Physical Interface Description	(Optional) Adds supplementary information about the physical T3 interface.	Type a text description of the T3 interface to more clearly identify it in monitoring displays.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the drop-down list, select the encapsulation for this T3 interface: ■ **PPP** ■ **Frame Relay** ■ **Cisco HDLC**
Enable CHAP	Enables or disables CHAP authentication on a T3 interface with PPP encapsulation only.	■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the T3 interface use the Services Router's system hostname in CHAP challenge and response packets.	■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this T3 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this T3 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.

Field	Function	Your Action
T3 Options		
MTU	Specifies the maximum transmission unit size for the T3 interface.	Type a value between 256 and 9192 bytes. The default MTU for T3 interfaces is 4474.
Clocking	Specifies the transmit clock source for the T3 line.	From the drop-down list, select one of the following: ■ **internal**—Services Router's own system clock (the default) ■ **external**—Clock received from the T3 interface
C-Bit Parity	Enables or disables C-bit parity mode, which controls the type of framing that is present on the transmitted T3 signal.	■ To enable, select the check box. ■ To disable, clear the check box.
Frame Checksum	Specifies the number of bits in the frame checksum. A 32-bit checksum provides more reliable packet verification, but is not supported by some older equipment.	Select **16** or **32**. The default value is 16.
Long Buildout	Specifies a short or long cable length for copper-cable-based T3 interfaces. A long cable is longer than 225 feet (68.6m).	■ To enable long buildout, select the check box. ■ To disable long buildout, clear the check box.

Configuring a Serial Interface with Quick Configuration

To configure properties on a serial interface:

1. From the Quick Configuration page, as shown in Figure 25, select the interface you want to configure.

   The properties you can configure on a serial interface are displayed, as shown in Figure 30.

Figure 30: Serial Interfaces Quick Configuration Page

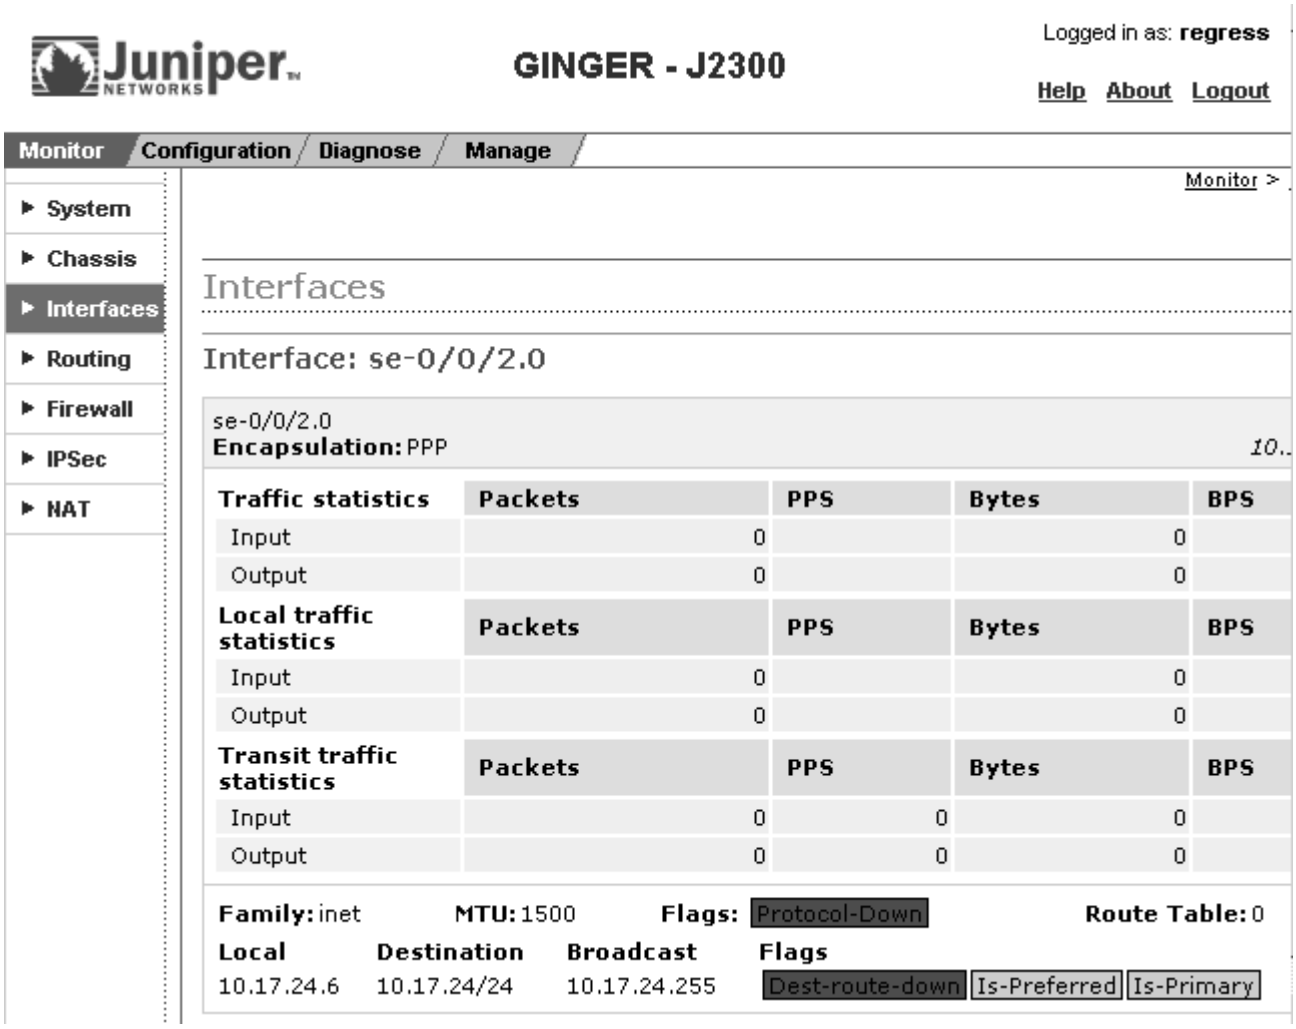

2. Enter information into the Quick Configuration page, as described in Table 32.
3. Click one of the following buttons:
   - To apply the configuration and stay in the Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the main configuration page, click **OK**.
   - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the serial interface is configured correctly, see "Verifying Interface Configuration" on page 104.

Table 32: Serial Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical serial interface. You must define at least one logical unit for a serial interface. You can define multiple units if the encapsulation type is Frame Relay.	Click **Add**.
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	1. Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 2. Click **Add**.
Physical Interface Description	(Optional) Adds supplementary information about the physical serial interface.	Type a text description of the serial interface to more clearly identify it in monitoring displays.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the drop-down list, select the encapsulation for this serial interface: ■ **PPP** ■ **Frame Relay** ■ **Cisco HDLC**
Enable CHAP	Enables or disables CHAP authentication on a serial interface with PPP encapsulation only.	■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the serial interface use the Services Router's system hostname in CHAP challenge and response packets.	■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this serial interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this serial interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.

Field	Function	Your Action
Serial Options		
MTU	Specifies the maximum transmission unit size for a serial interface.	Type a value between 256 and 9192 bytes. The default MTU for serial interfaces is 1504.
Clocking Mode	Specifies the clock source to determine the timing on serial interfaces.	From the drop-down list, select one of the following timing sources: ■ **dce**—Uses a transmit clock generated by the data circuit-terminating equipment (DCE) for the Services Router's DTE (the default). ■ **internal**—Uses the Services Router's internal clock. ■ **loop**—Uses the DCE's receive clock. This mode is the only one supported for X.21 serial interfaces.
Clock Rate	Specifies the line speed in kilohertz or megahertz for serial interfaces that use the DTE clocking mode.	From the drop-down list, select one of the following clock rates: ■ **1.2 KHz** ■ **2.4 KHz** ■ **9.6 KHz** ■ **19.2 KHz** ■ **38.4 KHz** ■ **56.0 KHz** ■ **64.0 KHz** ■ **72.0 KHz** ■ **125.0 KHz** ■ **148.0 KHz** ■ **250.0 KHz** ■ **500.0 KHz** ■ **800.0 KHz** ■ **1.0 MHz** ■ **1.3 MHz** ■ **2.0 MHz** ■ **4.0 MHz** ■ **8.0 MHz**

Configuring Network Interfaces with a Configuration Editor

To enable the interfaces installed on your Services Router to work properly, you must configure their properties. You can perform basic interface configuration using the J-Web Configuration Page, as described in "Configuring Network Interfaces with Quick Configuration" on page 84. You can perform the same configuration tasks using the J-Web or CLI configuration editors. In addition, you can configure a wider variety of options that are encountered less frequently.

You can perform the following tasks to configure interfaces:

- Adding a Network Interface with a Configuration Editor on page 102
- Deleting a Network Interface with a Configuration Editor on page 103

For information about using the J-Web and CLI configuration editors, see "Using J-series Configuration Tools" on page 127.

Adding a Network Interface with a Configuration Editor

To configure network interfaces for the Services Router:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 33.
3. When you are finished configuring the interface, click the **Commit** button or enter the commit command to commit the configuration.

Table 33: Adding an Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Interfaces** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Configuration > View and Edit > Edit Configuration**. 2. Next to Interfaces, click **Configure** or **Edit**.	From the top of the configuration hierarchy, enter edit interfaces
Create the new interface.	1. Next to Interface, click **Add new entry**. 2. Enter the name of the new interface in the Interface name box. Make sure the name conforms to the interface naming rules. For more information, see "Interfaces and Interface Naming" on page 82. 3. Click **OK**.	Create and name the interface: set *interface-name* Make sure the name conforms to the interface naming rules. For more information, see "Interfaces and Interface Naming" on page 82.

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the basic configuration for the new interface.	1. Under Interface Name in the table, click the name of the new interface. 2. Enter values in the other fields on this page if warranted. All these entries are optional, but you need to set values for Clocking and Encapsulation in particular if the default values are not suitable.	Enter values for physical interface properties as needed. Examples include changes to the default values for physical encapsulation or MTU.
Add values for interface-specific options. Most interface types have optional parameters that are specific to the interface type.	1. Under Nested configuration, click **Configure** for the appropriate interface type. 2. In the interface-specific page that appears, enter the values you need to supply or change the default values. 3. When you are finished, click **OK** to confirm your changes or **Cancel** to cancel them and return to the previous page.	1. From the [edit interfaces *interface-name*] hierarchy level, enter edit *interface*-options 2. Enter the statement for each interface-specific property for which you need to change the default value.
Add logical interfaces.	1. In the main Interface page for this interface, next to Unit, click **Add new entry**. 2. In the Unit page for logical interfaces that appears, enter a number from 0 through 16384 in the Interface unit number box. 3. Enter values in other fields as required for your network. 4. To configure protocol family values if needed, under Family, click **Configure** next to the appropriate protocol. 5. To access additional subordinate hierarchies under Nested configuration, click **Configure** next to any parameter you want to configure. 6. When you are finished, click **OK** to confirm your changes or **Cancel** to cancel them and return to the previous page.	1. From the [edit interfaces *interface-name*] hierarchy level, enter set unit *logical-unit-number* Replace *logical-unit-number* with a value from 0 through 16384. 2. Enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Deleting a Network Interface with a Configuration Editor

To delete an interface on a Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 34.

Table 34: Deleting an Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Interfaces** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Configuration > View and Edit > Edit Configuration**. 2. Next to Interfaces, click **Edit**.	From the top of the configuration hierarchy, enter edit interfaces
Select the interface you want to delete.	In the Interface table, under Interface name, click the name of the interface you want to delete.	Enter delete *interface-name*
Execute the selection.	1. Click the **Discard** button. 2. In the page that appears, select the appropriate radio button. If you have not made any previous changes, the only selection available is **Delete Configuration Below This Point**.	Commit the configuration change: commit

NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web Monitor and Quick Configuration pages.

Verifying Interface Configuration

To verify an interface configuration, perform these tasks:

- Verifying the Link State of All Interfaces on page 104
- Verifying Interface Properties on page 105

Verifying the Link State of All Interfaces

Purpose By using the ping tool on each peer address in the network, verify that all interfaces on the Services Router are operational.

Action For each interface on the Services Router:

1. In the J-Web interface, select **Diagnose > Ping Host**.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. Output appears on a separate page.

Sample Output

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

What It Means If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the time field. For more information about the output, see Table 82.

For more information about using the J-Web interface to ping a host, see "Using the J-Web Ping Host Tool" on page 218.

For information about the ping command, see "Using the ping Command" on page 226 or the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From the CLI, enter the show interfaces detail command.

Sample Output

```
user@host> show interfaces detail

Physical interface: fe-1/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 27, Generation: 17
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:87:44:9d, Hardware address: 00:90:69:87:44:9d
  Last flapped   : 2004-08-25 15:42:30 PDT (4w5d 22:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                   0                    0
    1 expedited-fo                   0                   0                    0
```

```
  2 assured-forw                 0                 0                 0
  3 network-cont                 0                 0                 0
Active alarms  : None
Active defects : None
```

What It Means The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do either of the following:
  - In the CLI configuration editor, delete the disable statement at the [edit interfaces *interface-name*] level of the configuration hierarchy.
  - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces >** *interface-name* page.
- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. The Last Flapped time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the clear interfaces statistics *interface-name* command.

For more information about show interfaces detail, see the *JUNOS Network and Services Interfaces Command Reference*.

Part 3

Using the J-series User Interfaces

Chapter 7

J-series User Interface Overview

You can use two user interfaces to monitor, configure, troubleshoot, and manage the Services Router—the J-Web interface and the JUNOS command-line interface (CLI). This chapter contains the following topics:

User Interface Overview

This section contains the following topics:

J-Web Overview

The J-Web graphical user interface (GUI) allows you to monitor, configure, troubleshoot, and manage the Services Router on a client by means of an HTTP- or HTTPS-enabled Web browser. The J-Web interface provides access to all the configuration statements supported by the router, so you can fully configure it without using the CLI.

The J-Web interface provides two methods of Services Router configuration:

- Quick Configuration
- Configuration editor

For more information, see “Comparison of Configuration Interfaces” on page 110.

In addition to configuration, you can use the J-Web interface to perform many monitoring, troubleshooting, and management tasks on the Services Router. For

example, to display a summary of routing table entries, click **Monitor** in the task bar, then click **Routing** in the side pane. The routes are displayed in the main pane.

For more information about the J-Web interface, see "Using the J-Web Interface" on page 112.

CLI Overview

The CLI is a straightforward command interface in which you type commands on a line and press Enter to execute them. The CLI provides command help, command completion, and Emacs-style keyboard sequences for moving around on the command line and scrolling through a buffer of recently executed commands.

The CLI has two modes:

- Operational mode—Complete set of commands to control the CLI environment, monitor and troubleshoot network connectivity, manage the Services Router, and enter configuration mode.
- Configuration mode—Complete set of commands to configure the Services Router. This guide refers to configuration mode as the *CLI configuration editor*. For more information, see "Comparison of Configuration Interfaces" on page 110.

For more information about the CLI, see "Using the Command-Line Interface" on page 117.

Comparison of Configuration Interfaces

Table 35 describes and compares the interfaces you can use to configure a Services Router.

Table 35: Services Router Configuration Interfaces

Services Router Interface	Description	Capabilities	Recommendations
J-Web Quick Configuration	Web browser pages for setting up the Services Router quickly and easily without configuring each statement individually. For example, use the Set Up Quick Configuration page to configure the Services Router for basic connectivity so you can manage it from the network.	Configure basic router services: ■ Setup ■ Secure Sockets Layer (SSL) ■ Interfaces ■ User access ■ SNMP notifications ■ Routing ■ Security firewalls and Network Address Translation (NAT) ■ IPSec tunnels	Use for basic configuration.
J-Web configuration editor	Web browser pages divided into panes in which you can do any of the following: ■ Expand the entire configuration hierarchy and click a configuration statement to view or edit. The main pane displays all the options for the statement, with a text box for each option. ■ Paste a complete configuration hierarchy into a scrollable text box, or edit individual lines. ■ Upload or download a complete configuration. ■ Roll back to a previous configuration.	Configure all router services: ■ System parameters ■ User access and accounting ■ Interfaces ■ SNMP network management ■ Routing options, including multicast routing ■ Routing protocols ■ Routing policies ■ Secure access ■ Service interfaces, including stateful firewalls and virtual private networks (VPNs) ■ Traffic engineering, including Multiprotocol Label Switching (MPLS) and class-of-service (CoS) packet prioritization ■ Chassis properties	Use for complete configuration if you are not familiar with the JUNOS CLI or prefer a graphical interface.
CLI configuration editor	Interface in which you do either of the following: ■ Type commands on a line and press Enter to create a hierarchy of configuration statements. ■ Create an ASCII text file that contains the statement hierarchy. ■ Upload a complete configuration, or roll back to a previous configuration.		Use for complete configuration if you know the JUNOS CLI or prefer a command interface.

Before You Begin

Before you start the user interface, you must perform the initial Services Router configuration described in "Establishing Basic Connectivity" on page 47. After the initial configuration, you use your username and password, and the hostname or IP address of the router, to start the user interface.

Using the J-Web Interface

This section contains the following topics:

Starting the J-Web Interface

To start the J-Web interface:

1. Launch your HTTP- or HTTPS-enabled Web browser.

   To use HTTPS, you must have installed the certificate provided by the Services Router. For more information, see "Managing J-series Licenses" on page 69.

2. After http:// or https:// in your Web browser, type the hostname or IP address of the Services Router and press Enter.

   The J-Web login page appears.

3. On the login page, type your username and password, and click **Log In**.

   To correct or change the username or password you typed, click **Reset**, type the new entry or entries, and click **Log In**.

   The J-Web **Quick Configuration > Set Up** (see Figure 31) or **Monitor > System** page appears.

To explicitly terminate a J-Web session at any time, click **Logout** in the top pane.

J-Web Layout

Each page of the J-Web interface is divided into the following panes shown in Figure 31 and Figure 32:

- Top pane—Displays identifying information and links.
- Main pane—Location where you monitor, configure, diagnose, and manage the Services Router by entering information in text boxes, making selections, and clicking buttons.
- Side pane—Displays suboptions of the Monitor, Configuration, Diagnose, or Manage task currently displayed in the main pane. Click a suboption to access it in the main pane.
- Bottom pane—Displays copyright and trademark information.

The layout of the panes allows you to quickly navigate through the interface. Table 36 summarizes the elements of the J-Web interface.

You navigate the J-Web interface, move forward and backward, scroll pages, and expand and collapse elements as you do in a typical Web browser interface.

Figure 31: J-Web Layout

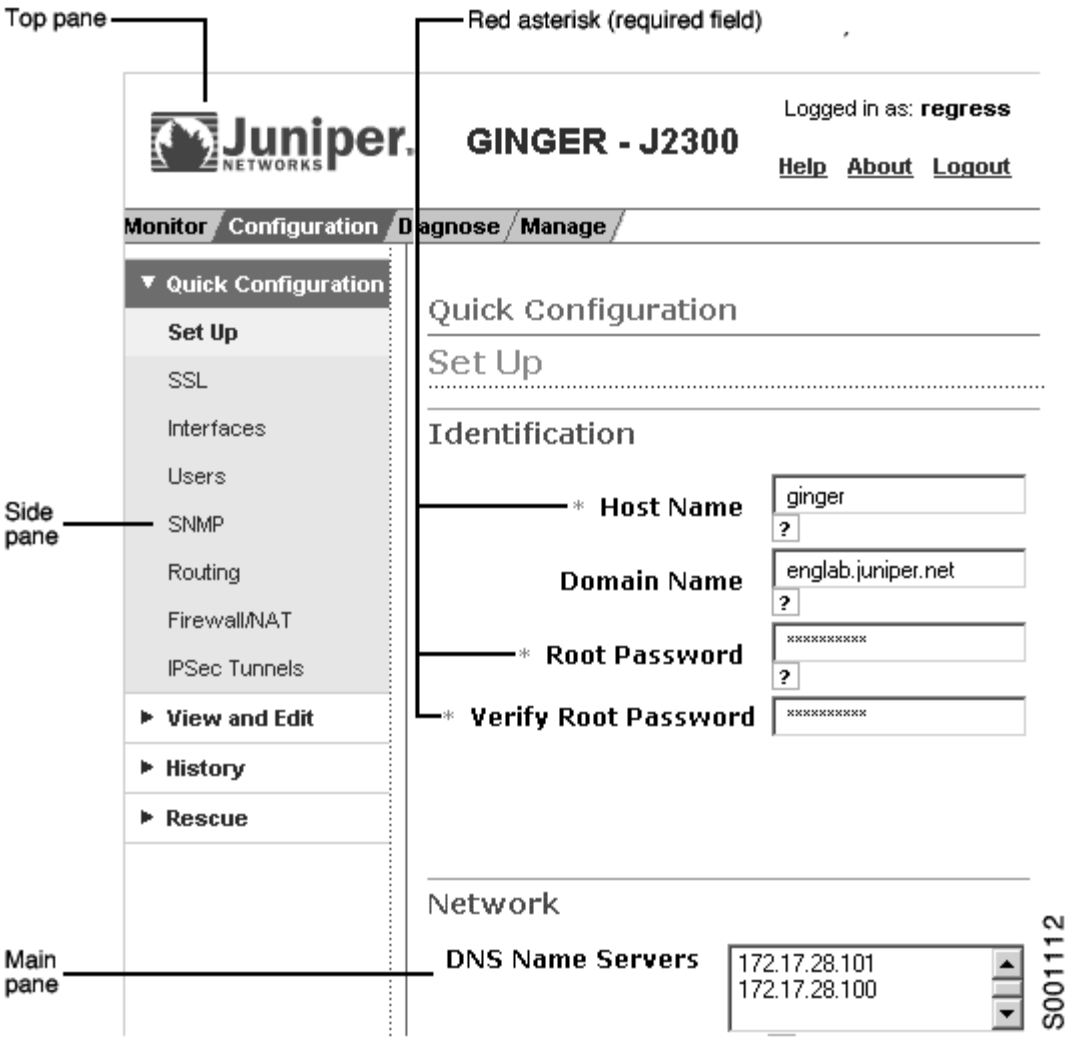

Figure 32: J-Web Layout—Configuration Editor

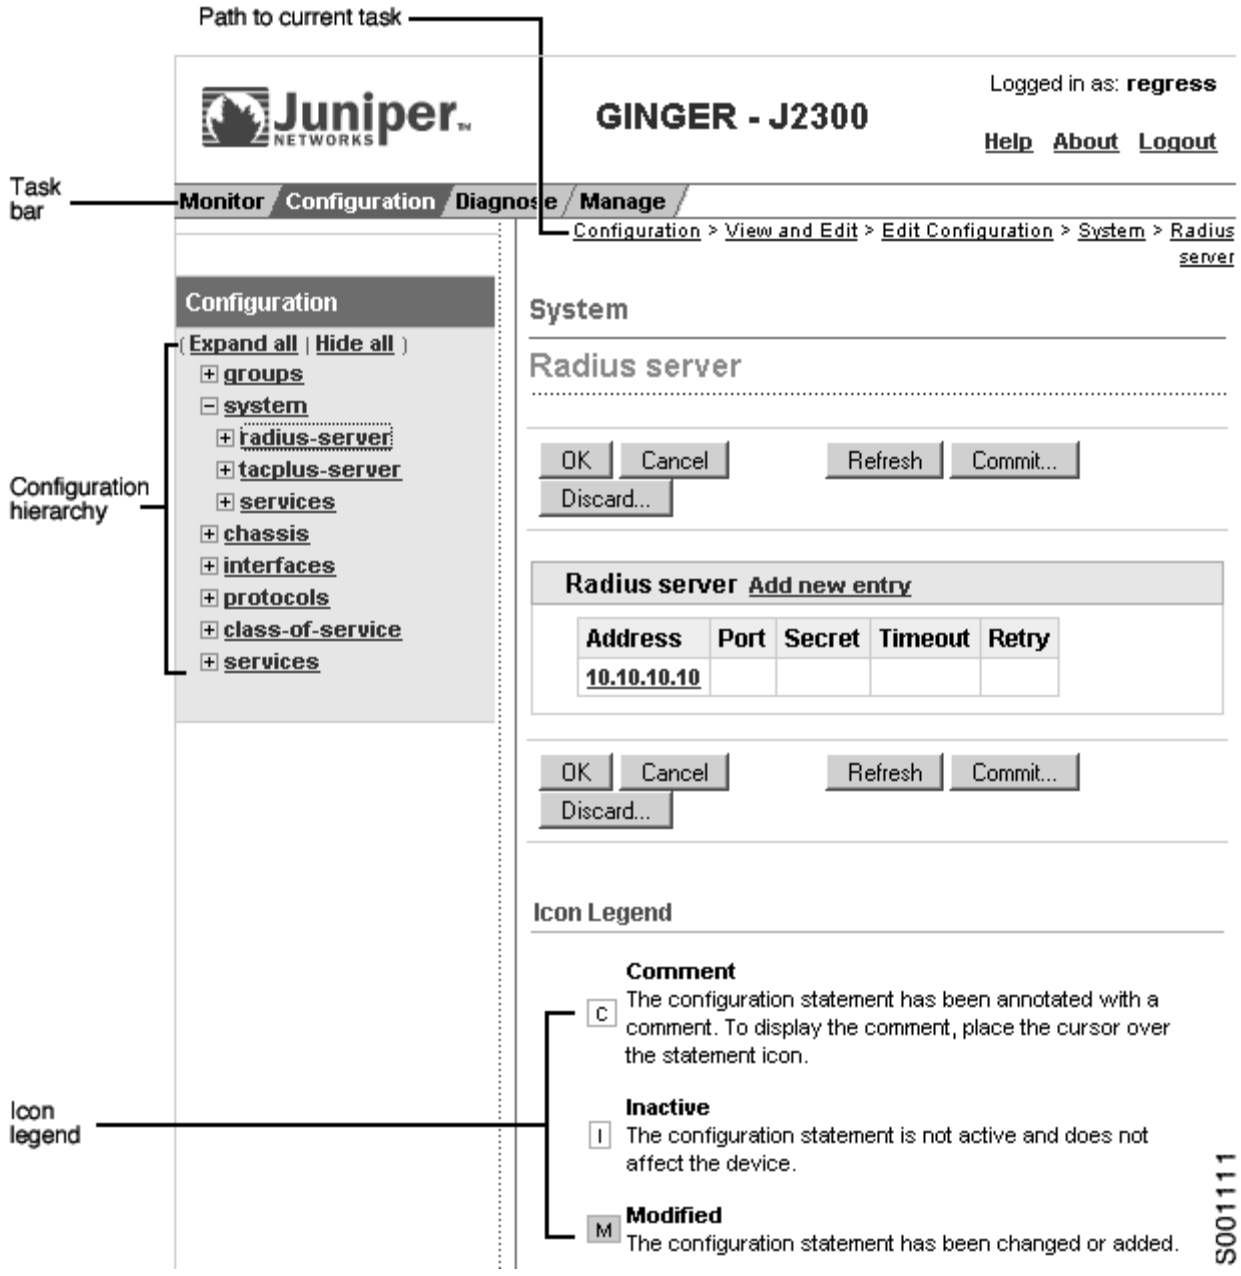

Table 36: Summary of J-Web Elements

J-Web Interface Element	Description
Top Pane	
Juniper Networks logo	Link to www.juniper.net in a new browser window.

J-Web Interface Element	Description
hostname – model	Hostname and model of the Services Router.
Logged in as: *username*	Username you used to log in to the Services Router.
Help	Link to context-sensitive help information.
About	Displays information about the J-Web Interface, such as the version number.
Logout	Ends your current login session with the Services Router and returns you to the login page.
Task bar	Menu of J-Web main options. Click to access. ■ **Monitor**—View information about configuration and hardware on the Services Router. ■ **Configuration**—Configure the Services Router with Quick Configuration or the configuration editor, and view configuration history. ■ **Diagnose**—Troubleshoot network connectivity problems. ■ **Manage**—Manage files and licenses, upgrade software, and reboot the Services Router.
Main Pane	
Help (?) icon	Displays useful information—such as the definition, format, and valid range of an option—when you move the cursor over the question mark.
Red asterisk (*)	Indicates a required field.
Path to current task	Path of main options and suboptions you selected to display the current main and side panes.
Icon Legend	(Applies to the configuration editor only) Explains icons that appear in the user interface to provide information about configuration statements: ■ C—Comment. Move your cursor over the icon to view a comment about the configuration statement. ■ I—Inactive. The configuration statement does not affect the Services Router. ■ M—Modified. The configuration statement is added or modified. ■ *—Mandatory. The configuration statement must have a value.
Side Pane	
Configuration hierarchy	(Applies to the configuration editor only) Displays the hierarchy of committed statements in the Services Router configuration. ■ Click **Expand all** to display the entire hierarchy. ■ Click **Hide all** to display only the statements at the top level. ■ Click plus signs (+) to expand individual items. ■ Click minus signs (-) to hide individual items.

J-Web Sessions

You establish a J-Web session with the Services Router through an HTTP- or HTTPS-enabled Web browser. The HTTPS protocol, which uses 128-bit encryption, is available only in domestic versions of the JUNOS software. To use HTTPS, you must have installed the certificate provided by the Services Router.

When you attempt to log in through the J-Web interface, the Services Router authenticates your username with the same methods used for telnet and SSH.

The Services Router supports only one J-Web session for a single username. Although you might be able to launch multiple Web browsers for multiple views of the same J-Web session, the session can have unpredictable results.

If the Services Router does not detect any activity through the J-Web interface for 15 minutes, the session times out and is terminated. You must log in again to begin a new session.

To explicitly terminate a J-Web session at any time, click **Logout** in the top pane.

Using the Command-Line Interface

This section contains the following topics:

For more information about the CLI, see the *JUNOS System Basics Configuration Guide*.

CLI Command Hierarchy

The CLI commands are organized hierarchically, with commands that perform a similar function grouped together under the same level. For example, all commands that display information about the Services Router system and system software are grouped under the show command, and all commands that display information about the routing table are grouped under the show route command. Figure 33 illustrates a portion of the show command hierarchy.

Figure 33: CLI Command Hierarchy Example

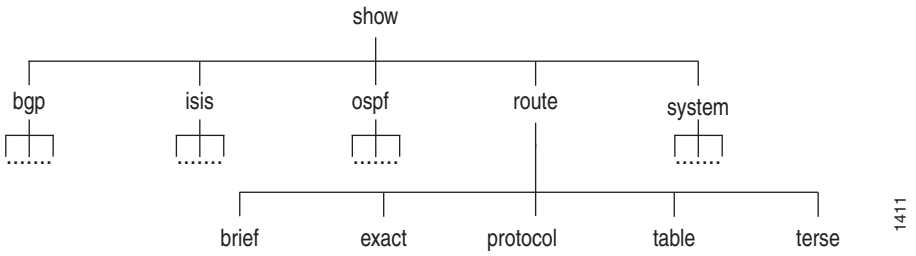

To execute a command, you enter the full command name, starting at the top level of the hierarchy. For example, to display a brief view of the routes in the routing table, use the command show route brief.

The hierarchical organization results in commands that have a regular syntax and provides the following features that simplify CLI use:

- Consistent command names—Commands that provide the same type of function have the same name, regardless of the portion of the software they are operating on. For example, all show commands display software information and statistics, and all clear commands erase various types of system information.
- Lists and short descriptions of available commands—Information about available commands is provided at each level of the CLI command hierarchy. If you type a question mark (?) at any level, you see a list of the available commands along with a short description of each command.
- Command completion—Command completion for command names (keywords) and command options is also available at each level of the hierarchy. If you type a partial command name followed immediately by a question mark (with no intervening space), you see a list of commands that match the partial name you typed.

Starting the CLI

To start the CLI:

1. Establish a connection with the Services Router:
   - To access the router remotely from the network, enter the command you typically use to establish a remote connection (such as telnet or ssh) using the router hostname.
   - To access the router through a management device attached to the console port, start the terminal application.
2. Log in using your username and password.

   After you log in, you enter a UNIX shell.

3. Start the CLI.

```
user# cli
user@host>
```

The presence of the angle bracket (>) prompt indicates the CLI has started. By default, the prompt is preceded by a string that contains your username and the hostname of the Services Router.

To exit the CLI and return to the UNIX shell, enter the quit command.

CLI Operational Mode

The CLI has two modes: *operational* and *configuration*. When you log in to the Services Router and the CLI starts, you are at the top level of operational mode.

To view a list of top-level operational mode commands, type a question mark (?) at the command-line prompt.

```
    user@host> ?

Possible completions:
  clear               Clear information in the system
  configure           Manipulate software configuration information
  file                Perform file operations
  help                Provide help information
  monitor             Show real-time debugging information
  mtrace              Trace multicast path from source to receiver
  ping                Ping remote target
  quit                Exit the management session
  request             Make system-level requests
  restart             Restart software process
  set                 Set CLI properties, date/time, craft interface message
  show                Show system information
  ssh                 Start secure shell on another host
  start               Start shell
  telnet              Telnet to another host
  test                Perform diagnostic debugging
  traceroute          Trace route to remote host
```

At the top level of operational mode are a number of broad groups of CLI commands that are used to perform the following functions:

- Control the CLI environment.
- Monitor and troubleshoot the router.
- Connect to other systems.
- Manage files and software images.
- Control software processes.
- Stop and reboot the router.
- Enter configuration mode.

To control the CLI environment, see "Configuring the CLI Environment" on page 124. To enter configuration mode, see "CLI Configuration Mode" on page 120. For information about the other CLI operational mode functions, see "Monitoring and Diagnosing a Services Router" on page 197 and "Managing Users and Operations" on page 163.

CLI Configuration Mode

To configure the Services Router, including system parameters, routing protocols, interfaces, network management, and user access, you must enter configuration mode. In configuration mode, the CLI provides commands to configure the router, load a text (ASCII) file that contains the router configuration, activate a configuration, and save the configuration to a text file.

You enter configuration mode by entering the configure operational mode command. The CLI prompt changes from user@host> to user@host#.

To view a list of configuration mode commands, type a question mark (?) at the command-line prompt. (You do not need to press Enter after typing the question mark.)

```
    user@host# ?

Possible completions:
  Enter                Execute this command
  activate             Remove the inactive tag from a statement
  annotate             Annotate the statement with a comment
  commit               Commit current set of changes
  copy                 Copy a statement
  deactivate           Add the inactive tag to a statement
  delete               Delete a data element
  edit                 Edit a sub-element
  exit                 Exit from this level
  help                 Provide help information
  insert               Insert a new ordered data element
  load                 Load configuration from ASCII file
  quit                 Quit from this level
```

```
rename              Rename a statement
rollback            Roll back to previous committed configuration
run                 Run an operational-mode command
save                Save configuration to ASCII file
set                 Set a parameter
show                Show a parameter
status              Show users currently editing configuration
top                 Exit to top level of configuration
up                  Exit one level of configuration
wildcard            Wildcard operations
```

The JUNOS software configuration consists of a hierarchy of *statements*. There are two types of statements: *container statements*, which contain other statements, and *leaf statements*, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.

Each statement consists of a fixed keyword and, optionally, an identifier that you define, such as the name of an interface or a username.

To configure the Services Router or to modify an existing configuration, you add statements to the configuration with the edit and set configuration mode commands. For more information about the CLI configuration editor and configuration mode, see "Using the CLI Configuration Editor" on page 146 and the JUNOS software configuration guides.

CLI Basics

This section contains the following topics:

Editing Keystrokes

In the CLI, you use keystrokes to move around on and edit the command line, and to scroll through a list of recently executed commands. Table 37 lists some typical CLI editing tasks and the keystrokes that perform them.

Table 37: CLI Editing Keystrokes

Task Category	Action	Keyboard Sequence
Move the cursor.	Move the cursor back one character.	Ctrl-b
	Move the cursor back one word.	Esc b
	Move the cursor forward one character.	Ctrl-f
	Move the cursor forward one word.	Esc f
	Move the cursor to the end of the command line.	Ctrl-e
Delete characters.	Delete the character before the cursor.	Ctrl-h, Delete, or Backspace
	Delete the character at the cursor.	Ctrl-d
	Delete all characters from the cursor to the end of the command line.	Ctrl-k
	Delete all characters on the command line.	Ctrl-u or Ctrl-x
	Delete the word before the cursor.	Ctrl-w or Esc Backspace
	Delete the word after the cursor.	Esc d
Insert recently deleted text.	Insert the most recently deleted text at the cursor.	Ctrl-y
Redraw the screen.	Redraw the current line.	Ctrl-l
Display previous command lines.	Scroll backward through the list of recently executed commands.	Ctrl-p
	Scroll forward through the list of recently executed commands.	Ctrl-n
	Search the CLI history in reverse order for lines matching the search string.	Ctrl-r
	Search the CLI history by typing some text at the prompt, followed by the keyboard sequence. The CLI attempts to expand the text into the most recent word in the history for which the text is a prefix.	Esc /
Repeat keyboard sequences.	Specify the number of times to execute a keyboard sequence. Replace *number* with a number from 1 through 9, and replace *sequence* with a keyboard sequence in this table.	Esc *number sequence*

Command Completion

You do not always have to remember or type the full command or option name for the CLI to recognize it. To display all possible command or option completions, type the partial command followed immediately by a question mark (?).

To complete a command or option that you have partially typed, press Tab or Spacebar. If the partially typed letters uniquely identify a command, the complete command name appears. Otherwise, a message indicates that your entry is ambiguous or invalid. Possible command completions are displayed if your entry is ambiguous.

You can also use command completion on filenames and usernames. To display all possible values, type one or more characters followed immediately by a question mark. To complete these partial entries, press Tab only. Pressing Spacebar does not work.

Online Help

The CLI provides context-sensitive help at every level of the command hierarchy. The help information tells you which commands are available at the current level in the hierarchy and provides a brief description of each.

To get help while in the CLI, type a question mark (?) in one of the following ways:

- Type a question mark at the command-line prompt—The CLI lists the available commands and options. For examples, see "CLI Operational Mode" on page 119 and "CLI Configuration Mode" on page 120.

- Type a question mark after entering the complete name of a command or command option—The CLI lists the available commands and options, then redisplays the command names and options that you typed:

```
user@host> request ?

Possible completions:
  chassis              Perform chassis-specific operations
  ipsec                Perform IP Security operations
  message              Send text message to other users
  routing-engine       Log in to Routing Engine
  security             Perform security-level operations
  services             Perform service application operations
  support              Perform JUNOS support tasks
  system               Perform system-level operations
user@host> request
```

- Type a question mark in the middle of a command name—The CLI lists possible command completions that match the letters you have entered so far, then redisplays the letters that you typed. For example, to list all operational mode commands that start with the letter **s**, type the following:

```
user@host> s?

Possible completions:
  set                  Set CLI properties, date/time, craft interface message
  show                 Show system information
  ssh                  Start secure shell on another host
  start                Start shell
user@host> s
```

The CLI also provides usage guidelines and summary information for text contained in configuration statements if you enter the help topic and help reference commands. For example, to display usage guidelines for the OSPF hello interval, enter the command help topic ospf hello-interval. You can enter help commands in operational or configuration mode.

Configuring the CLI Environment

You can configure the CLI environment for your current login session. Your settings are not retained when you exit the CLI.

To display the current CLI settings, enter the show cli command:

```
    user@host> show cli

CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 49
CLI screen-width set to 132
CLI terminal is 'vt100'
CLI is operating in enhanced mode
CLI working directory is '/cf/var/home/remote'
```

To change the CLI environment, use the set cli operational mode command:

```
    user@host> set cli ?

Possible completions:
  complete-on-space    Set whether typing space completes current word
  directory            Set working directory
  idle-timeout         Set maximum idle time before login session ends
  prompt               Set CLI command prompt string
  restart-on-upgrade   Set whether CLI prompts to restart after software upgrade
  screen-length        Set number of lines on screen
  screen-width         Set number of characters on a line
  terminal             Set terminal type
```

Table 38 shows how you can change the CLI environment features.

Table 38: Configuring the CLI Environment

Environment Feature	CLI Command	Default Setting	Options
Command completion	set cli complete-on-space (on | off)	on—Pressing Tab or Spacebar completes a command.	■ Set off to allow only Tab for command completion. ■ Set on to re-enable Tab and Spacebar for command completion.
Your working directory	set cli directory *path* 8	/cf/var/home/remote	Replace *path* with the directory you want to enter when you log in to the Services Router.

Environment Feature	CLI Command	Default Setting	Options
Minutes of idle time	set cli idle-time *minutes*	Your session never times out unless your login class specifies a timeout.	■ To enable the timeout feature, replace *timeout* with a value between 1 and 100,000. ■ To disable the timeout feature, replace *timeout* with 0.
Your session prompt	set cli prompt *string*	*user@host* >	Replace *string* with the prompt you want. If the prompt contains spaces or special characters, enclose *string* in quotation marks (“ “).
Restart after upgrade prompt	set cli restart-on-upgrade (on | off)	CLI prompts you to restart the Services Router after a software upgrade.	■ Set off to disable the prompt for the session. ■ Set on to re-enable the prompt.
Number of CLI output line displayed at once	set cli screen-length *length*	Variable (depends on terminal type).	■ To change the number of lines displayed on the screen, replace *length* with a value between 1 and 100,000. ■ To disable the display of a set number of lines, replace *length* with 0. (This feature can be useful when you are issuing CLI commands from scripts.)
Number of CLI characters displayed on a line	set cli screen-width *width*	Variable (depends on terminal type).	To change the number of characters displayed on a line, replace *width* with a value between 0 and 100,000.
Your terminal type.	set cli terminal *terminal-type*	unknown, or set by console.	Replace *terminal-type* with one of the following values: ■ ansi ■ vt100 ■ small-xterm ■ xterm

Chapter 8

Using J-series Configuration Tools

Use J-series configuration tools to configure all services on a J-series Services Router, including system parameters, routing protocols, interfaces, network management, and user access.

This chapter contains the following topics:

Configuration Tools Terms

Before using the J-series configuration tools, become familiar with the terms defined in Table 39.

Table 39: Configuration Tools Terms

Term	Definition
candidate configuration	A working copy of the configuration that can be edited without affecting the Services Router until it is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.
commit a configuration	Have the candidate configuration checked for proper syntax, activated, and marked as the current configuration file running on the Services Router.

Term	Definition
configuration hierarchy	The JUNOS software configuration consists of a hierarchy of *statements*. There are two types of statements: *container statements*, which contain other statements, and *leaf statements*, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.
rescue configuration	Configuration that recovers a Services Router from a configuration that denies management access. You set a current committed configuration through the J-Web interface or CLI for emergency use. To load and commit the rescue configuration, you press and release the CONFIG button.
roll back a configuration	Return to a previously committed configuration.

Configuration Tools Overview

The J-Web interface provides a Quick Configuration tool for basic configuration and a configuration editor for complete configuration. You can also use the JUNOS CLI configuration mode as a configuration editor to create and modify a complete configuration hierarchy. For a comparison of configuration interfaces, see Table 35.

This section contains the following topics:

Editing and Committing a Configuration

When you edit a configuration, you work in a copy of the current configuration to create a candidate configuration. The changes you make to the candidate configuration are visible through the user interface immediately, but do not take effect on the Services Router until you *commit* the changes. When you commit the configuration, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration when you commit the candidate configuration, all changes made by all the users take effect.

If you are editing the configuration with the CLI, you can edit an *exclusive* or *private* candidate configuration. For more information, see "Entering and Exiting Configuration Mode" on page 146.

When you commit a configuration, the Services Router saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0 (the current operational version), and the oldest saved configuration is version 49. You can roll back the configuration to any saved version. Version 0 is stored in the file juniper.conf, and the last three committed configurations are stored in the files juniper.conf.1.gz, juniper.conf.2.gz, and juniper.conf.3.gz. These four files are located in the /config directory, and the

remaining 46 previous versions of committed configurations—files juniper.conf.4.gz through juniper.conf.49.gz—are stored in the /var/db/config directory.

J-Web Configuration Options

You access the J-Web interface configuration tools by selecting **Configuration** in the task bar. Table 40 describes the J-Web configuration options.

Table 40: J-Web Configuration Options

Option	Purpose	Description
Quick Configuration	Basic configuration	Displays options for quick Services Router configuration—Set Up, SSL, Interfaces, Users, SNMP, Routing, Firewall/NAT, and IPSec Tunnels. You can access these options in both the side and main panes. For more information, see "Using J-Web Quick Configuration" on page 131.
View and Edit	Complete configuration	Displays the configuration editor options—View Configuration, Edit Configuration, Edit Configuration Text, and Upload Configuration File. For more information, see "Using the J-Web Configuration Editor" on page 132.
History	File management	Displays the Services Router configuration history and a list of users currently editing the configuration. You can compare, roll back, or download specific versions of the configuration. For more information, see "Managing Configuration Files with the J-Web Interface" on page 139.
Rescue	Configuration recovery	Displays options for setting the current configuration as the rescue configuration, and for viewing and deleting the rescue configuration. For more information, see "Setting a Rescue Configuration" on page 145.

CLI Configuration Commands

The CLI configuration commands allow you to perform the same configuration tasks you can perform using the J-Web interface. Instead of invoking the tools through a graphical interface, you enter configuration mode commands to perform the tasks.

Table 41 provides a summary of the top-level CLI configuration commands.

Table 41: Top-Level CLI Configuration Commands

Command	Function
Managing the Configuration and Configuration Files	
commit	Commit the set of configuration changes in the candidate configuration to take operational effect.

Command	Function
load	Load a configuration from an ASCII configuration file or from terminal input.
rollback	Return to a previously committed configuration.
save	Save the configuration to an ASCII file.
Modifying the Configuration and Its Statements	
activate	Activate a previously deactivated statement or identifier.
annotate	Add a comment to a statement.
copy	Copy and add a statement to the configuration.
deactivate	Deactivate a statement or identifier.
delete	Delete a statement or identifier from the configuration.
insert	Insert an identifier into an existing hierarchy.
rename	Rename an existing statement or identifier.
set	Create a statement hierarchy and set identifier values.
Navigating the Configuration Hierarchy	
edit	Move inside the specified statement hierarchy.
exit	Exit the current level of the statement hierarchy (same function as quit).
quit	Exit the current level of the statement hierarchy (same function as exit).
top	Return to the top level of configuration mode.
up	Move up one level in the statement hierarchy.
Miscellaneous	
help	Provide help about statements.
run	Issue an operational mode command without leaving configuration mode.
show	Display the current configuration.
status	Display the users currently editing the configuration.

For more information about CLI configuration mode commands, see the JUNOS software configuration guides.

Filtering Configuration Command Output

Certain configuration commands, such as show commands, display output. You can filter or redirect the output to a file by including a vertical bar (|), called a *pipe*, when you enter the command. For more information, see "Monitoring and Diagnosing a Services Router" on page 197.

Before You Begin

To use the J-Web interface and CLI configuration tools, you must have the appropriate access privileges. For more information about configuring access privilege levels, see "Adding New Users" on page 175 and the *JUNOS System Basics Configuration Guide*.

Using J-Web Quick Configuration

Use J-Web Quick Configuration to quickly and easily configure the Services Router for basic operation. To access Quick Configuration, select **Configuration > Quick Configuration**. You can select a Quick Configuration option from either the side pane or the main pane (see Figure 34). To configure the Services Router using Quick Configuration, see the configuration sections in this manual.

Figure 34: J-Web Quick Configuration Options

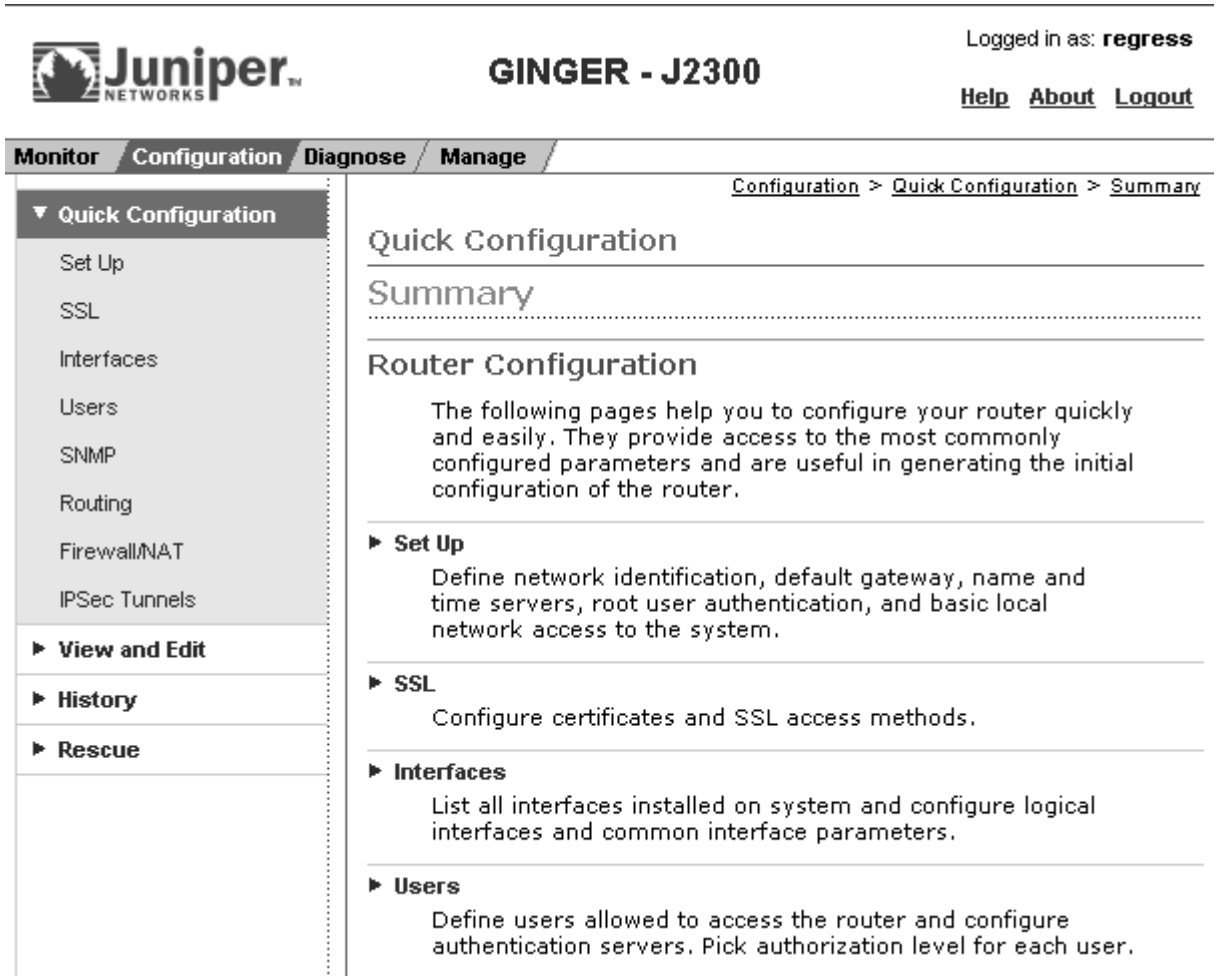

Table 42 describes the functions of the buttons that appear in the J-Web Quick Configuration pages.

Table 42: J-Web Quick Configuration Buttons

Button	Function
Add	Adds statements or identifiers to the configuration.
Delete	Deletes statements or identifiers from the configuration.
OK	Commits your entries into the configuration, and returns you one level up in the configuration hierarchy.
Cancel	Clears the entries you have not yet applied to the configuration, and returns you one level up in the configuration hierarchy.
Apply	Commits your entries into the configuration, and stays at the same level in the configuration hierarchy.

Using the J-Web Configuration Editor

You can use the J-Web configuration editor to perform the following tasks:

- Editing and Committing the Clickable Configuration on page 132
- Viewing the Configuration Text on page 136
- Editing and Committing the Configuration Text on page 137
- Uploading a Configuration File on page 138

Editing and Committing the Clickable Configuration

Use the J-Web configuration editor's clickable interface to perform the following configuration tasks on a Services Router:

- Editing the Clickable Configuration on page 132
- Discarding Parts of a Candidate Configuration on page 135
- Committing a Clickable Configuration on page 136

Editing the Clickable Configuration

To edit the configuration on a series of pages of clickable options that steps you through the hierarchy, select **Configuration > View and Edit > Edit Configuration**. The side pane displays the top level of the configured hierarchy, and the main pane displays configured hierarchy options and the Icon Legend (see Figure 35).

Figure 35: Edit Configuration Page (Clickable)

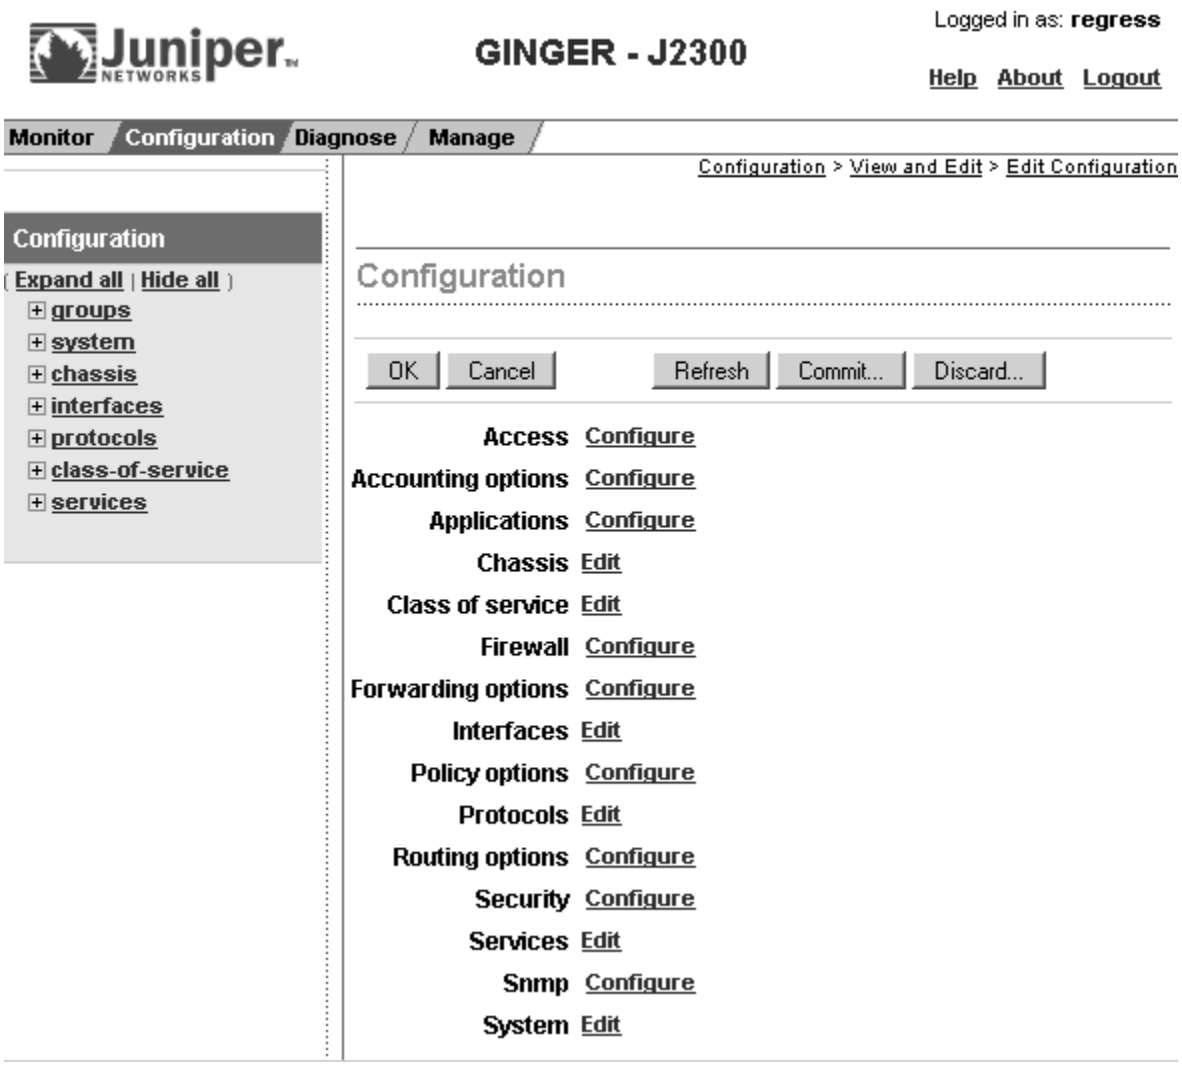

To expand or hide the hierarchy of all the statements in the side pane, click **Expand all** or **Hide all**. To expand or hide an individual statement in the hierarchy, click the expand (+) or collapse (–) icon to the left of the statement.

NOTE: Only those statements included in the committed configuration are displayed in the hierarchy.

The configuration information in the main pane consists of configuration options that correspond to configuration statements. Configuration options that contain subordinate statements are identified by the term *Nested configuration* .

To include or edit statements in the candidate configuration, click one of the links described in Table 43 in the main pane. Then specify configuration information by typing into a field, selecting a value from a list, or clicking a check box (toggle).

Table 43: J-Web Edit Clickable Configuration Links

Link	Function
Edit	Displays information for a configuration option that has already been configured, allowing you to edit a statement.
Configure	Displays information for a configuration option that has not been configured, allowing you to include a statement.
Add new entry	Displays fields and drop-down menus for a statement identifier, allowing you to add a new identifier to a statement.
identifier	Displays fields and drop-down menus for an existing statement identifier, allowing you to edit the identifier.

As you navigate through the configuration, the hierarchy level is displayed at the top of the main pane. You can click a statement or identifier in the hierarchy to display the corresponding configuration options in the main pane.

The main pane includes icons that display information about statements and identifiers when you place your cursor over them. Table 44 describes the meaning of these icons.

Table 44: J-Web Edit Clickable Configuration Icons

Icon	Meaning
C	Displays a comment about a statement.
I	Indicates that a statement is inactive.
M	Indicates that a statement has been added or modified, but has not been committed.
*	Indicates that the statement or identifier is required in the configuration.
?	Provides help information.

NOTE: You can annotate statements with comments or make them inactive only through the CLI. For more information, see "Deactivating a Statement or Identifier" on page 153 and the *JUNOS System Basics Configuration Guide*.

After typing or selecting your configuration edits, click a button in the main pane (described in Table 45) to apply your changes or cancel them, refresh the display, or discard parts of the candidate configuration. An updated configuration does not take effect until you commit it.

Table 45: J-Web Edit Clickable Configuration Buttons

Button	Function
OK	Applies edits to the candidate configuration, and returns you one level up in the configuration hierarchy.
Cancel	Clears the entries you have not yet applied to the candidate configuration, and returns you one level up in the configuration hierarchy.
Refresh	Updates the display with any changes to the configuration made by other users.
Commit	Verifies edits and applies them to the current configuration file running on the Services Router. For details, see "Committing a Clickable Configuration" on page 136.
Discard	Removes edits applied to, or deletes existing statements or identifiers from, the candidate configuration. For details, see "Discarding Parts of a Candidate Configuration" on page 135.

Discarding Parts of a Candidate Configuration

Before committing a candidate configuration, you can discard changes you applied or delete existing statements or identifiers.

To discard parts of a candidate configuration:

1. Navigate to the level of the hierarchy you want to edit and click **Discard**.

   The main pane displays a list of target statements based on the hierarchy level and the changes you have made.

2. Select a radio button to specify the appropriate discard operation or deletion. (Not all buttons appear in all situations.)

   - **Discard Changes Below This Point**—Discards changes made to the candidate configuration at the displayed hierarchy level and below. All subordinate statements and identifiers contained within a discarded statement are also discarded.

   - **Discard All Changes**—Discards all changes made to the candidate configuration.

   - **Delete Configuration Below This Point**—Deletes all changes and statements in the candidate configuration at the displayed hierarchy level and below. All subordinate statements and identifiers contained within a deleted statement are also deleted.

3. To confirm the discard operation or deletion, click **OK**.

   To cancel a discard operation or deletion, click **Cancel**.

The updated candidate configuration does not take effect on the Services Router until you commit it.

Committing a Clickable Configuration

When you finish making changes to a candidate configuration with the J-Web configuration editor's clickable interface, you must commit the changes to use them in the current operational software running on the Services Router.

If another user is editing an exclusive candidate configuration with the CLI, you cannot commit a configuration until the user has committed the configuration. To display a list of users, see "Displaying Users Editing the Configuration" on page 142. For more information about editing an exclusive candidate configuration, see "Entering and Exiting Configuration Mode" on page 146.

To commit a candidate configuration:

1. In the J-Web configuration editor's clickable interface, click **Commit**.

   The main pane displays a summary of your changes in statement form.

2. To confirm the commit operation, click **OK**.

   To cancel a commit operation, click **Cancel**.

   If multiple users are editing the configuration when you commit the candidate configuration, all changes made by all users take effect.

3. To display all the edits applied to the running configuration, click **Refresh**.

Viewing the Configuration Text

To view the entire configuration in text format, select **Configuration > View and Edit > View Configuration Text**. The main pane displays the configuration in text format (see Figure 36).

Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, using an open brace ({) at the beginning of each hierarchy level and a closing brace (}) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon (;), as does the last statement in the hierarchy.

This indented representation is used when the configuration is displayed or saved as an ASCII file. However, when you load an ASCII configuration file, the format of the file is not so strict. The braces and semicolons are required, but the indention and use of new lines are not required in ASCII configuration files.

Figure 36: View Configuration Text Page

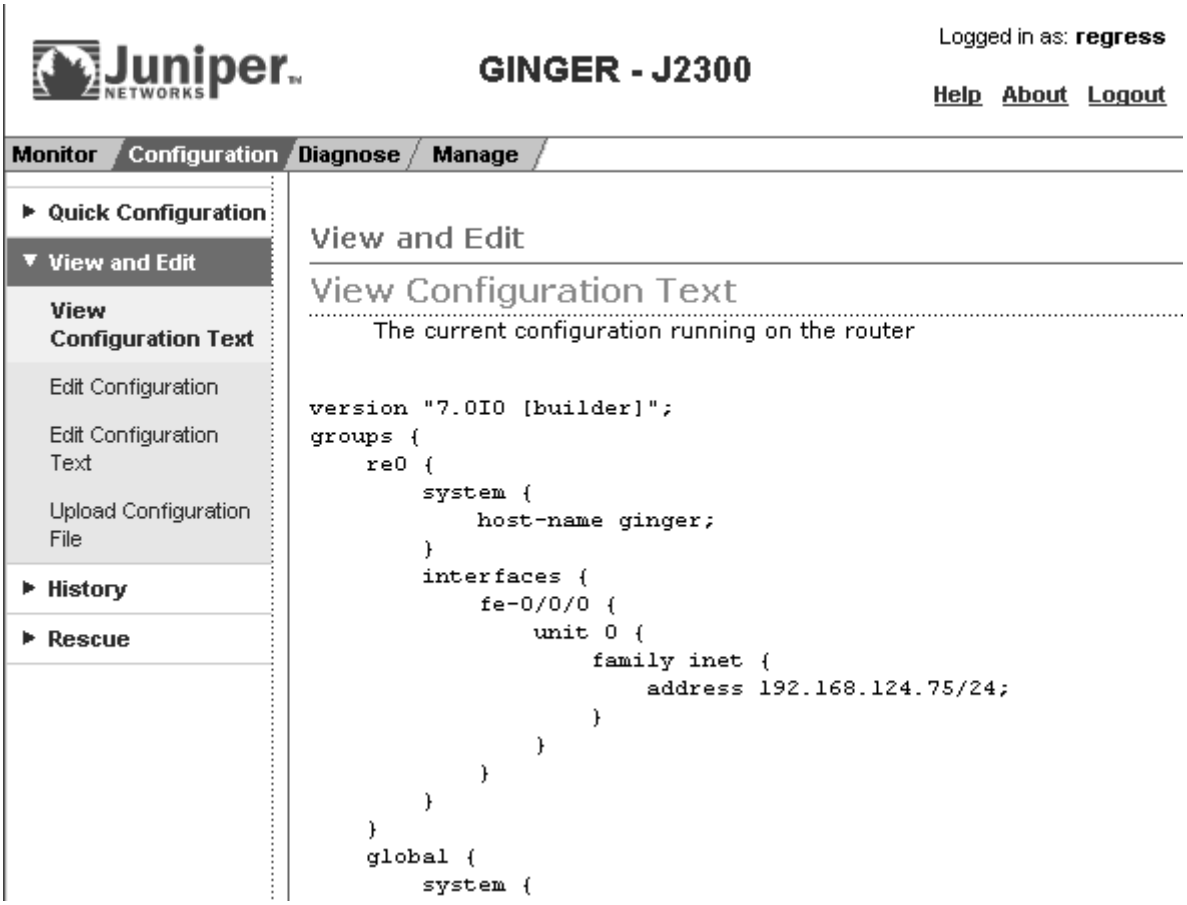

Editing and Committing the Configuration Text

To edit the entire configuration in text format:

CAUTION: We recommend that you use this method to edit and commit the configuration only if you have experience editing configurations through the CLI.

1. Select **Configuration > View and Edit > Edit Configuration Text**. The main pane displays the configuration in a text editor (see Figure 37).

   For more information about the format of an ASCII configuration file, see "Viewing the Configuration Text" on page 136.

2. Navigate to the hierarchy level you want to edit.

You can edit the candidate configuration using standard text editor operations—insert lines (by using the Enter key), delete lines, and modify, copy, and paste text.

3. Click **OK** to load and commit the configuration.

   The Services Router checks the configuration for the correct syntax before committing it.

Figure 37: Edit Configuration Text Page

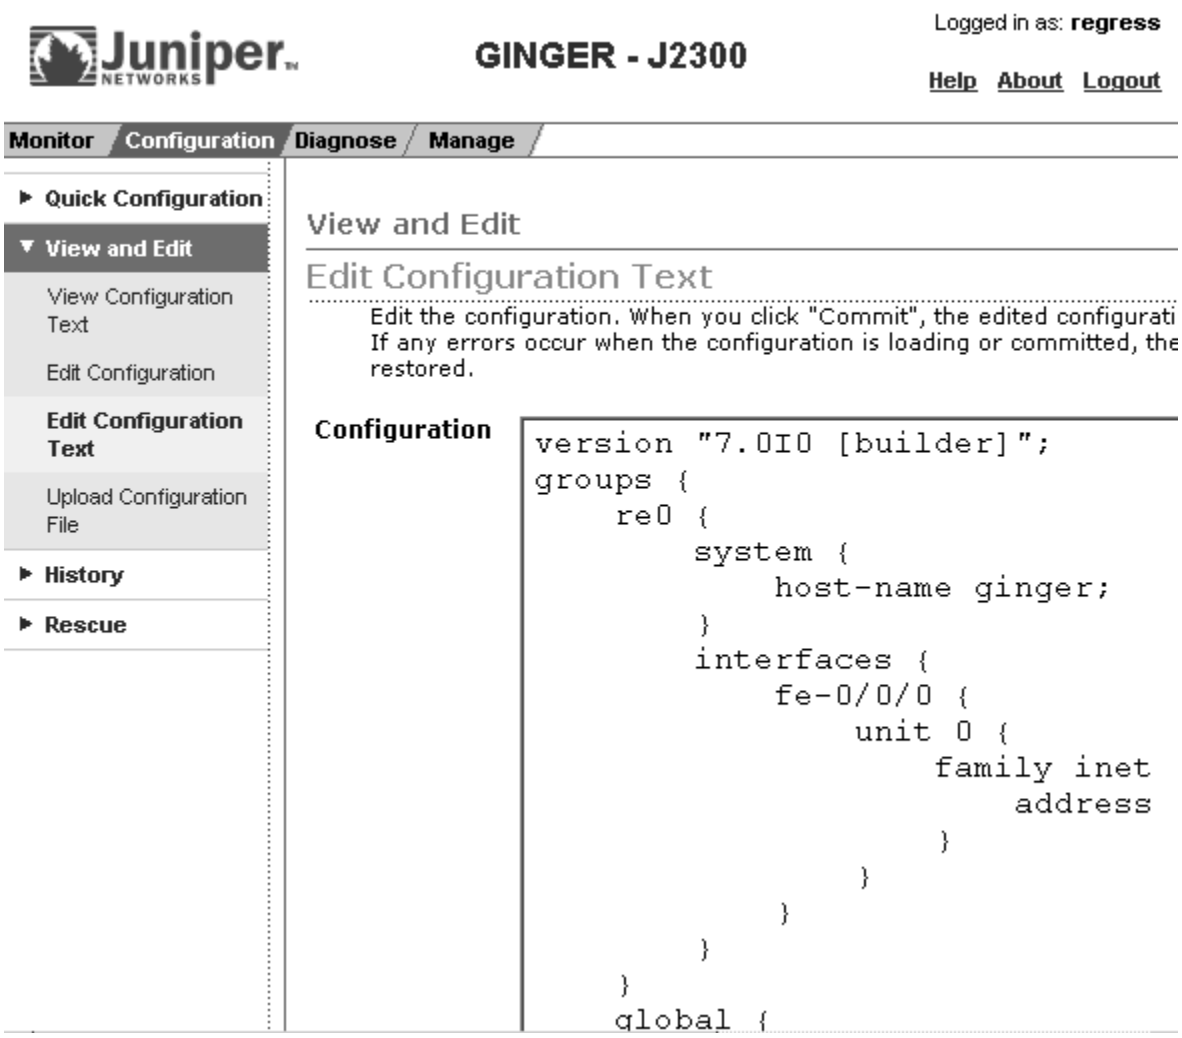

Uploading a Configuration File

To upload a configuration file from your local system:

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

   The main pane displays the File to Upload box (see Figure 38).

2. Specify the name of the file to upload using one of the following methods:

   - Type the absolute path and filename in the File to Upload box.
   - Click **Browse** to navigate to the file.

3. Click **OK** to upload and commit the configuration.

   The Services Router checks the configuration for the correct syntax before committing it.

Figure 38: J-Web Upload Configuration File Page

Managing Configuration Files with the J-Web Interface

The J-Web interface provides configuration database and history information that allows you to manage configuration files. This section contains the following topics:

Configuration Database and History Overview

When you commit a configuration, the Services Router saves the current operational version and the previous 49 versions of committed configurations. To manage these configuration files with the J-Web interface, select **Configuration > History**. The main pane displays Database Information and Configuration History (see Figure 39).

Table 46 and Table 47 summarize the contents of the display.

Figure 39: Configuration Database and History Page

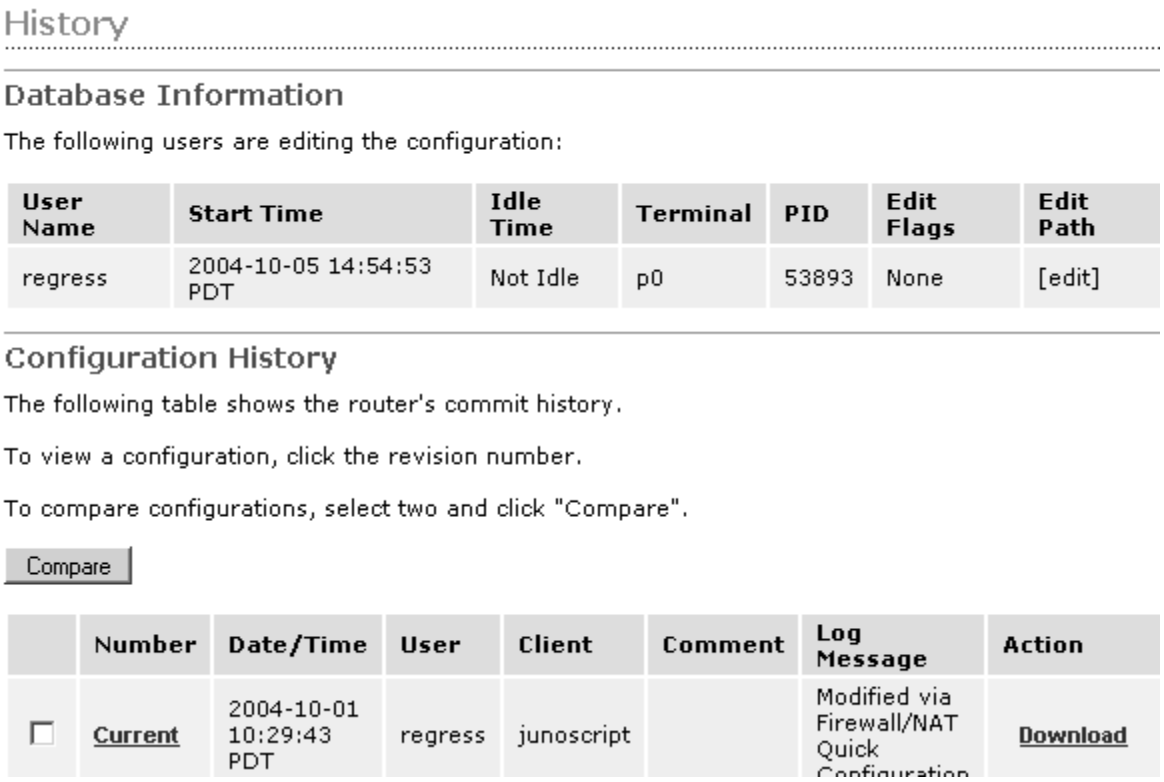

Table 46: J-Web Configuration Database Information Summary

Field	Description
User Name	Name of user editing the configuration.
Start Time	Time of day the user logged in to the Services Router.
Idle Time	Elapsed time since the user issued a configuration command from the CLI.
Terminal	Terminal on which the user is logged in.
PID	Process identifier assigned to the user by the Services Router.
Edit Flags	Designates a private or exclusive edit.
Edit Path	Level of the configuration hierarchy that the user is editing.

Table 47: J-Web Configuration History Summary

Field	Description
Number	Version of the configuration file.
Date/Time	Date and time the configuration was committed.
User	Name of the user who committed the configuration.
Client	Method by which the configuration was committed: ■ cli—A user entered a JUNOS command-line interface command. ■ junoscript—A JUNOScript client performed the operation. Commit operations performed by users through the J-Web interface are identified in this way. ■ snmp—An SNMP set request started the operation. ■ button—The CONFIG button on the router was pressed to commit the rescue configuration (if set) or to clear all configurations except the factory configuration. ■ autoinstall—Autoinstallation was performed. ■ other—Another method was used to commit the configuration.
Comment	Comment.

Field	Description
Log Message	Method used to edit the configuration: ■ Imported via paste—Configuration was edited and loaded with the **Configuration > View and Edit > Edit Configuration Text** option. For more information, see "Editing and Committing the Configuration Text" on page 137. ■ Imported upload [*filename*]—Configuration was uploaded with the **Configuration > View and Edit > Upload Configuration File** option. For more information, see "Uploading a Configuration File" on page 138. ■ Modified via *quick-configuration*—Configuration was modified using the J-Web Quick Configuration tool specified by *quick-configuration*. For more information, see "Using J-Web Quick Configuration" on page 131. ■ Rolled back via *user-interface*—Configuration was rolled back to a previous version through the user interface specified by *user-interface*, which can be Web Interface or CLI. For more information, see "Loading a Previous Configuration File" on page 145.
Action	Action to perform with the configuration file. The action can be **Download** or **Rollback**. For more information, see "Downloading a Configuration File" on page 144 and "Loading a Previous Configuration File" on page 145.

The configuration history display allows you to perform the following operations:

- View a configuration.
- Compare two configurations.
- Download a configuration file to your local system.
- Roll back the configuration to any of the previous versions stored on the Services Router.

For more information about saved versions of configuration files, see "Editing and Committing a Configuration" on page 128.

Displaying Users Editing the Configuration

To display a list of users editing the Services Router configuration, select **Configuration > History**. The list is displayed as Database Information in the main pane (see Figure 39). Table 46 summarizes the Database Information display.

Comparing Configuration Files

To compare any two of the past 50 committed configuration files:

1. Select **Configuration > History**.

   A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 39). Table 47 summarizes the Configuration History display.

2. Click two of the check boxes to the left of the configuration versions you want to compare.

3. Click **Compare**.

   The main pane displays the differences between the two configuration files at each hierarchy level as follows (see Figure 40):

   - Lines that have changed are highlighted side by side in green.
   - Lines that exist only in the more recent configuration file are displayed in red on the left.
   - Lines that exist only in the least recent configuration file are displayed in blue on the right.

Figure 40: J-Web Configuration File Comparison Results

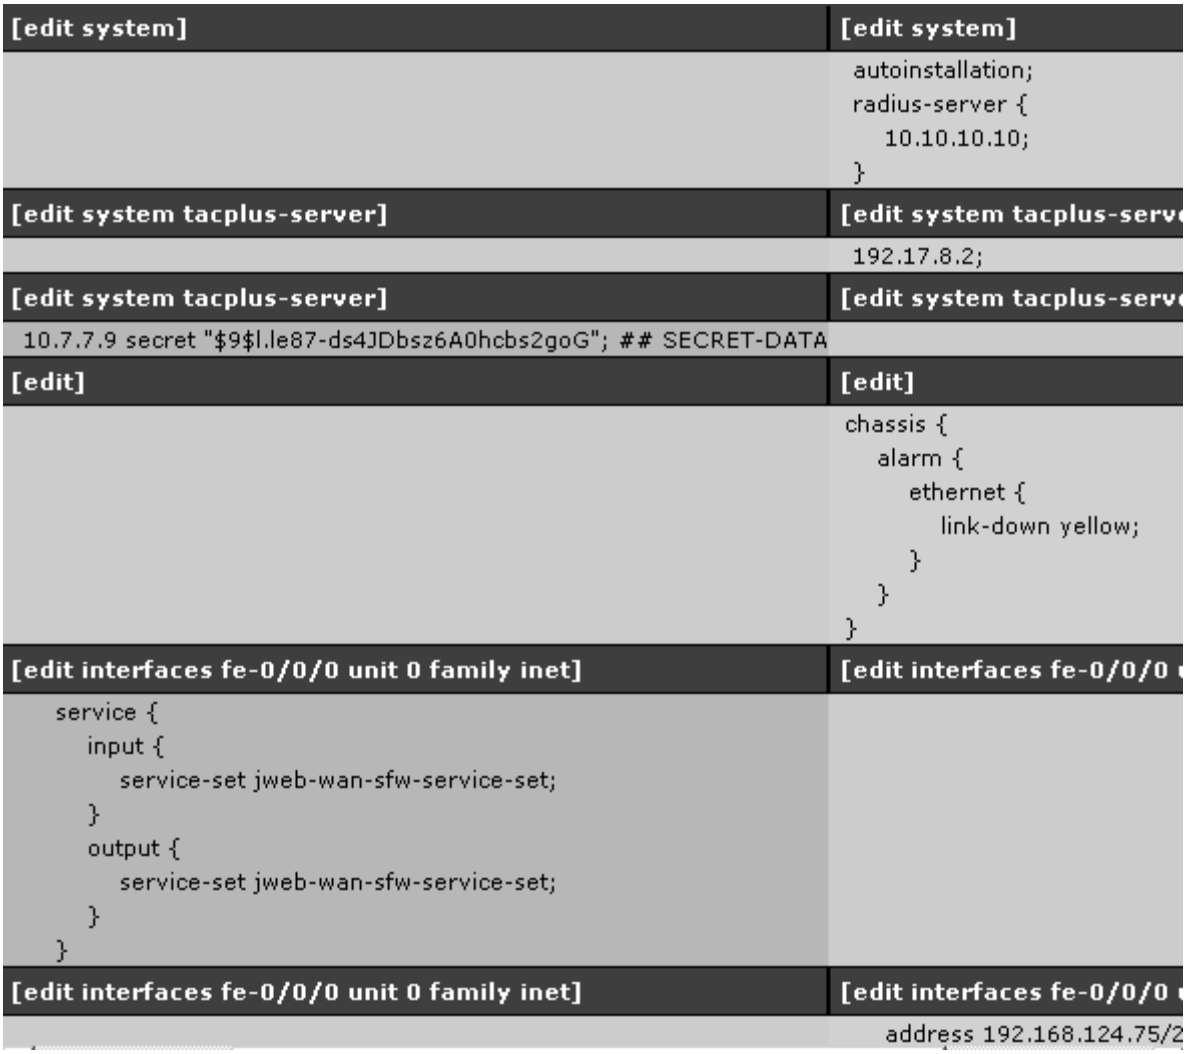

Downloading a Configuration File

To download a configuration file from the Services Router to your local system:

1. Select **Configuration > History**.

   A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 39). Table 47 summarizes the Configuration History display.

2. In the Action column, click **Download** for the version of the configuration you want to download.

3. Select the options your Web browser provides that allow you to save the configuration file to a target directory on your local system.

   The file is saved as an ASCII file.

Loading a Previous Configuration File

To download a configuration file from the Services Router to your local system:

To load (roll back) and commit a previous configuration file stored on the Services Router:

1. Select **Configuration > History**.

   A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 39). Table 47 summarizes the Configuration History display.

2. In the Action column, click **Rollback** for the version of the configuration you want to load.

   The main pane displays the results of the rollback operation.

NOTE: When you click **Rollback**, the Services Router loads and commits the selected configuration. This behavior is different from entering the rollback configuration mode command from the CLI, where the configuration is loaded, but not committed.

Setting a Rescue Configuration

If someone inadvertently commits a configuration that denies management access to the Services Router, you can delete the invalid configuration and replace it with a rescue configuration by pressing the CONFIG button on the router. You must have previously set the rescue configuration through the J-Web interface or the CLI. The rescue configuration is a previously committed, valid configuration.

CAUTION: Pressing and holding the CONFIG button for longer than 15 seconds deletes *all* configurations on the router, including the backup configurations and rescue configuration, and loads and commits the factory configuration.

To view, set, or delete the rescue configuration, select **Configuration > Rescue**. On the Rescue page, you can perform the following tasks:

- View the current rescue configuration—Click **View rescue configuration**.
- Set the current running configuration as the rescue configuration—Click **Set rescue configuration**.
- Delete the current rescue configuration—Click **Delete rescue configuration**.

Using the CLI Configuration Editor

You can use the CLI configuration editor to perform the following tasks:

- Entering and Exiting Configuration Mode on page 146
- Navigating the Configuration Hierarchy on page 148
- Modifying the Configuration on page 149
- Committing a Configuration with the CLI on page 154
- Entering Operational Mode Commands During Configuration on page 157

Entering and Exiting Configuration Mode

You must have access privileges to edit the configuration. For more information, see "Before You Begin" on page 130.

To enter and exit configuration mode:

1. At the CLI prompt, enter the configure operational mode command.

   Select the form of the configure command (see Table 48) that is appropriate for the way you want to edit and commit the candidate configuration. For example:

   ```
   user@host> configure
   user@host#
   ```

2. To display the users currently editing the configuration, enter the status command:

   ```
   user@host# status

   Users currently editing the configuration:
     user1 terminal p1 (pid 66847) on since 2004-04-19 12:32:56 PDT
         [edit]
     user2 terminal p2 (pid 85743) on since 2004-04-19 11:44:06 PDT
         [edit interfaces]
   ```

For each user, the CLI displays the username, terminal, process identifier, login date and time, and hierarchy level being edited. You can specify the terminal and process identifier in the request system logout command.

3. To exit configuration mode and return to operational mode:

   - For the top level, enter the following command:

     user@host# **exit**

   - From any level, enter the following command:

     user@host# **exit configuration-mode**

For more information about the configure command, including restrictions on entering and exiting the various configuration modes, see the *JUNOS System Basics Configuration Guide*.

Table 48: Forms of the configure Command

Command	Edit Access	Commit Access
configure	■ No one can lock the configuration. All users can make configuration changes. ■ When you enter configuration mode, the CLI displays the following information: ▫ A list of the other users editing the configuration. ▫ Hierarchy levels the users are viewing or editing. ▫ Whether the configuration has been changed, but not committed.	■ No one can lock the configuration. All users can commit all changes to the candidate configuration. ■ If you and another user make changes and the other user commits changes, your changes are committed as well.
configure exclusive	■ One user locks the configuration and makes changes without interference from other users. ■ Other users can enter and exit configuration mode, but they cannot change the configuration. ■ If you enter configuration mode while another user has locked the configuration, the CLI displays the user and the hierarchy level the user is viewing or editing. ■ If you enter configuration mode while another user has locked the configuration, you can forcibly log out that user with the request system logout operational mode command. (For details, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.)	
configure private	■ Multiple users can edit the configuration at the same time. ■ Each user has a private candidate configuration to edit independently of other users.	■ When you commit the configuration, the Services Router verifies that the operational (running) configuration has not been modified by another user before accepting your private candidate configuration as the new operational configuration. ■ If the configuration has been modified by another user, you can merge the modifications into your private candidate configuration and attempt to commit again.

Navigating the Configuration Hierarchy

When you first enter configuration mode, you are at the top level of the configuration command hierarchy, which is indicated by the [edit] banner. To move down through an existing configuration command hierarchy, or to create a hierarchy and move down to that level, use the edit command, specifying the hierarchy level at which you want to be:

user@host# **edit** <*statement-path*> <*identifier*>

Replace *statement-path* with the hierarchy level and *identifier* with a string that identifies an instance of a statement. (Not all statements require identifiers.) If the identifier contains a space, you must enclose the identifier in quotation marks (” “).

After you enter an edit command, the banner changes to indicate your current level in the hierarchy:

```
[edit]
user@host# edit protocols ospf

[edit protocols ospf]
user@host#
```

To move back up to the previous hierarchy level, enter the exit command. This command is, in effect, the opposite of the edit command. For example:

```
[edit]
user@host# edit protocols ospf

[edit protocols ospf]
user@host# edit area 0.0.0.0

[edit protocols ospf area 0.0.0.0]
user@host# exit

[edit protocols ospf]
user@host# exit

[edit]
user@host#
```

To move up one level, enter the up command. For example:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0

[edit protocols ospf area 0.0.0.0]
user@host# up

[edit protocols ospf]
user@host# up

[edit protocols]
user@host# up

[edit]
user@host#
```

To move directly to the top level of the hierarchy, enter the top command. For example:

```
[edit protocols ospf area 0.0.0.0]
user@host# top

[edit]
user@host#
```

To display the configuration, enter the show command:

show <*statement-path*>

The configuration at the current hierarchy level, or at the level specified by *statement-path*, is displayed. For example, entering the show command in each of the following cases displays the same level of the configuration:

```
[edit]
user@host# show interfaces fe-0/0/0
unit 0 {
   family inet {
                address 192.168.4.1/30;
   }
}

[edit]
user@host# edit interfaces fe-0/0/0
[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
   family inet {
                address 192.168.4.1/30;
   }
}
```

Modifying the Configuration

You can modify the configuration by performing the following operations:

- Adding or Modifying a Statement or Identifier on page 150
- Deleting a Statement or Identifier on page 150
- Copying a Statement on page 151
- Renaming an Identifier on page 151
- Inserting an Identifier on page 152
- Deactivating a Statement or Identifier on page 153

Adding or Modifying a Statement or Identifier

To add or modify statements in a configuration, use the set command:

set <*statement-path*> *statement* <*identifier*>

Replace *statement-path* with the path to the statement from the current hierarchy level, and *statement* with the statement itself. Replace *identifier* with a string that identifies an instance of a statement. (Not all statements require identifiers.) If the identifier contains a space, you must enclose the identifier in quotation marks (" ").

If the statement or identifier does not exist in the configuration hierarchy, it is added. If the statement or identifier already exists, it is modified (unless multiple occurrences of the same statement or identifier are allowed in the configuration, in which case another instance is added to the configuration). After you enter the set command, you remain at the same level in the hierarchy.

You can enter a single set command from the top level of the hierarchy. Alternatively, you can enter the edit command to move to the target hierarchy level, from which you can enter the set command. In either case, the CLI creates the hierarchy level if it does not exist. For example, to set the OSPF hello interval from the top level of the hierarchy, enter the set command as follows:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface t1-0/0/0 hello-interval 5
```

Alternatively, use the edit command to create and move to the [edit protocols ospf area 0.0.0.0 interface t1-0/0/0] hierarchy level, then enter a set command to set the value of the hello-interval statement:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0 interface t1-0/0/0

[edit protocols ospf area 0.0.0.0 interface t1-0/0/0]
user@host# set hello-interval 5
```

Deleting a Statement or Identifier

To delete a statement or identifier from the configuration, enter the delete command:

delete <*statement-path*> <*identifier*>

When you delete a statement, the statement and all its subordinate statements and identifiers are removed from the configuration and revert to their default values. To delete the entire hierarchy starting at the current level, enter the delete command without specifying a statement or an identifier. You are prompted to confirm the deletion.

As with the set command, you can enter a single delete command from the top level of the hierarchy, or you can use the edit command to move to the target hierarchy level, from which you can enter the delete command.

Copying a Statement

To make a copy of an existing statement in the configuration, use the copy command:

copy *existing-statement* **to** *new-statement*

The existing statement and all its subordinate statements are copied and added to the configuration. After you enter the copy command, the configuration might not be valid. If necessary, modify the existing statement or the new statement to ensure the configuration is valid.

The following example shows how to copy a unit configured at the [edit interfaces fe-0/0/0] hierarchy level:

```
[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
   family inet {
                address 10.14.1.1/24;
   }
}

[edit interfaces fe-0/0/0]
user@host# copy unit 0 to unit 1

[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
   family inet {
                address 10.14.1.1/24;
   }
}
unit 1 {
   family inet {
                address 10.14.1.1/24;
   }
}
```

In this example, after you enter the copy command, unit 0 and unit 1 have the same IP address in the candidate configuration. To modify the IP address of unit 1 before committing the configuration, use the rename command as described in "Renaming an Identifier" on page 151.

Renaming an Identifier

There are two ways to rename an identifier that already exists in a configuration:

- Delete the identifier with the delete command, then add it back into the configuration with the set command.
- Rename the identifier with the rename command:

  rename <*statement-path*> *identifier1* **to** *identifier2*

In the example provided in “Copying a Statement” on page 151, to rename the IP address of unit 1 from 10.14.1.1/24 to 10.14.2.1/24, enter the rename command as follows:

```
user@host# rename interfaces fe-0/0/0 unit 1 family inet address 10.14.1.1/24 to
address 10.14.2.1/24
```

Inserting an Identifier

To insert an identifier into a specific location within the configuration, use the insert command:

insert <*statement-path*> *identifier1* (**before** | **after**) *identifier2*

Generally, you can add most identifiers into the configuration in any order. However, when you are inserting identifiers that must be analyzed in order—such as terms in a routing policy or firewall filter—you must specify before or after. If you do not specify where to insert an identifier with the insert command, the identifier is placed at the end of the list of similar identifiers.

In the following example, the firewall filter terms were added to the configuration in the following order: term1, term3, term2. The insert command is used to insert term2 before term3.

```
[edit]
user@host# show firewall
family inet {
   filter filter1 {
      term term1 {
         from {
            address {
                        192.168.0.0/16;
            }
         }
         then {
                     reject;
         }
      }
      term term3 {
         then {
                     reject;
         }
      }
      term term2 {
         from {
                     destination-port ssh;
         }
                  then accept;
      }
   }
}

[edit]
user@host# insert firewall family inet filter filter1 term term2 before term term3
```

```
[edit]
user@host# show firewall
family inet {
   filter filter1 {
      term term1 {
         from {
            address {
                        192.168.0.0/16;
            }
         }
         then {
                     reject;
         }
      }
      term term2 {
         from {
                     destination-port ssh;
         }
                  then accept;
      }
      term term3 {
         then {
                     reject;
         }
      }
   }
}
```

Deactivating a Statement or Identifier

You can deactivate a statement or identifier so that it does not take effect when you enter the commit command. Any deactivated statements and identifiers are marked with the inactive: tag and remain in the configuration.

To deactivate a statement or identifier, use the deactivate command:

deactivate (*statement* | *identifier*)

To reactivate a statement or identifier, use the reactivate command:

reactivate (*statement* | *identifier*)

Reactivate removes the inactive: tag so that a statement or identifier takes effect when you commit the configuration.

In both commands, *statement* or *identifier* must be at the current hierarchy level.

The following example shows how to deactivate interface fe-0/0/0 at the [edit interfaces] hierarchy level:

```
[edit interfaces]
user@host# deactivate fe-0/0/0
```

```
[edit interfaces]
user@host# show
inactive: fe–0/0/0 {
   unit 0 {
      family inet {
                  address 10.14.1.1/24;
      }
   }
}
```

Committing a Configuration with the CLI

To save candidate configuration changes to the configuration database and activate the configuration on the Services Router, enter the commit command from any hierarchy level:

```
[edit]
user@host# commit
commit complete
```

If more than one user is modifying the configuration, committing it saves and activates the changes made by all the users.

The Services Router checks the configuration for syntax errors. If the syntax is correct, the configuration is activated and becomes the current, operational configuration running on the Services Router. If the configuration contains syntax errors, the router sends a message indicating the location of the error and does not activate the configuration. The error message has the following format:

```
[edit edit-path ]
   offending-statement ;
      error-message
```

You can specify one or more options within the commit command—or use it with the rollback command—to perform the following operations:

- Verifying a Configuration on page 154
- Committing a Configuration and Exiting Configuration Mode on page 155
- Committing a Configuration That Requires Confirmation on page 155
- Scheduling and Canceling a Commit on page 155
- Loading a Previous Configuration File on page 156

Verifying a Configuration

To verify that the syntax of a configuration is correct, enter the commit check command:

```
[edit]
user@host# commit check
```

```
configuration check succeeds
```

If the configuration contains syntax errors, a message indicates the location of the error.

Committing a Configuration and Exiting Configuration Mode

To save candidate configuration changes, activate the configuration on the Services Router, and exit configuration mode, enter the commit and-quit command:

```
[edit]
user@host# commit and-quit
commit complete
exiting configuration mode
user@host>
```

If the configuration contains syntax errors, a message indicates the location of the error.

Committing a Configuration That Requires Confirmation

You can commit the current candidate configuration but require an explicit confirmation for the committed configuration to become permanent. This commit process is useful for verifying that a configuration change works correctly and does not prevent management access to the Services Router. If the change prevents access or causes other errors, an automatic rollback to the previous configuration restores access after the rollback confirmation timeout expires.

To commit the current candidate configuration, but require an explicit confirmation for the commit to become permanent, use the commit confirmed command:

commit confirmed <*minutes*>

Replace *minutes* with the number of minutes to allow for the timeout period. The default value is 10 minutes.

To make the new configuration permanent, enter the commit or commit check command within the timeout period specified in the commit confirmed command. If the commit is not confirmed within the timeout period, the Services Router automatically rolls back to the previous configuration.

If the configuration contains syntax errors, a message indicates the location of the error.

Scheduling and Canceling a Commit

To schedule a candidate configuration for a commit operation at a future time or the next time the Services Router is rebooted, use the commit at command:

commit at *string*

Replace *string* with reboot or the time at which the configuration is to be committed, in one of the following formats:

- *hh* : *mm* < : ss > —Hours, minutes, and seconds (optional), in 24-hour format. For example, 20:30 is 8:30 PM.
- *yyyy* - *mm* - *dd* *hh* : *mm* < : ss > —Year, month, date, hours, minutes, and seconds (optional), in 24-hour format. For example, 2004-09-05 08:00 is September 5, 2004 at 8:00 AM.

The Services Router checks the configuration. If the result of the check is successful, the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit operation can be performed until the scheduled one is completed. If the configuration contains syntax errors, a message indicates the location of the error.

To cancel a pending commit operation, use the clear system commit operational mode command. For more information, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Loading a Previous Configuration File

To load, or *roll back*, a previous configuration file stored on the Services Router without activating it, use the rollback command:

rollback <*string*>

Replace *string* with a value from 0 through 49, or rescue (if a rescue configuration exists). The default value is 0.

When you commit a configuration, the Services Router saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0 (the current operational version), and the oldest saved configuration is version 49.

If you have defined a rescue configuration (by using the request system configuration rescue save operational mode command), you can roll back to this configuration by entering rollback rescue. (You can also roll back to the rescue configuration or the default factory configuration by pressing the CONFIG button on the Services Router. For more information, see “CONFIG Button and LED” on page 12.)

To set a rescue configuration with the J-Web interface, see “Setting a Rescue Configuration” on page 145.

For more information about saved versions of configuration files, see “Editing and Committing a Configuration” on page 128.

To activate the configuration you loaded, you must commit it:

```
[edit]
user@host# rollback 2
load complete
[edit]
```

user@host# **commit**

To display previous configurations, including the rollback number, date, time, name of the user who committed changes, and commit method, use the rollback ? command:

```
    user@host# rollback ?

Possible completions:
  <[Enter]>           Execute this command
  0                   2004-05-27 14:50:05 PDT by root via junoscript
  1                   2004-05-27 14:00:14 PDT by root via cli
  2                   2004-05-27 13:16:19 PDT by snmpset via snmp
  ...
  28                  2004-05-21 16:56:25 PDT by root via cli
  rescue              2004-05-27 14:30:23 PDT by root via cli
  |                   Pipe through a command
```

The access privilege level for using the rollback command is controlled by the rollback permission bit. Users for whom this permission bit is not set can return only to the most recently committed configuration. Users for whom this bit is set can return to any prior committed configuration. For more information, see the *JUNOS System Basics Configuration Guide*.

Entering Operational Mode Commands During Configuration

While in configuration mode, you might need to enter an operational mode command, such as show or request. To enter a single operational mode command, first enter the run command and then specify the operational mode command as follows:

user@host# **run** *operational-mode-command*

For example, to display a pending system reboot while in configuration mode, enter the show system reboot operational mode command as follows:

```
[edit]
user@host# run show system reboot
No shutdown/reboot scheduled.
```

If you are in operational mode, the show cli history command displays the history of the operational mode commands issued. To display the history of the configuration mode commands issued, enter the show cli history command from configuration mode as follows:

```
[edit]
user@host# run show cli history
15:32:51 – exit
15:52:02 – load merge terminal
17:07:57 – run show ospf statistics
17:09:12 – exit
17:18:49 – run show cli history
```

Managing Configuration Files with the CLI

This section contains the following topics:

Loading a New Configuration File

You can create a configuration file, copy the file to the Services Router, and then load the file into the CLI. After you load the file, you can commit it to activate the configuration on the router, or you can edit the configuration interactively with the CLI and commit it at a later time.

You can also create a configuration while typing at the terminal and then load it. Loading a configuration from the terminal is generally useful when you are cutting existing portions of the configuration and pasting them elsewhere in the configuration.

To load an existing configuration file that is located on the router, use the following version of the load command:

load (**merge** | **override** | **patch** | **replace** | **update**) *filename* <**relative**>

To load a configuration from the terminal, use the following version of the load command:

load (**merge** | **override** | **patch** | **replace** | **update**) **terminal** <**relative**>

Use the load command options provided in Table 49. (The *incoming configuration* is the configuration in *filename* or the one that you type at the terminal). For more information about loading a configuration, see the *JUNOS System Basics Configuration Guide*.

Table 49: Load Configuration File Options

Option	Function
merge	Combines the current configuration and the incoming configuration. A merge operation is useful when you are adding a new section to an existing configuration. If the existing configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the existing configuration.
override	Discards the current candidate configuration and loads the incoming configuration.
patch	Changes part of the configuration with the incoming configuration and marks only those parts as changed.
relative	Allows you to use the merge, replace, and update options without specifying the full hierarchy level.

Option	Function
replace	Replaces portions of the configuration based on the **replace:** tags in the incoming configuration. The Services Router searches for the **replace:** tags, deletes the existing statements of the same name (if any), and replaces them with the incoming configuration. If no statement of the same name exists in the configuration, the replace operation adds it to the configuration. If you are performing a replace operation and the incoming configuration does not contain any **replace:** tags, the replace operation is equivalent to a merge operation. If you are running automated scripts and cannot know in advance whether the scripts need to perform a replace or a merge operation, the scripts can use the replace operation to cover either case. If you are performing an override or merge operation and the incoming configuration contains **replace:** tags, the tags are ignored and the override or merge operation is performed.
update	Replaces only the configuration that has changed. An update operation compares the current configuration to the current candidate configuration, and loads only the changes between these configurations in the incoming configuration.

Figure 41 through Figure 43 show the results of override, replace, and merge operations.

Figure 41: Loading a Configuration with the Override Operation

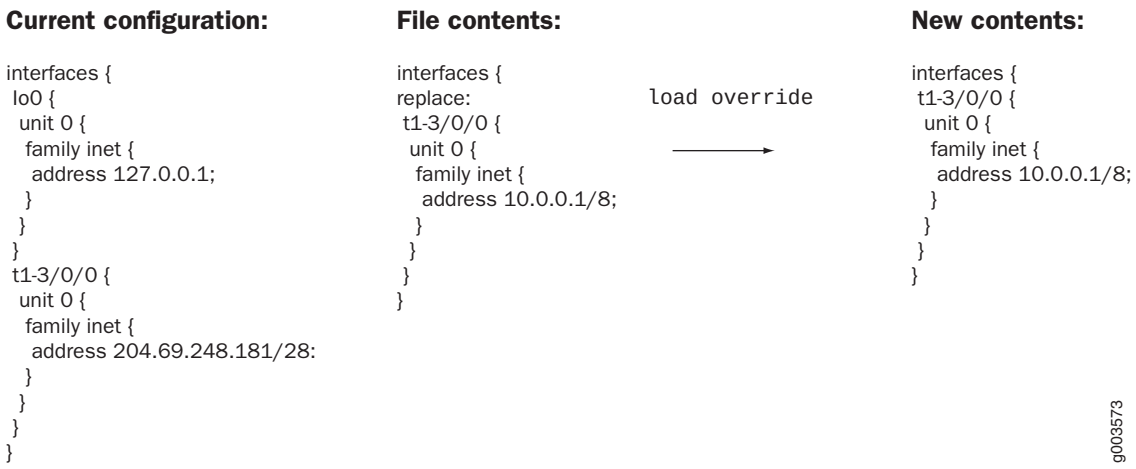

Figure 42: Loading a Configuration with the Replace Operation

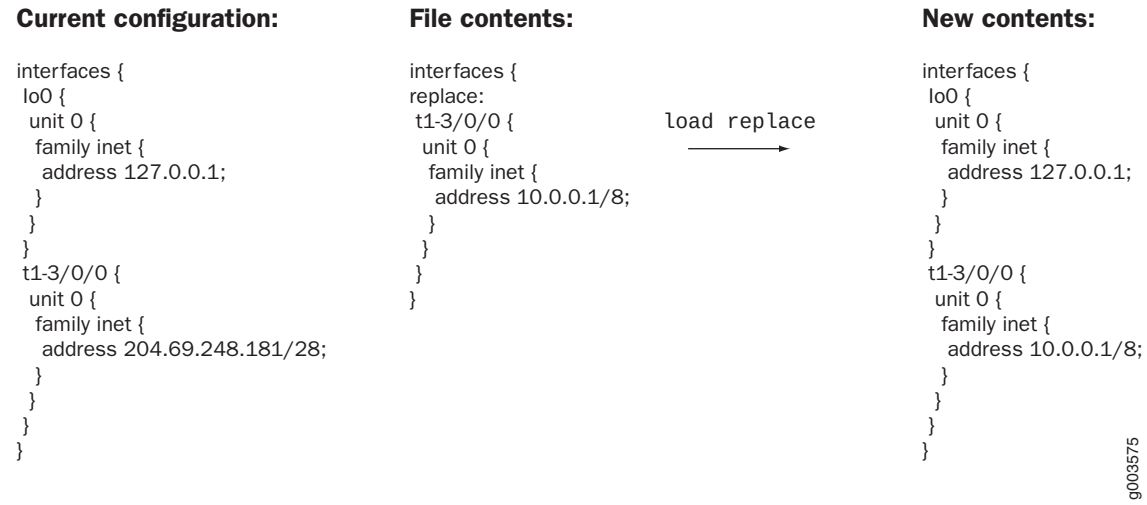

Figure 43: Loading a Configuration with the Merge Operation

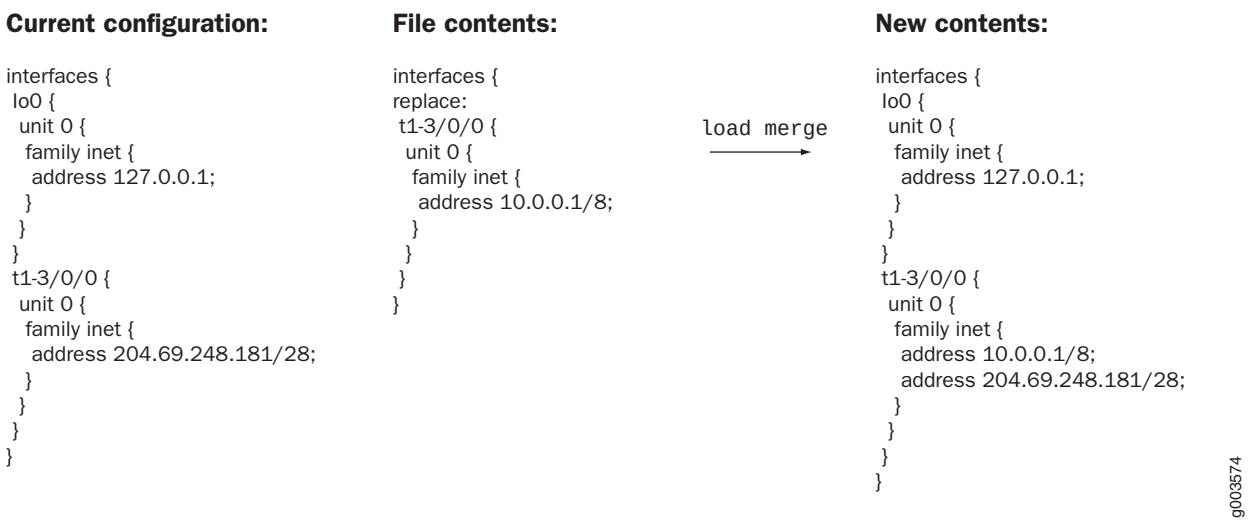

Saving a Configuration File

To save your current configuration to an ASCII file, including any uncommitted changes made by you and all users, issue the save command:

save *filename*

By default, the configuration is saved to a file in your home directory. For information about specifying filenames, see the *JUNOS System Basics Configuration Guide*.

Part 4

Managing the Services Router

Chapter 9
Managing Users and Operations

You can use either J-Web Quick Configuration or a configuration editor to manage system functions, including RADIUS and TACACS+ servers, user login accounts, routine file operations, and system log messages.

This chapter contains the following topics. For more information about system management, see the *JUNOS System Basics Configuration Guide*.

System Management Terms

Before performing system management tasks, become familiar with the terms defined in Table 50.

Table 50: System Management Terms

Term	Definition
Remote Authentication Dial-In User Service (RADIUS)	Authentication method for validating users who attempt to access one or more Services Routers by means of telnet. RADIUS is a multivendor IETF standard whose features are more widely accepted than those of TACACS+ or other proprietary systems. All one-time-password system vendors support RADIUS.
Terminal Access Controller Access Control System Plus (TACACS+)	Authentication method for validating users who attempt to access one or more Services Routers by means of telnet.

System Management Overview

This section contains the following topics:

System Authentication

The JUNOS software supports three methods of user authentication: local password authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+).

With local password authentication, you configure a password for each user allowed to log into the Services Router.

RADIUS and TACACS+ are authentication methods for validating users who attempt to access the router using telnet. Both are distributed client/server systems—the RADIUS and TACACS+ clients run on the router, and the server runs on a remote network system.

You can configure the router to use RADIUS or TACACS+ authentication, or both, to validate users who attempt to access the router. If you set up both authentication methods, you also can configure which the router will try first.

User Accounts

User accounts provide one way for users to access the Services Router. Users can access the router without accounts if you configured RADIUS or TACACS+ servers, as described in "Managing Users with Quick Configuration" on page 169 and "Managing Users and Files with a Configuration Editor" on page 182. After you have created an account, the router creates a home directory for the user. An account for the user root is always present in the configuration. For information about configuring the password for the user root, see "Establishing Basic Connectivity" on page 47. For each user account, you can define the following:

- Username—Name that identifies the user. It must be unique within the router. Do not include spaces, colons, or commas in the username.
- User's full name—If the full name contains spaces, enclose it in quotation marks (" "). Do not include colons or commas.
- User identifier (UID)—Numeric identifier that is associated with the user account name. The identifier must be in the range 100 through 64000 and

must be unique within the router. If you do not assign a UID to a username, the software assigns one when you commit the configuration, preferring the lowest available number.

- User's access privilege—You can create login classes with specific permission bits or use one of the default classes listed in Table 52.

- Authentication method or methods and passwords that the user can use to access the router—You can use SSH or an MD5 password, or you can enter a plain-text password that the JUNOS software encrypts using MD5-style encryption before entering it in the password database. If you configure the plain-text-password option, you are prompted to enter and confirm the password.

Login Classes

All users who log into the Services Router must be in a login class. With login classes, you define the following:

- Access privileges users have when they are logged into the router. For more information, see "Permission Bits" on page 165.

- Commands and statements that users can and cannot specify. For more information, see "Denying or Allowing Individual Commands" on page 167.

- How long a login session can be idle before it times out and the user is logged off.

You can define any number of login classes. You then apply one login class to an individual user account. The software contains a few predefined login classes, which are listed in Table 52. The predefined login classes cannot be modified.

Permission Bits

Each top-level command-line interface (CLI) command and each configuration statement has an access privilege level associated with it. Users can execute only those commands and configure and view only those statements for which they have access privileges. The access privileges for each login class are defined by one or more permission bits (see Table 51).

Two forms for the permissions control the individual parts of the configuration:

- "Plain" form—Provides read-only capability for that permission type. An example is interface.

- Form that ends in -control—Provides read and write capability for that permission type. An example is interface-control.

Table 51: Permission Bits for Login Classes

Permission Bit	Access
admin	Can view user account information in configuration mode and with the **show configuration** command.
admin-control	Can view user accounts and configure them (at the **[edit system login]** hierarchy level).
access	Can view the access configuration in configuration mode and with the **show configuration** operational mode command.
access-control	Can view and configure access information (at the **[edit access]** hierarchy level).
all	Has all permissions.
clear	Can clear (delete) information learned from the network that is stored in various network databases (using the **clear** commands).
configure	Can enter configuration mode (using the **configure** command) and commit configurations (using the **commit** command).
control	Can perform all control-level operations (all operations configured with the **-control** permission bits).
field	Reserved for field (debugging) support.
firewall	Can view the firewall filter configuration in configuration mode.
firewall-control	Can view and configure firewall filter information (at the **[edit firewall]** hierarchy level).
floppy	Can read from and write to the removable media.
interface	Can view the interface configuration in configuration mode and with the **show configuration** operational mode command.
interface-control	Can view chassis, class of service, groups, forwarding options, and interfaces configuration information. Can configure chassis, class of service, groups, forwarding options, and interfaces (at the **[edit]** hierarchy).
maintenance	Can perform system maintenance, including starting a local shell on the router and becoming the superuser in the shell (by issuing the **su root** command), and can halt and reboot the router (using the **request system** commands).
network	Can access the network by entering the **ping**, **ssh**, **telnet**, and **traceroute** commands.
reset	Can restart software processes using the **restart** command and can configure whether software processes are enabled or disabled (at the **[edit system processes]** hierarchy level).
rollback	Can use the **rollback** command to return to a previously committed configuration other than the most recently committed one.
routing	Can view general routing, routing protocol, and routing policy configuration information in configuration and operational modes.
routing-control	Can view general routing, routing protocol, and routing policy configuration information and configure general routing (at the **[edit routing-options]** hierarchy level), routing protocols (at the **[edit protocols]** hierarchy level), and routing policy (at the **[edit policy-options]** hierarchy level).
secret	Can view passwords and other authentication keys in the configuration.
secret-control	Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.
security	Can view security configuration in configuration mode and with the **show configuration** operational mode command.

Permission Bit	Access
security-control	Can view and configure security information (at the [edit security] hierarchy level).
shell	Can start a local shell on the router by entering the start shell command.
snmp	Can view SNMP configuration information in configuration and operational modes.
snmp-control	Can view SNMP configuration information and configure SNMP (at the [edit snmp] hierarchy level).
system	Can view system-level information in configuration and operational modes.
system-control	Can view system-level configuration information and configure it (at the [edit system] hierarchy level).
trace	Can view trace file settings in configuration and operational modes.
trace-control	Can view trace file settings and configure trace file properties.
view	Can use various commands to display current systemwide, routing table, and protocol-specific values and statistics.

Table 52: Predefined Login Classes

Login Class	Permission Bits Set
operator	clear, network, reset, trace, view
read-only	view
super-user and superuser	all
unauthorized	None

Denying or Allowing Individual Commands

By default, all top-level CLI commands have associated access privilege levels. Users can execute only those commands and view only those statements for which they have access privileges. For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that are otherwise permitted or not allowed by a permission bit.

Template Accounts

You use local user template accounts when you need different types of templates. Each template can define a different set of permissions appropriate for the group of users who use that template. These templates are defined locally on the Services Router and referenced by the TACACS+ and RADIUS authentication servers.

When you configure local user templates and a user logs in, the JUNOS software issues a request to the authentication server to authenticate the user's login name. If a user is authenticated, the server returns the local username to the router, which then determines whether a local username is specified for that login name (local-username for TACACS+, Juniper-Local-User for RADIUS). If

so, the router selects the appropriate local user template locally configured on the router. If a local user template does not exist for the authenticated user, the router defaults to the remote template.

For more information, see "Setting Up Template Accounts" on page 189.

System Log Files

The JUNOS software generates system log messages (also called syslog messages) to record events that occur on the Services Router, including the following:

- Routine operations, such as creation of an Open Shortest Path First (OSPF) protocol adjacency or a user login into the configuration database
- Failure and error conditions, such as failure to access a configuration file or unexpected closure of a connection to a child or peer process
- Emergency or critical conditions, such as router power-off due to excessive temperature

The JUNOS system logging utility is similar to the UNIX syslogd utility. Each system log message identifies the software process that generated the message and briefly describes the operation or error that occurred.

When you configure system logging, you can direct messages to one or more destinations:

- To a named file in a local file system
- To the terminal session of one or more specific users (or all users) when they are logged into the router
- To the router console
- To a remote machine that is running the UNIX syslogd utility

Each system log message belongs to a facility, which is a group of messages that are either generated by the same software process or concern a similar condition or activity (such as authentication attempts).

Reboot requests are recorded to the system log files, which you can view with the show log command. Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the show system processes command.

Before You Begin

Before you perform any system management tasks, you must perform the initial Services Router configuration described in "Establishing Basic Connectivity" on page 47.

Managing Users and Files with the J-Web Interface

This section contains the following topics:

Managing Users with Quick Configuration

This section contains the following topics:

Adding a RADIUS Server for Authentication

You can use the Users Quick Configuration page for RADIUS servers to configure a RADIUS server for system authentication. This Quick Configuration page allows you to specify the IP address and secret (password) of the RADIUS server.

Figure 44 shows the Users Quick Configuration page for RADIUS servers.

Figure 44: Users Quick Configuration Page for RADIUS Servers

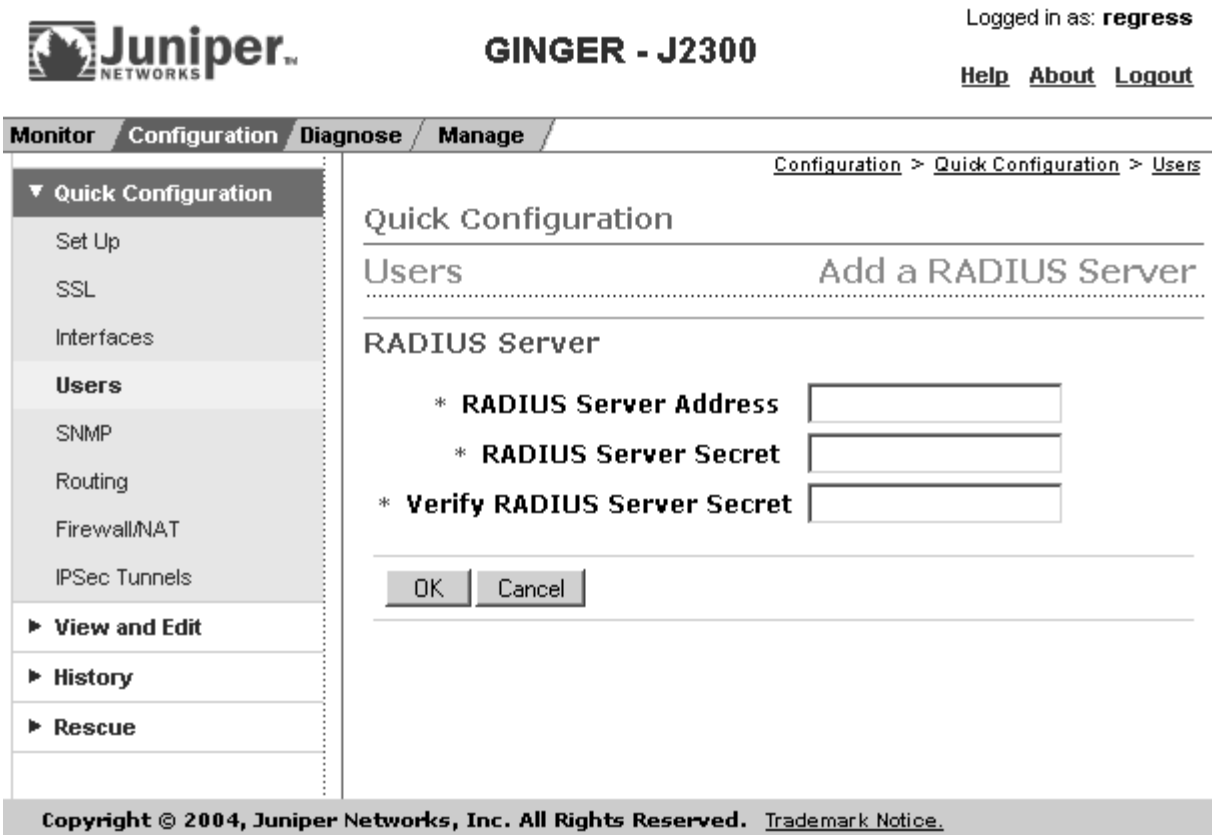

To configure a RADIUS server with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Users**.
2. Under RADIUS servers, click **Add** to configure a RADIUS server.
3. Enter information into the Users Quick Configuration page for RADIUS servers, as described in Table 53.
4. Click one of the following buttons on the Users Quick Configuration page for RADIUS servers:
   - To apply the configuration and return to the Users Quick Configuration page, click **OK**.
   - To cancel your entries and return to the Users Quick Configuration page, click **Cancel**.

Table 53: Users Quick Configuration for RADIUS Servers Summary

Field	Function	Your Action
RADIUS Server		
RADIUS Server Address (required)	Identifies the IP address of the RADIUS server.	Type the RADIUS server's 32-bit IP address, in dotted decimal notation.
RADIUS Server Secret (required)	The secret (password) of the RADIUS server.	Type the secret (password) of the RADIUS server. Secrets can contain spaces. The secret used must match that used by the RADIUS server.
Verify RADIUS Server Secret (required)	Verifies the secret (password) of the RADIUS server is entered correctly.	Retype the secret of the RADIUS server.

Adding a TACACS+ Server for Authentication

You can use the Users Quick Configuration page for TACACS + servers to configure a TACACS + server for system authentication. This Quick Configuration page allows you to specify the IP address and secret of the TACACS + server.

Figure 45 shows the Users Quick Configuration page for TACACS + servers.

Figure 45: Users Quick Configuration Page for TACACS+ Servers

To configure a TACACS+ server with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Users**.
2. Under TACACS+ servers, click **Add** to configure a TACACS+ server.
3. Enter information into the Users Quick Configuration page for TACACS+ servers, as described in Table 54.
4. Click one of the following buttons on the Users Quick Configuration page for TACACS+ servers:
   - To apply the configuration and return to the Users Quick Configuration page, click **OK**.
   - To cancel your entries and return to the Users Quick Configuration page, click **Cancel**.

Table 54: Users Quick Configuration for TACACS+ Servers Summary

Field	Function	Your Action
TACACS+ Server		
TACACS + Server Address (required)	Identifies the IP address of the TACACS + server.	Type the TACACS + server's 32-bit IP address, in dotted decimal notation.
TACACS + Server Secret (required)	The secret (password) of the TACACS + server.	Type the secret (password) of the TACACS + server. Secrets can contain spaces. The secret used must match that used by the TACACS + server.
Verify TACACS + Server Secret (required)	Verifies the secret (password) of the TACACS + server is entered correctly.	Retype the secret of the TACACS + server.

Configuring System Authentication

On the Users Quick Configuration page, you can configure the authentication methods the Services Router uses to verify that a user can gain access. For each login attempt, the router tries the authentication methods in order, starting with the first one, until the password matches.

If you do not configure system authentication, users are verified based on their configured local passwords.

Figure 46 shows the Users Quick Configuration page.

Figure 46: Users Quick Configuration Page

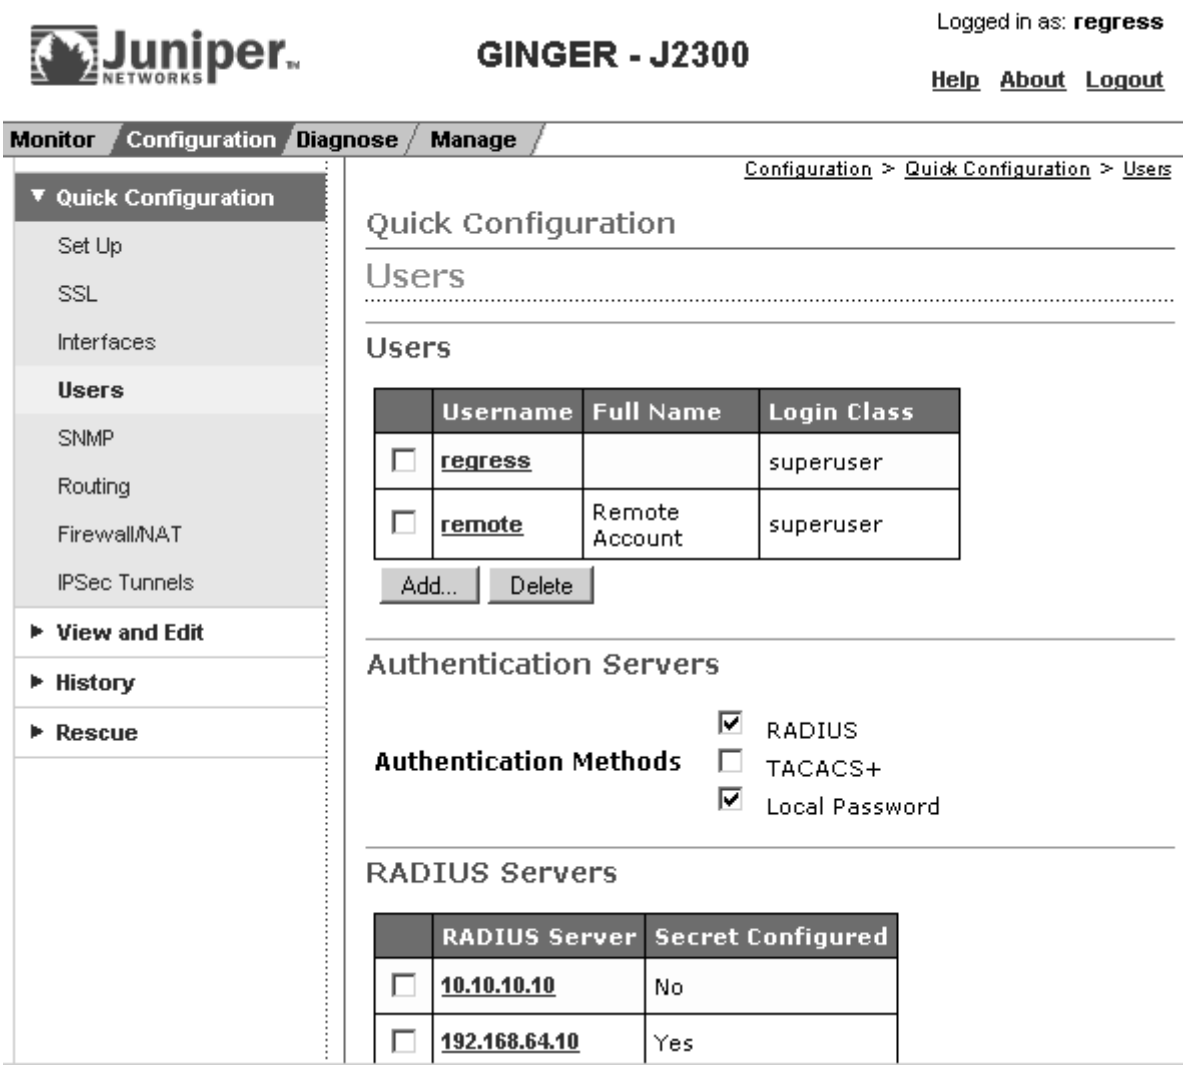

To configure system authentication with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Users**.
2. Under Authentication Servers, select the check box next to each authentication method the router must use when users log in:
   - RADIUS
   - TACACS +
   - Local Password

3. Click one of the following buttons on the Users Quick Configuration page:

   - To apply the configuration and stay in the Users Quick Configuration page, click **Apply**.
   - To apply the configuration and return to the Quick Configuration page, click **OK**.
   - To cancel your entries and return to the Quick Configuration page, click **Cancel**.

Adding New Users

You can use the Users Quick Configuration page for user information to add new users to a Services Router. For each account, you define a login name and password for the user and specify a login class for access privileges.

Figure 47 shows the Quick Configuration page for adding a user.

Figure 47: Add a User Quick Configuration Page

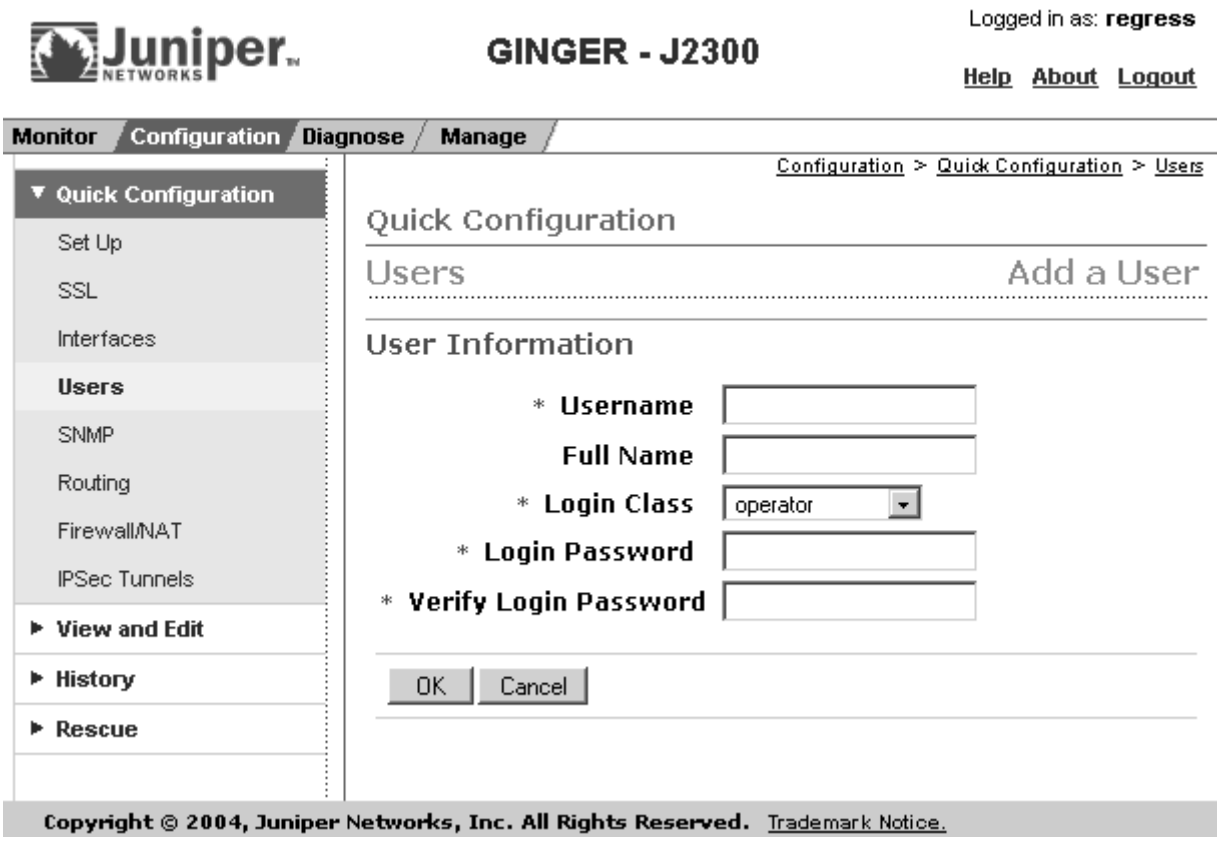

To configure users with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Users**.
2. Under Users, click **Add** to add a new user.
3. Enter information into the Add a User Quick Configuration page, as described in Table 55.
4. Click one of the following buttons on the Add a User Quick Configuration page:
   - To apply the configuration and return to the Users Quick Configuration page, click **OK**.
   - To cancel your entries and return to the Users Quick Configuration page, click **Cancel**.

Table 55: Add a User Quick Configuration Page Summary

Field	Function	Your Action
User Information		
Username (required)	Name that identifies the user.	Type the username. It must be unique within the router. Do not include spaces, colons, or commas in the username.
Full Name	The user's full name.	Type the user's full name. If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
Login Class (required)	Defines the user's access privilege.	From the drop-down list, select the user's login class: ■ **operator** ■ **read-only** ■ **super-user/superuser** ■ **unauthorized** This list also includes any user-defined login classes. For more information, see "Login Classes" on page 165.

Field	Function	Your Action
Login Password (required)	The login password for this user.	Type the login password for this user. The login password must meet the following criteria: ■ The password must be at least 6 characters long. ■ You can include most character classes in a password (alphabetic, numeric, and special characters), except control characters. ■ The password must contain at least one change of case or character class.
Verify Login Password (required)	Verifies the login password for this user.	Retype the login password for this user.

Managing Files with the J-Web Interface

This section contains the following topics:

Cleaning Up Files

You can use the J-Web interface to rotate and delete files on the Services Router. If you are running low on storage space, the file cleanup procedure quickly identifies files that can be deleted.

The file cleanup procedure performs the following tasks:

- Rotates log files—All information in the current log files is archived, and fresh log files are created.
- Deletes log files in /cf/var/log—Any files that are not currently being written to are deleted.
- Deletes temporary files in /cf/var/tmp—Any files that have not been accessed within two days are deleted.
- Deletes all crash files in /cf/var/crash—Any core files that the router has written during an error are deleted.

Figure 48 shows the Clean Up Files page.

Figure 48: Clean Up Files Page

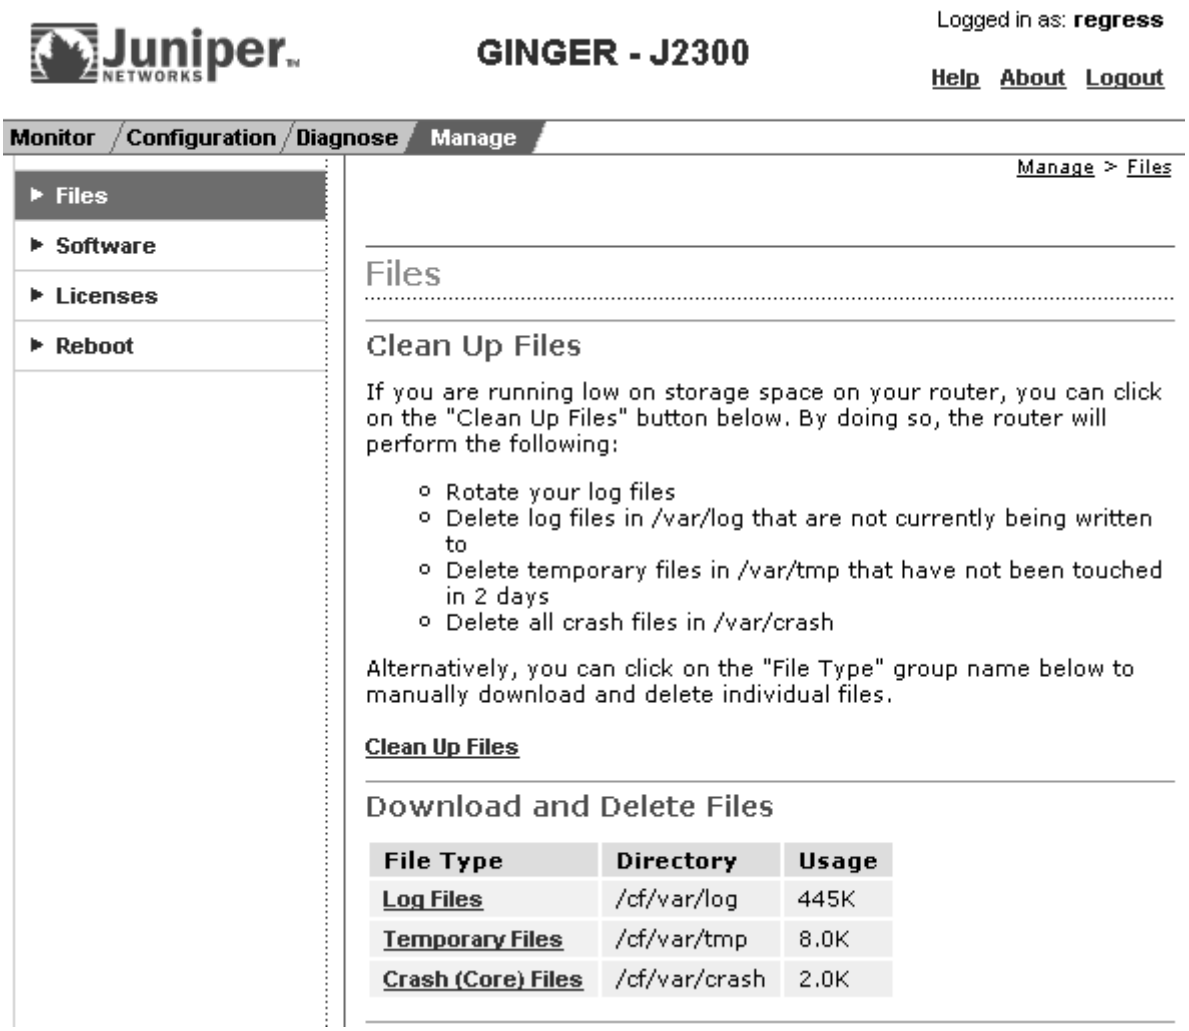

To rotate and delete files with the J-Web interface:

1. In the J-Web interface, select **Manage > Files**.
2. In the Clean Up Files section, click **Clean Up Files**. The router rotates log files and identifies the files that can be safely deleted.
3. The J-Web interface displays the files that you can delete and the amount of space that will be freed on the file system.
4. Click one of the following buttons on the confirmation page:
   - To delete the files and return to the Files page, click **OK**.

- To cancel your entries and return to the list of files in the directory, click **Cancel**.

Downloading Files

You can use the J-Web interface to download a copy of an individual file from the Services Router. When you download a file, it is not deleted from the file system.

Figure 49 shows the J-Web page from which you can download log files.

Figure 49: Log Files Page (Download)

To download files with the J-Web interface:

1. In the J-Web interface, select **Manage > Files**.
2. In the Download and Delete Files section, click one of the following file types:
   - **Log Files**—Lists the log files located in the /cf/var/log directory on the router.
   - **Temporary Files**—Lists the temporary files located in the /cf/var/tmp directory on the router.
   - **Crash (Core) Files**—Lists the core files located in the /cf/var/crash directory on the router.
3. The J-Web interface displays the files located in the directory.
4. To download an individual file, click **Download**.
5. Choose a location for the browser to save the file.

   The file is saved as a text file, with a .txt file extension.
6. To view the file, open it with a text editor.

Deleting Files

You can use the J-Web interface to delete an individual file from the Services Router. When you delete the file, it is permanently removed from the file system.

CAUTION: If you are unsure whether to delete a file from the router, we recommend using the **Cleanup Files** tool described in “Cleaning Up Files” on page 177. This tool determines which files can be safely deleted from the file system.

Figure 50 shows the J-Web page on which you confirm the deletion of files.

Figure 50: Confirm File Delete Page

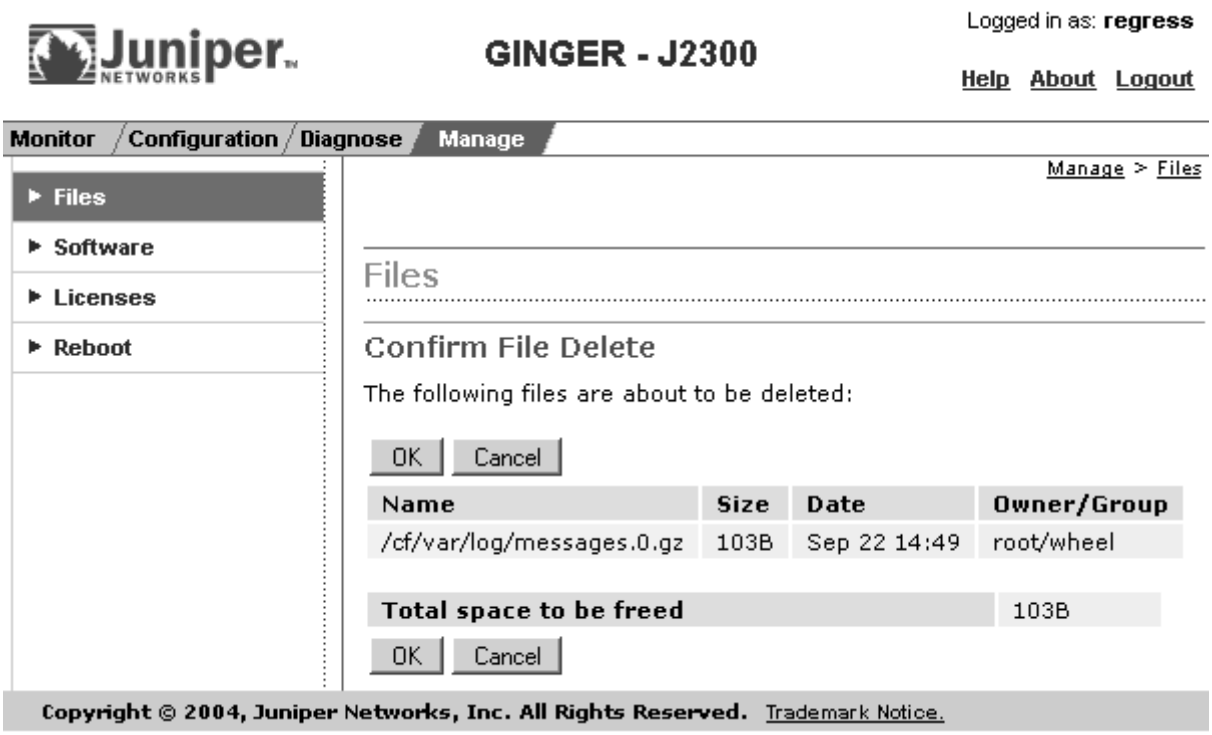

To rotate and delete files with the J-Web interface:

1. In the J-Web interface, select **Manage > Files**.
2. In the Download and Delete Files section, click one of the following file types:
   - **Log Files**—Lists the log files located in the /cf/var/log directory on the router.
   - **Temporary Files**—Lists the temporary files located in the /cf/var/tmp directory on the router.
   - **Crash (Core) Files**—Lists the core files located in the /cf/var/crash directory on the router.
3. The J-Web interface displays the files located in the directory.
4. Check the box next to each file you plan to delete.
5. Click **Delete**.

   The J-Web interface displays the files you can delete and the amount of space that will be freed on the file system.

6. Click one of the following buttons on the confirmation page:

   - To delete the files and return to the Files page, click **OK**.
   - To cancel your entries and return to the list of files in the directory, click **Cancel**.

Managing Users and Files with a Configuration Editor

This section contains the following topics:

Setting Up RADIUS Authentication

To use RADIUS authentication, you must configure at least one RADIUS server.

The procedure provided in this section identifies the RADIUS server, specifies the secret (password) of the RADIUS server, and sets the source address of the Services Router's RADIUS requests to the loopback address of the router. The procedure uses the following sample values:

- The RADIUS server's IP address is 172.16.98.1.
- The RADIUS server's secret is Radiussecret1.
- The loopback address of the router is 10.0.0.1.

To configure RADIUS authentication:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 56.
3. If you are finished configuring the network, commit the configuration.

   To completely set up RADIUS authentication, you must create user template accounts and specify a system authentication order.
4. Go on to one of the following procedures:

- To specify a system authentication order, see “Configuring Authentication Order” on page 185.
- To configure a remote user template account, see “Creating a Remote Template Account” on page 189.
- To configure local user template accounts, see “Creating a Local Template Account” on page 190.

Table 56: Setting Up RADIUS Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System**.	From the top of the configuration hierarchy enter edit system
Add a new RADIUS server	1. In the Radius server box, click **Add new entry**. 2. In the Address box, type the IP address of the RADIUS server: 172.16.98.1	Set the IP address of the RADIUS server: set radius-server address 172.16.98.1
Specify the shared secret (password) of the RADIUS server. The secret is stored as an encrypted value in the configuration database.	In the Secret box, type the shared secret of the RADIUS server: Radiussecret1	Set the shared secret of the RADIUS server: set radius-server 172.16.98.1 secret Radiussecret1
Specify the source address to be included in the RADIUS server requests by the router. In most cases, you can use the loopback address of the router.	In the Source address box, type the loopback address of the router: 10.0.0.1	Set the router’s loopback address as the source address: set radius-server 172.16.98.1 source-address 10.0.0.1

Setting Up TACACS+ Authentication

To use TACACS+ authentication, you must configure at least one TACACS+ server.

The procedure provided in this section identifies the TACACS+ server, specifies the secret (password) of the TACACS+ server, and sets the source address of the Services Router’s TACACS+ requests to the loopback address of the router. This procedure uses the following sample values:

- The TACACS+ server’s IP address is 172.16.98.24.
- The TACACS+ server’s secret is Tacacssecret1.
- The loopback address of the router is 10.0.0.1.

To configure TACACS+ authentication:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 57.
3. If you are finished configuring the network, commit the configuration.

   To completely set up TACACS+ authentication, you must create user template accounts and specify a system authentication order.
4. Go on to one of the following procedures:

   - To specify a system authentication order, see "Configuring Authentication Order" on page 185.
   - To configure a remote user template account, see "Creating a Remote Template Account" on page 189.
   - To configure local user template accounts, see "Creating a Local Template Account" on page 190.

Table 57: Setting Up TACACS+ Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System**.	From the top of the configuration hierarchy enter edit system
Add a new TACACS+ server	1. In the Tacplus server box, click **Add new entry**. 2. In the Address box, type the IP address of the TACACS+ server: 172.16.98.24	Set the IP address of the TACACS+ server: set tacplus-server address 172.16.98.24
Specify the shared secret (password) of the TACACS+ server. The secret is stored as an encrypted value in the configuration database.	In the Secret box, type the shared secret of the TACACS+ server: Tacacssecret1	Set the shared secret of the TACACS+ server: set tacplus-server 172.16.98.24 secret Tacacssecret1
Specify the source address to be included in the TACACS+ server requests by the router. In most cases, you can use the loopback address of the router.	In the Source address box, type the loopback address of the router: 10.0.0.1	Set the router's loopback address as the source address: set tacplus-server 172.16.98.24 source-address 10.0.0.1

Configuring Authentication Order

The procedure provided in this section configures the Services Router to attempt user authentication with the local password first, then with the RADIUS server, and finally with the TACACS+ server.

To configure authentication order:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 58.
3. If you are finished configuring the network, commit the configuration.

   To completely set up RADIUS or TACACS+ authentication, you must configure at least one RADIUS or TACACS+ server and create user template accounts.
4. Go on to one of the following procedures:

   - To configure a RADIUS server, see "Setting Up RADIUS Authentication" on page 182.
   - To configure a TACACS+ server, see "Setting Up TACACS+ Authentication" on page 183.
   - To configure a remote user template account, see "Creating a Remote Template Account" on page 189.
   - To configure local user template accounts, see "Creating a Local Template Account" on page 190.

Table 58: Configuring Authentication Order

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System**.	From the top of the configuration hierarchy enter edit system
Add RADIUS authentication to the authentication order.	1. In the Authentication order box, click **Add new entry**. 2. In the drop-down list, select **radius**. 3. Click **OK**.	Insert the radius statement in the authentication order: insert system authentication-order radius after password
Add TACACS+ authentication to the authentication order.	1. In the Authentication Order box, click **Add new entry**. 2. In the drop-down list, select **tacplus**. 3. Click **OK**.	Insert the tacplus statement in the authentication order: insert system authentication-order tacplus after radius

Controlling User Access

This section contains the following topics:

Defining Login Classes

You can define any number of login classes. You then apply one login class to an individual user account, as described in "Creating User Accounts" on page 188 and "Setting Up Template Accounts" on page 189.

The procedure provided in this section creates a sample login class named operator-and-boot with the following privileges:

- The operator-and-boot login class can reboot the Services Router using the request system reboot command.
- The operator-and-boot login class can also use commands defined in the clear, network, reset, trace, and view permission bits. For more information, see "Permission Bits" on page 165.

To define login classes:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 59.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To create user accounts, see "Creating User Accounts" on page 188.
   - To create shared user accounts, see "Setting Up Template Accounts" on page 189.

Table 59: Defining Login Classes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System Login** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System > Login**.	From the top of the configuration hierarchy enter edit system login

Task	J-Web Configuration Editor	CLI Configuration Editor
Create a login class named operator-and-boot with the ability to reboot the router.	1. Next to Class, click **Add new entry**. 2. Type the name of the login class: operator-and-boot 3. In the Allow commands box, type the request system reboot command enclosed in quotation marks: "request system reboot" 4. Click **OK**.	Set the name of the login class and the ability to use the request system reboot command: set class operator-and-boot allow-commands "request system reboot"
Give the operator-and-boot login class operator privileges.	1. Next to Permissions, click **Add new entry**. 2. In the Value drop-down list, select **clear**. 3. Click **OK**. 4. Next to Permissions, click **Add new entry**. 5. In the Value drop-down list, select **network**. 6. Click **OK**. 7. Next to Permissions, click **Add new entry**. 8. In the Value drop-down list, select **reset**. 9. Click **OK**. 10. Next to Permissions, click **Add new entry**. 11. In the Value drop-down list, select **trace**. 12. Click **OK**. 13. Next to Permissions, click **Add new entry**. 14. In the Value drop-down list, select **view**. 15. Click **OK**.	Set the permission bits for the operator-and-boot login class: set class operator-and-boot permissions [clear network reset trace view]

Creating User Accounts

User accounts provide one way for users to access the Services Router. (Users can access the router without accounts if you configured RADIUS or TACACS+ servers, as described in "Setting Up RADIUS Authentication" on page 182 and "Setting Up TACACS+ Authentication" on page 183.)

The procedure provided in this section creates a sample user named cmartin with the following characteristics:

- The user cmartin belongs to the superuser login class.
- The user cmartin uses an encrypted password, $1$14c5.$sBopasdFFdssdfFFdsdfs0.

To create user accounts:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 60.
3. If you are finished configuring the network, commit the configuration.

Table 60: Creating User Accounts

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System Login** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System > Login**.	From the top of the configuration hierarchy enter edit system login
Create a user named cmartin who belongs to the superuser login class.	1. Next to User, click **Add new entry**. 2. In the User name box, type cmartin. 3. In the Class box, type superuser. 4. Click **OK**.	Set the username and the login class for the user: **set user cmartin class superuser**
Define the encrypted password for cmartin.	1. Next to Authentication, click **Configure**. 2. In the Encrypted password box, type $1$14c5.$sBopasdFFdssdfFFdsdfs0 3. Click **OK**.	Set the encrypted password for cmartin. set user cmartin authentication encrypted-password $1$14c5.$sBopasdFFdssdfFFdsdfs0

Setting Up Template Accounts

You can create template accounts that are shared by a set of users when you are using RADIUS or TACACS+ authentication. When a user is authenticated by a template account, the CLI username is the login name, and the privileges, file ownership, and effective user ID are inherited from the template account.

This section contains the following topics:

Creating a Remote Template Account

You can create a remote template that is applied to users authenticated by RADIUS or TACACS+ that do not belong to a local template account.

By default, the JUNOS software uses the remote template account when

- The authenticated user does not exist locally on the Services Router.
- The authenticated user's record in the RADIUS or TACACS+ server specifies local user, or the specified local user does not exist locally on the router.

The procedure provided in this section creates a sample user named remote that belongs to the operator login class.

To create a remote template account:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 61.
3. If you are finished configuring the network, commit the configuration.

   To completely set up RADIUS or TACACS+ authentication, you must configure at least one RADIUS or TACACS+ server and specify a system authentication order.
4. Go on to one of the following procedures:
   - To configure a RADIUS server, see "Setting Up RADIUS Authentication" on page 182.
   - To configure a TACACS+ server, see "Setting Up TACACS+ Authentication" on page 183.
   - To specify a system authentication order, see "Configuring Authentication Order" on page 185.

Table 61: Creating a Remote Template Account

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System Login** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System > Login**.	From the top of the configuration hierarchy enter edit system login
Create a user named remote who belongs to the operator login class.	1. Next to User, click **Add new entry**. 2. In the User name box, type remote. 3. In the Class box, type operator. 4. Click **OK**.	Set the username and the login class for the user: set user remote class operator

Creating a Local Template Account

You can create a local template that is applied to users authenticated by RADIUS or TACACS+ that are assigned to the local template account. You use local template accounts when you need different types of templates. Each template can define a different set of permissions appropriate for the group of users who use that template.

The procedure provided in this section creates a sample user named admin that belongs to the superuser login class.

To create a local template account:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 62.
3. If you are finished configuring the network, commit the configuration.

   To completely set up RADIUS or TACACS+ authentication, you must configure at least one RADIUS or TACACS+ server and specify a system authentication order
4. Go on to one of the following procedures:

   - To configure a RADIUS server, see “Setting Up RADIUS Authentication” on page 182.
   - To configure a TACACS+ server, see “Setting Up TACACS+ Authentication” on page 183.
   - To configure a system authentication order, see “Configuring Authentication Order” on page 185.

Table 62: Creating a Local Template Account

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System Login** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System > Login**.	From the top of the configuration hierarchy enter edit system login
Create a user named admin who belongs to the superuser login class.	1. Next to User, click **Add new entry**. 2. In the User name box, type admin. 3. In the Class box, type superuser. 4. Click **OK**.	Set the username and the login class for the user: set user admin class superuser

Using System Logs

You can send system logging information to one or more destinations. The destinations can be one or more files, one or more remote hosts, the terminals of one or more users if they are logged in, and the system console.

For each place where you can send system logging information, you specify the class (facility) of messages to log and the minimum severity level (level) of the message.

Table 63 lists the system logging facilities, and Table 64 lists the system logging severity levels. For more information about system log messages, see the *JUNOS System Log Messages Reference*.

Table 63: System Logging Facilities

Facility	Description
any	Any facility
authorization	Any authorization attempt
change-log	Any change to the configuration
cron	Cron scheduling process
daemon	Various system processes
interactive-commands	Commands executed in the CLI
kernel	Messages generated by the JUNOS kernel
user	Messages from random user processes

Table 64: System Logging Severity Levels

Severity Level (from Highest to Lowest Severity)	Description
emergency	Panic or other conditions that cause the system to become unusable.
alert	Conditions that must be corrected immediately, such as a corrupted system database.
critical	Critical conditions, such as hard drive errors.
error	Standard error conditions.
warning	System warning messages.
notice	Conditions that are not error conditions, but that might warrant special handling.
info	Informational messages. This is the default.
debug	Software debugging messages.

This section contains the following topics:

Sending System Log Messages to a File

You can direct system log messages to a file on the compact flash drive. The default directory for log files is /var/log. To specify a different directory on the compact flash drive, include the complete pathname. For the list of logging facilities and severity levels, see Table 63 and Table 64.

For information about archiving log files, see "Archiving System Logs" on page 194.

The procedure provided in this section sends all security-related information to the sample file named security.

To send messages to a file:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 65.
3. If you are finished configuring the network, commit the configuration.

Table 65: Sending Messages to a File

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System Syslog** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System > Syslog**.	From the top of the configuration hierarchy enter edit system syslog
Create a file named security, and send log messages of the authorization class at the severity level info to the file.	1. Next to File, click **Add new entry**. 2. In the File name box, type security. 3. Next to Contents, click **Add new entry**. 4. In the Facility drop-down menu, select **authorization**. 5. In the Level drop-down menu, select **info**.	Set the filename and the facility and severity level: set file security authorization info

Sending System Log Messages to a User Terminal

To direct system log messages to the terminal session of one or more specific users (or all users) when they are logged into the local Routing Engine, specify one or more JUNOS usernames. Separate multiple values with spaces, or use the asterisk (*) to indicate all users who are logged into the local Routing Engine. For the list of logging facilities and severity levels, see Table 63 and Table 64.

The procedure provided in this section sends send any critical messages to the terminal of the sample user frank, if he is logged in.

To send messages to a user terminal:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 66.
3. If you are finished configuring the network, commit the configuration.

Table 66: Sending Messages to a User Terminal

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **System Syslog** level in the configuration hierarchy.	In the configuration editor hierarchy, select **System > Syslog**.	From the top of the configuration hierarchy enter edit system syslog
Send all critical messages to the user frank.	1. Next to User, click **Add new entry**. 2. In the User name box, type frank. 3. Next to Contents, click **Add new entry**. 4. In the Facility drop-down menu, select **any**. 5. In the Level drop-down menu, select **critical**.	Set the filename and the facility and severity level: set user frank any critical

Archiving System Logs

By default, the JUNOS logging utility stops writing messages to a log file when the file reaches 128 KB in size. It closes the file and adds a numerical suffix, then opens and directs messages to a new file with the original name. By default, the logging utility creates up to 10 files before it begins overwriting the contents of the oldest file. The logging utility by default also limits the users who can read log files to the root user and users who have the JUNOS maintenance permission.

To enable all users to read log files, include the world-readable statement at the [edit system syslog archive] hierarchy level. To restore the default permissions, include the no-world-readable statement. You can include the archive statement at the [edit system syslog file *filename*] hierarchy level to configure the number of files, file size, and permissions for the specified log file. For configuration details, see the information about archiving log files in the *JUNOS System Basics Configuration Guide*.

Disabling System Logs

To disable logging of the messages from a facility, use the facility none configuration statement. This statement is useful when, for example, you want to log messages of the same severity level from all but a few facilities. Instead of including a configuration statement for each facility you want to log, you can configure the any level statement and then a facility none statement for each facility you do not want to log. For configuration details, see the information about disabling logging in the *JUNOS System Basics Configuration Guide*.

Accessing Remote Devices with the CLI

This section contains the following topics:

Using the telnet Command

You can use the CLI telnet command to open a telnet session to a remote device:

```
user@host> telnet host <8bit> <bypass-routing> <inet>
<interface interface-name> <no-resolve> <port port>
<routing-instance routing-instance-name> <source address>
```

To escape from the telnet session to the telnet command prompt, press Ctrl-]. To exit from the telnet session and return to the CLI command prompt, enter quit.

Table 67 describes the telnet command options. For more information, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Table 67: CLI telnet Command Options

Option	Description
8bit	Use an 8-bit data path.
bypass-routing	Bypass the routing tables and open a telnet session only to hosts on directly attached interfaces. If the host is not on a directly attached interface, an error message is returned.
host	Open a telnet session to the specified hostname or IP address.
inet	Force the telnet session to an IPv4 destination.
interface *source-interface*	Open a telnet session to a host on the specified interface. If you do not include this option, all interfaces are used.
no-resolve	Suppress the display of symbolic names.
port *port*	Specify the port number or service name on the host.
routing-instance *routing-instance-name*	Use the specified routing instance for the telnet session.
source *address*	Use the specified source address for the telnet session.

Using the ssh Command

You can use the CLI ssh command to use the secure shell (SSH) program to open a connection to a remote device:

```
user@host> ssh host <bypass-routing> <inet>
<interface interface-name> <logical-router logical-router-name>
<routing-instance routing-instance-name> <source address> <v1> <v2>
```

Table 68 describes the ssh command options. For more information, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Table 68: CLI ssh Command Options

Option	Description
bypass-routing	Bypass the routing tables and open an SSH connection only to hosts on directly attached interfaces. If the host is not on a directly attached interface, an error message is returned.
host	Open an SSH connection to the specified hostname or IP address.
inet	Force the SSH connection to an IPv4 destination.
interface *source-interface*	Open an SSH connection to a host on the specified interface. If you do not include this option, all interfaces are used.
routing-instance *routing-instance-name*	Use the specified routing instance for the SSH connection.
source *address*	Use the specified source address for the SSH connection.
v1	Force SSH to use version 1 for the connection.
v2	Force SSH to use version 2 for the connection.

Chapter 10

Monitoring and Diagnosing a Services Router

J-series Services Routers support a suite of J-Web tools and CLI operational mode commands for monitoring and managing system health and performance. Monitoring tools and commands display the current state of the router. Diagnostic tools and commands test the connectivity and reachability of hosts in the network.

This chapter contains the following topics. For complete descriptions of CLI operational mode commands, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference* and the *JUNOS Network and Services Interfaces Command Reference*.

Monitoring and Diagnostic Terms

Before monitoring and diagnosing J-series Services Routers, become familiar with the terms defined in Table 69.

Table 69: J-series Monitoring and Diagnostic Terms

Term	Definition
autonomous system (AS)	Network of nodes that route packets based on a shared map of the network topology stored in their local databases.
Don't Fragment (DF) bit	Bit in the IP header that instructs routers not to fragment a packet. You might set this bit if the destination host cannot reassemble the packet or if you want to test the path maximum transmission unit (MTU) for a destination host.

Term	Definition
Internet Control Message Protocol (ICMP)	TCP/IP protocol used to send error and information messages.
routing instance	Collection of routing tables, interfaces, and routing protocol interfaces. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.
loose source routing	Option in the IP header used to route a packet based on information supplied by the source. A gateway or host must route the packet using the routers specified by this information, but the packet can use other routers along the way.
routing table	Database of routes learned from one or more protocols.
strict source routing	Option in the IP header used to route a packet based on information supplied by the source. A gateway or host must route the packet exactly as specified by this information.
time to live (TTL)	Value (octet) in the IP header that is (usually) decremented by 1 for each hop the packet passes through. If the field reaches zero, the packet is discarded and a corresponding error message is sent to the source of the packet.
type of service (TOS)	Value (octet) in the IP header that defines the service the source host requests, such as the packet's priority and the preferred delay, throughput, and reliability.

Monitoring and Diagnostic Tools Overview

Use the J-Web Monitor, Manage, and Diagnose options to monitor and diagnose a Services Router. J-Web results are displayed in the browser.

You can also monitor and diagnose the router with CLI operational mode commands. CLI command output appears on the screen of your console or management device, or you can filter the output to a file.

This section contains the following topics:

Monitoring Tools Overview

J-Web monitoring tools consist of the options that appear when you select **Monitor** in the task bar. The Monitor options display diagnostic information about the Services Router.

Alternatively, you can enter show commands from the CLI to display the same information, and often greater detail. CLI show commands display the current configuration and information about interfaces, routing protocols, routing

tables, routing policy filters, and the chassis. Use the CLI clear command to clear statistics and protocol database information.

Table 70 describes the function of each J-Web Monitor option and lists the corresponding CLI show commands.

Table 70: J-Web Monitor Options and CLI show Commands

Monitor Option	Function	Corresponding CLI Commands
System	Displays Services Router system properties, such as the system identification and uptime, users, and resource usage. For details, see "Monitoring System Properties" on page 204.	■ show system uptime ■ show system users ■ show system storage ■ show system processes
Chassis	Displays alarm, environment, and hardware information. For details, see "Monitoring the Chassis" on page 206.	■ show chassis alarms ■ show chassis environment ■ show chassis hardware
Interfaces	Hierarchically displays all Services Router physical and logical interfaces, including state and configuration information. For details, see "Monitoring the Interfaces" on page 208.	■ show interfaces terse ■ show interfaces detail ■ show interfaces *interface-name*
Routing	Displays routing information through the following options: ■ Route Information—Displays all routes in the routing table, including protocol, state, and parameter information. You can narrow the list of routes displayed by specifying search criteria. ■ OSPF Information—Displays a summary of OSPF neighbors, interfaces, and statistics. ■ BGP Information—Displays a summary of BGP routing and neighbor information. ■ RIP Information—Displays a summary of RIP neighbors and statistics. For details, see "Monitoring Routing Information" on page 210.	■ Route information ▫ show route terse ▫ show route detail ■ OSPF information ▫ show ospf neighbors ▫ show ospf interfaces ▫ show ospf statistics ■ BGP information ▫ show bgp summary ▫ show bgp neighbor ■ RIP information ▫ show rip statistics ▫ show rip neighbors

Monitor Option	Function	Corresponding CLI Commands
Firewall	Displays firewall and intrusion detection service (IDS) information through the following options: ■ Stateful Firewall—Displays the stateful firewall configuration. ■ IDS Information—Displays information about the configured IDS. For details, see "Monitoring Firewalls " on page 214.	■ Stateful firewall information ▫ show services stateful-firewall conversations ▫ show services stateful-firewall flows ■ IDS information ▫ show services ids destination-table ▫ show services ids source-table ▫ show services ids pair-table
IPSec	Displays configured IPSec tunnels and statistics, and IKE security associations. For details, see "Monitoring IPSec Tunnels" on page 216.	■ show services ipsec-vpn ipsec statistics ■ show services ipsec-vpn ike security-associations
NAT	Displays configured NAT pools. For details, see "Monitoring NAT Pools" on page 217.	■ show services nat pool

J-Web Diagnostic Tools Overview

The J-Web diagnostic tools consist of the options that appear when you select **Diagnose** and **Manage** in the task bar. Table 71 describes the functions of the Diagnose and Manage options.

Table 71: J-Web Interface Diagnose and Manage Options

Option	Function
Diagnose Options	
Ping Host	Allows you to ping a remote host. You can configure advanced options for the ping operation. For details, see "Using the J-Web Ping Host Tool" on page 218.
Traceroute	Allows you to trace a route between the Services Router and a remote host. You can configure advanced options for the traceroute operation. For details, see "Using the J-Web Traceroute Tool" on page 222.
Manage Options	
Files	Allows you manage log, temporary, and core files on the Services Router. For details, see "Managing Files with the J-Web Interface" on page 177.
Upgrade	Allows you to upgrade and manage Services Router software packages. For details, see "Performing Software Upgrades and Reboots" on page 501.

Option	Function
Licenses	Displays a summary of the licenses needed and used for each feature that requires a license. Allows you to add licenses. For details, see "Managing J-series Licenses with the J-Web Interface" on page 71.
Reboot	Allows you to reboot the Services Router at a specified time. For details, see "Rebooting or Halting a Services Router with the J-Web Interface" on page 512.

CLI Diagnostic Commands Overview

The CLI commands available in operational mode allow you to perform the same monitoring, troubleshooting, and management tasks you can perform with the J-Web interface. Instead of invoking the tools through a graphical interface, you use operational mode commands to perform the tasks.

Because the CLI is a superset of the J-Web interface, you can perform certain tasks only through the CLI. For example, you can use the mtrace command to display trace information about a multicast path from a source to a receiver, which is a feature available only through the CLI.

To view a list of top-level operational mode commands, type a question mark (?) at the command-line prompt. (See "CLI Operational Mode" on page 119.)

At the top level of operational mode are the broad groups of CLI diagnostic commands listed in Table 72.

Table 72: CLI Diagnostic Command Summary

Command	Function
Controlling the CLI Environment	
set *option*	Configures the CLI display.
Diagnosis and Troubleshooting	
clear	Clears statistics and protocol database information.
mtrace	Traces information about multicast paths from source to receiver. For details, see "Using mtrace Commands" on page 235.
monitor	Performs real-time debugging of various software components, including the routing protocols and interfaces. For details, see the following sections: ■ "Using the monitor interface Command" on page 229 ■ "Using the monitor traffic Command" on page 231 ■ "Using the monitor file Command" on page 235

Command	Function
ping	Determines the reachability of a remote network host. For details, see "Using the ping Command" on page 226.
test	Tests the configuration and application of policy filters and AS path regular expressions.
traceroute	Traces the route to a remote network host. For details, see "Using the traceroute Command" on page 228.
Connecting to Other Network Systems	
ssh	Opens secure shell connections. For details, see "Using the ssh Command" on page 195.
telnet	Opens telnet sessions to other hosts on the network. For details, see "Using the telnet Command" on page 195.
Management	
copy	Copies files from one location on the Services Router to another, from the router to a remote system, or from a remote system to the router.
restart *option*	Restarts the various JUNOS software processes, including the routing protocol, interface, and SNMP processes.
request	Performs system-level operations, including stopping and rebooting the Services Router and loading JUNOS software images.
start	Exits the CLI and starts a UNIX shell.
configuration	Enters configuration mode. For details, see "CLI Configuration Mode" on page 120.
quit	Exits the CLI and returns to the UNIX shell.

Filtering Command Output

For operational commands that display output, such as the show commands, you can redirect the output into a filter or a file. When you display help about these commands, one of the options listed is |, called a *pipe*, which allows you to filter the command output.

For example, if you enter the show configuration command, the complete Services Router configuration is displayed on the screen. To limit the display to only those lines of the configuration that contain address, issue the show configuration command using a pipe into the match filter:

```
user@host> show configuration | match address

address-range low 192.168.3.2 high 192.168.3.254;
address-range low 192.168.71.71 high 192.168.71.254;
address 192.168.71.70/21;
address 192.168.2.1/24;
```

```
address 127.0.0.1/32;
```

For a complete list of the filters, type a command, followed by the pipe, followed by a question mark (?):

```
    user@host> show configuration | ?

Possible completions:
  compare              Compare configuration changes with prior version
  count                Count occurrences
  display              Show additional kinds of information
  except               Show only text that does not match a pattern
  find                 Search for first occurrence of pattern
  hold                 Hold text without exiting the --More-- prompt
  last                 Display end of output only
  match                Show only text that matches a pattern
  no-more              Don't paginate output
  request              Make system-level requests
  resolve              Resolve IP addresses
  save                 Save output text to file
  trim                 Trim specified number of columns from start of line
```

You can specify complex expressions as an option for the match and except filters. For more information about command output filtering and creating match expressions, see the *JUNOS System Basics Configuration Guide*.

NOTE: To filter the output of configuration mode commands, use the filter commands provided for the operational mode commands. In configuration mode, an additional filter is supported. See "Filtering Configuration Command Output" on page 130.

Before You Begin

To use the J-Web interface and CLI operational tools, you must have the appropriate access privileges. For more information about configuring access privilege levels, see "Adding New Users" on page 175 and the *JUNOS System Basics Configuration Guide*.

Using the Monitoring Tools

This section describes the monitoring tools in detail. It contains the following topics:

- Monitoring Routing Information on page 210
- Monitoring Firewalls on page 214
- Monitoring IPSec Tunnels on page 216
- Monitoring NAT Pools on page 217

Monitoring System Properties

The system properties include everything from the name and IP address of the Services Router to the resource usage on the Routing Engine. To view these system properties, select **Monitor > System** in the J-Web interface, or enter the following CLI show commands:

- show system uptime
- show system users
- show system storage
- show system processes

Table 73 summarizes key output fields in system properties displays.

Table 73: Summary of Key System Properties Output Fields

Field	Values	Additional Information
System Identification		
Serial Number	Serial number for the J-series Services Router.	
JUNOS Software Version	Version of JUNOS software active on the Services Router.	
Router Hostname	Hostname of the Services Router, as defined with the set system hostname command.	
Router IP Address	IP address, in dotted decimal notation, of the Ethernet management port (fe-0/0/0), as defined with the set interfaces fe-0/0/0 command.	
Loopback Addresses	IP address, in dotted decimal notation, of the loopback address, as defined with the set interfaces lo0 command.	
Domain Name Servers	IP addresses, in dotted decimal notation, of the domain name servers, as defined with the set system name-server command.	
Time Zone	Time zone of the Services Router, as defined with the set system time-zone command.	
System Time		

Field	Values	Additional Information
Current Time	Current system time, in Coordinated Universal Time (UTC).	
System Booted Time	Date and time when the router was last booted and how long it has been running.	
Protocol Started Time	Date and time when the routing protocols were last started and how long they have been running.	
Last Configured Time	Date and time when a configuration was last committed. This field also shows the name of the user who issued the last commit command, through either the J-Web interface or the CLI.	
Users		
User	Username of any user logged in to the Services Router.	
TTY	Terminal through which the user is logged in.	
From	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.	
Login Time	Time when the user logged in.	This is the LOGIN@ field in show system users command output.
Idle Time	How long the user has been idle.	
Command	Processes that the user is running.	This is the WHAT field in show system users command output.
Memory Usage		
Total Memory Available	Total RAM available on the Services Router.	
Total Memory Used	Total RAM currently being consumed by processes actively running on the Services Router, displayed both as a quantity of memory and as a percentage of the total RAM on the router.	
Process ID	Process identifier.	This is the PID field in show system processes command output.
Process Owner	Name of the process owner.	
Process Name	Command that is currently running.	Individual processes on the Services Router are listed here. Because each process within JUNOS operates in a protected memory environment, you can diagnose whether a particular process is consuming abnormally high amounts of resources. If a software process is using too much CPU or memory, you can restart the process by entering the restart command from the CLI.
CPU Usage	Percentage of the CPU that is being used by the process.	

Field	Values	Additional Information
Memory Usage	Percentage of the installed RAM that is being used by the process.	
CPU Usage		
Total CPU Used	Sum of CPU usages by all processes, expressed as a percentage of total CPU available.	
Process ID	Process identifier.	This is the PID field in show system processes command output.
Process Owner	Name of the process' owner.	
Process Name	Command that is currently running.	Individual processes on the Services Router are listed here. Because each process within JUNOS operates in a protected memory environment, you can diagnose whether a particular process is consuming an abnormal amount of resources. If a software process is using too much CPU or memory, you can restart the process by entering the restart command from the CLI.
CPU Usage	Percentage of the CPU that is being used by the process.	
Memory Usage	Percentage of the installed RAM that is being used by the process.	
System Storage		
Total Flash Size	Total size, in megabytes, of the primary flash device.	
Usable Flash Size	Total usable memory, in megabytes, of the primary flash device.	The total usable flash memory is the total memory minus the size of the JUNOS image installed on the Services Router.
Flash Used	Total flash memory used, in megabytes and as a percentage of the total usable flash size, of the primary flash device.	
Log Files	Total size, in kilobytes, of the log files on the Services Router.	This is the sum of file sizes in the /var/log directory.
Temporary Files	Total size, in kilobytes, of the temporary files on the Services Router.	This is the sum of the file sizes in the /var/tmp directory.
Crash (Core) Files	Total size, in kilobytes, of the core files on the Services Router.	This is the sum of the file sizes in the /var/crash directory.
Database Files	Total size, in kilobytes, of the configuration database files on the Services Router.	This is the sum of the file sizes in the /var/db directory.

Monitoring the Chassis

The chassis properties include the status of any alarms on the Services Router, environment measurements, and a summary of the field-replaceable units (FRUs)

on the router. To view these chassis properties, select **Monitor > Chassis** in the J-Web interface, or enter the following CLI show commands:

- show chassis alarms
- show chassis environment
- show chassis hardware

Table 74 summarizes key output fields in chassis displays.

Table 74: Summary of Key Chassis Output Fields

Field	Values	Additional Information
Alarm Summary		
Alarm Time	Date and time alarm was first recorded.	
Alarm Class	Severity class for this alarm: Minor or Major.	JUNOS has system-defined alarms and configurable alarms. System-defined alarms include FRU detection alarms (power supplies removed, for instance) and environmental alarms. The values for these alarms are defined within JUNOS. Configurable alarms are set in either of the following ways: ■ In the J-Web configuration editor, on the **Chassis > Alarm >** *interface-type* page ■ In the CLI configuration editor, with the alarm statement at the [edit chassis] level of the configuration hierarchy
Alarm Description	A brief synopsis of the alarm.	
Environment Information		
Name	Chassis component. For J-series Services Routers, the chassis components are the Routing Engine, flexible PIM concentrator (FPC), and physical interface module (PIM)—identified in the display as a PIC.	On Services Routers, an FPC and a PIM are the same physical unit.
Gauge Status	Status of the temperature gauge on the specified hardware component.	
Temperature	Temperature of the air flowing past the hardware component.	
Hardware Summary		
Name	Chassis component. For J-series Services Routers, the chassis components are the Routing Engine, FPC, and PIM—identified in the display as a PIC.	On Services Routers, an FPC and a PIM are the same physical unit.
Version	Revision level of the specified hardware component.	Supply the version number when reporting any hardware problems to customer support.

Field	Values	Additional Information
Part Number	Part number of the chassis component.	
Serial Number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis.	Use this serial number when you need to contact customer support about the router chassis.
Description	Brief description of the hardware item.	For PIMs, the description lists the number and type of the ports on the PIM—identified in the display as a PIC.

Monitoring the Interfaces

The interface information is divided into multiple parts. To view general interface information such as available interfaces, operation states of the interfaces, and descriptions of the configured interfaces, select **Monitor > Interfaces** in the J-Web interface. To view interface-specific properties such as administrative state or traffic statistics in the J-Web interface, select the interface name on the Interfaces page.

Alternatively, enter the following CLI show commands:

- show interfaces terse
- show interfaces detail
- show interfaces *interface-name*

Table 75 summarizes key output fields in interfaces displays.

Table 75: Summary of Key Interfaces Output Fields

Field	Values	Additional Information
Interface Summary		
Interface Name	Name of interface.	Click an interface name to see more information about the interface.
Oper State	Link state of the interface: Up or Down.	The operational state is the physical state of the interface. If the interface is physically operational, even if it is not configured, the operational state is Up. An operational state of Down indicates a problem with the physical interface.
Admin State	Whether the interface is enabled up (Up) or disabled (Down).	Interfaces are enabled by default. To disable an interface: ■ In the J-Web configuration editor, select the **Disable** check box on the **Interfaces >** *interfaces-name* page. ■ In the CLI configuration editor, add the disable statement at the [edit interfaces *interfaces-name*] level of the configuration hierarchy

Field	Values	Additional Information
Description	Configured description for the interface.	
Interface: *interface-name*		
State	Link state of the interface: Up or Down.	The operational state is the physical state of the interface. If the interface is physically operational, even if it is not configured, the operational state is Up. An operational state of Down indicates a problem with the physical interface.
Admin State	Whether the interface is enabled up (Up) or disabled (Down).	Interfaces are enabled by default. To disable an interface: ■ In the J-Web configuration editor, select the **Disable** check box on the **Interfaces >** *interfaces-name* page. ■ In the CLI configuration editor, add the disable statement at the [edit interfaces *interfaces-name*] level of the configuration hierarchy
MTU	Maximum transmission unit (MTU) size on the physical interface.	
Speed	Speed at which the interface is running.	
Current Address	Configured media access control (MAC) address.	
Hardware Address	Hardware MAC address.	
Last Flapped	Date, time, and how long ago the interface changed state from Down to Up.	
Active Alarms	List of any active alarms on the interface.	Configure alarms on interfaces as follows: ■ In the J-Web configuration editor, on the **Chassis > Alarm >** *interface-type* page ■ In the CLI configuration editor, with the alarm statement at the [edit chassis] level of the configuration hierarchy
Traffic Statistics	Number of packets and bytes received and transmitted on the physical interface.	
Input Errors	Input errors on the interface. (See the following rows of this table for specific error types.)	
Drops	Number of packets dropped by the output queue.	If the interface is saturated, this number increments once for every packet that is dropped by the Services Router's random early detection (RED) mechanism.
Framing errors	Sum of ATM Adaptation Layer (AAL5) packets that have frame check sequence (FCS) errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.	
Policed discards	Number of packets dropped as a result of routing policies configured on the interface.	

Monitoring Routing Information

Routing information is divided into multiple parts:

- To view the inet.0 (IPv4) routing table in the J-Web interface, select **Monitor > Routing > Route Information**, or enter the following CLI commands:
  - show route terse
  - show route detail
- To view BGP routing information, select **Monitor > Routing > BGP Information**, or enter the following CLI commands:
  - show bgp summary
  - show bgp neighbor
- To view OSPF routing information, select **Monitor > Routing > OSPF Information**, or enter the following CLI commands:
  - show ospf neighbors
  - show ospf interfaces
  - show ospf statistics
- To view RIP routing information, select **Monitor > Routing > RIP Information**, or enter the following CLI commands:
  - show rip statistics
  - show rip neighbors

Table 76 summarizes key output fields in routing displays.

Table 76: Summary of Key Routing Output Fields

Field	Values	Additional Information
Route Information		
n destinations	Number of destinations for which there are routes in the routing table.	

Field	Values	Additional Information
n routes	Number of routes in the routing table: ■ active—Number of routes that are active. ■ holddown—Number of routes that are in hold-down state (neither advertised nor updated) before being declared inactive. ■ hidden—Number of routes not used because of routing policies configured on the Services Router.	
Destination	Destination address of the route.	
Protocol/ Preference	Protocol from which the route was learned: Static, Direct, Local, or the name of a particular protocol. The preference is the individual preference value for the route.	The route preference is used as one of the route selection criteria.
Next-Hop	Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	If a next hop is listed as Discard, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the discard attribute has been set. If a next hop is listed as Reject, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected. If a next hop is listed as Local, the destination is an address on the host (either the loopback address or the Ethernet management port address, for example).
Age	How long the route has been known.	
State	Flags for this route.	There are many possible flags. For a complete description, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.
AS Path	AS path through which the route was learned. The letters of the AS path indicate the path origin: ■ I — IGP. ■ E — EGP. ■ ? — Incomplete. Typically, the AS path was aggregated.	
BGP Summary		
Groups	Number of BGP groups.	
Peers	Number of BGP peers.	

Field	Values	Additional Information
Down Peers	Number of unavailable BGP peers.	
Peer	Address of each BGP peer.	
InPkt	Number of packets received from the peer,	
OutPkt	Number of packets sent to the peer.	
Flaps	Number of times a BGP session has changed state from Down to Up.	A high number of flaps might indicate a problem with the interface on which the BGP session is enabled.
Last Up/Down	Last time that a session became available or unavailable, since the neighbor transitioned to or from the established state.	If the BGP session is unavailable, this time might be useful in determining when the problem occurred.
State	A multipurpose field that displays information about BGP peer sessions. The contents of this field depend upon whether a session is established. ■ If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. ■ If a BGP session is established, the field shows the number of active, received, and damped routes that are received from a neighbor. For example, 2/4/0 indicates two active routes, four received routes, and no damped routes.	
BGP Neighbors		
Peer	Address of the BGP neighbor.	
AS	AS number of the peer.	
Type	Type of peer: Internal or External.	
State	Current state of the BGP session: ■ Active—BGP is initiating a TCP connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message. ■ Connect—BGP is waiting for the TCP connection to become complete. ■ Established—The BGP session has been established, and the peers are exchanging BGP update messages. ■ Idle—This is the first stage of a connection. BGP is waiting for a Start event. ■ OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. ■ OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.	Generally, the most common states are Active, which indicates a problem establishing the BGP conenction, and Established, which indicates a successful session setup. The other states are transition states, and BGP sessions normally do not stay in those states for extended periods of time.

Field	Values	Additional Information
Export	Names of any export policies configured on the peer.	
Import	Names of any import policies configured on the peer.	
Number of flaps	Number of times the BGP sessions has changed state from Down to Up.	A high number of flaps might indicate a problem with the interface on which the session is established.
OSPF Neighbors		
Address	Address of the neighbor.	
Interface	Interface through which the neighbor is reachable.	
State	State of the neighbor: Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2way.	Generally, only the Down state, indicating a failed OSPF adjacency, and the Full state, indicating a functional adjacency, are maintained for more than a few seconds. The other states are transitional states that a neighbor is in only briefly while an OSPF adjacency is being established.
ID	Router ID of the neighbor.	
Priority	Priority of the neighbor to become the designated router.	
Dead	Number of seconds until the neighbor becomes unreachable.	
OSPF Interfaces		
Interface	Name of the interface running OSPF.	
State	State of the interface: BDR, Down, DR, DRother, Loop, PtToPt, or Waiting.	The Down state, indicating that the interface is not functioning, and PtToPt state, indicating that a point-to-point connection has been established, are the most common states.
Area	Number of the area that the interface is in.	
DR ID	Address of the area's designated router.	
BDR ID	Address of the area's backup designated router.	
Nbrs	Number of neighbors on this interface.	
OSPF Statistics		
Packet Type	Type of OSPF packet.	
Total Sent/Total Received	Total number of packets sent and received.	
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.	
Receive errors	Number and type of receive errors.	
RIP Statistics		

Field	Values	Additional Information
Rip info	Information about RIP on the specified interface, including UDP port number, hold-down interval (during which routes are neither advertised nor updated), and timeout interval.	
Logical interface	Name of the logical interface on which RIP is configured.	
Routes learned	Number of RIP routes learned on the logical interface.	
Routes advertised	Number of RIP routes advertised on the logical interface.	
RIP Neighbors		
Neighbor	Name of the RIP neighbor.	This value is the name of the interface on which RIP is enabled. The name is set in either of the following ways: ■ In the J-Web configuration editor, on the **Protocols > RIP > Group >** *group-name* **> Neighbor** page ■ In the CLI configuration editor, with the neighbor *neighbor-name* statement at the [edit protocols rip group *group-name*] level of the configuration hierarchy
State	State of the RIP connection: Up or Dn (Down).	
Source Address	Local source address.	This value is the configured address of the interface on which RIP is enabled.
Destination Address	Destination address.	This value is the configured address of the immediate RIP adjacency.
In Met	Value of the incoming metric configured for the RIP neighbor.	

Monitoring Firewalls

Firewall information is divided into multiple parts:

- To view stateful firewall information in the J-Web interface, select **Monitor > Firewall > Stateful Firewall**. To display firewall information for a particular address prefix, port, or other characteristic, type or select information in one or more of the Narrow Search boxes, and click **OK**.

Alternatively, enter the following CLI show commands:

- show services stateful-firewall conversations
- show services stateful-firewall flows

- To view intrusion detection service (IDS) information, select **Monitor > Firewall > IDS Information**. Click one of the following criteria to order the display accordingly:
  - **Bytes** (received bytes)
  - **Packets** (received packets)
  - **Flows**
  - **Anomalies**

  To limit the display of IDS information, type or select information in one or more of the Narrow Search boxes listed in Table 77, and click **OK**.

Table 77: IDS Search-Narrowing Characteristics

Narrow Search Box	Entry or Selection
Destination Address	Type a destination address prefix to display IDS information for only that prefix.
IDS Table	Select one of the following: ■ **Destination**—Displays information for an address under attack. ■ **Pair**—Displays information for a suspected attack source and destination pair. ■ **Source**—Displays information for an address that is a suspected attacker.
Number of IDS Entries to Display	Select a number between **25** and **500** to display only a particular number of entries.
Threshold	Type a number to display events with only that number of bytes, packets, flows, or anomalies—whichever you selected to order the display. For example, to display all events with more than 100 flows, click **Flows** and then type 100 in the Threshold box.
Service Set	Select a service set to display information for only the set.

Alternatively, enter the following CLI show commands:

- show services ids destination-table
- show services ids source-table
- show services ids pair-table

Table 78 summarizes key output fields in firewall and IDS displays.

Table 78: Summary of Key Firewall and IDS Output Fields

Field	Values
Stateful Firewall	
Protocol	Protocol used for the specified stateful firewall flow.
Source IP	Source prefix of the stateful firewall flow.
Source Port	Source port number of stateful firewall flow.
Destination IP	Destination prefix of the stateful firewall flow.
Destination Port	Destination port number of the stateful firewall flow.
Flow State	Status of the stateful firewall flow: ■ Drop—Drop all packets in the flow without response. ■ Forward—Forward the packet in the flow without inspecting it. ■ Reject—Drop all packets in the flow with response. ■ Watch—Inspect packets in the flow.
Direction	Direction of the flow: I (input) or O (output).
Frames	Number of frames in the flow.
IDS Information	
Source Address	Source address for the event.
Destination address	Destination address for the event.
Time	Total time the information has been in the IDS table.
Bytes	Total number of bytes sent from the source to the destination address, in thousands (k) or millions (m).
Packets	Total number of packets sent from the source to the destination address, in thousands (k) or millions (m).
Flows	Total number of flows of packets sent from the source to the destination address, in thousands (k) or millions (m).
Anomalies	Total number of anomalies in the anomaly table, in thousands (k) or millions (m).
Application	Configured application, such as FTP or telnet.

Monitoring IPSec Tunnels

IPSec tunnel information includes information about active IPSec tunnels configured on the Services Router, as well as traffic statistics through the tunnels. To view IPSec tunnel information, select **Monitor > IPSec** in the J-Web interface, or enter the following CLI show commands:

- show services ipsec-vpn ipsec statistics
- show services ipsec-vpn ike security-associations

Table 79 summarizes key output fields in IPSec displays.

Table 79: Summary of Key IPSec Output Fields

Field	Values
IPSec Tunnels	
Service Set	Name of the service set for which the IPSec tunnel is defined.
Rule	Name of the rule set applied to the IPSec tunnel.
Term	Name of the IPSec term applied to the IPSec tunnel.
Local Gateway	Gateway address of the local system.
Remote Gateway	Gateway address of the remote system.
Direction	Direction of the IPSec tunnel: Inbound or Outbound.
Protocol	Protocol supported: either Encapsulation Security Protocol (ESP) or Authentication Header and ESP (AH+ESP).
Tunnel Index	Numeric identifier of the IPSec tunnel.
Tunnel Local Identity	Prefix and port number of the local endpoint of the IPSec tunnel.
Tunnel Remote Identity	Prefix and port number of the remote endpoint of the IPSec tunnel.
IPSec Statistics	
Service Set	Name of the service set for which the IPSec tunnel is defined.
Local Gateway	Gateway address of the local system.
Remote Gateway	Gateway address of the remote system.
ESP Encrypted Bytes	Total number of bytes encrypted by the local system across the IPSec tunnel.
ESP Decrypted Bytes	Total number of bytes decrypted by the local system across the IPSec tunnel.
AH Input Bytes	Total number of bytes received by the local system across the IPSec tunnel.
AH Output Bytes	Total number of bytes transmitted by the local system across the IPSec tunnel.

Monitoring NAT Pools

NAT pool information includes information about the address ranges configured within the pool on the Services Router. To view NAT pool information, select **Monitor > NAT** in the J-Web interface, or enter the following CLI show command:

- show services nat pool

Table 80 summarizes key output fields in NAT displays.

Table 80: Summary of Key NAT Output Fields

Field	Values
NAT Pools	
NAT Pool	Name of the NAT pool.
Pool Start Address	Lower address in the NAT pool address range.

Field	Values
Pool Address End	Upper address in the NAT pool address range.
Port High	Upper port in the NAT pool port range.
Port Low	Lower port in the NAT pool port range.
Ports In Use	Number of ports allocated in this NAT pool.

Using J-Web Diagnostic Tools

This section contains the following topics:

Using the J-Web Ping Host Tool

You can use the ping host diagnostic tool to verify that a host can be reached over the network. The output is useful for diagnosing host and network connectivity problems. The Services Router sends a series of ICMP echo (ping) requests to a specified host and expects to receive ICMP echo responses.

Alternatively, you can use the CLI ping command. (See "Using the ping Command" on page 226.)

To use the ping host tool:

1. Select **Diagnose** from the task bar.
2. Next to Advanced options, click the expand icon (see Figure 51).
3. Enter information into the Ping Host page, as described in Table 81.

   The Remote Host field is the only required field.
4. Click **Start**.

   The results of the ping operation are displayed in the main pane (see Figure 52). If no options are specified, each ping response is in the following format:

   ```
   bytes bytes from ip-address:  icmp_seq=number ttl=number time=time
   ```

   Table 82 summarizes the output fields of the display.
5. To stop the ping operation before it is complete, click **OK**.

Figure 51: Ping Host Page

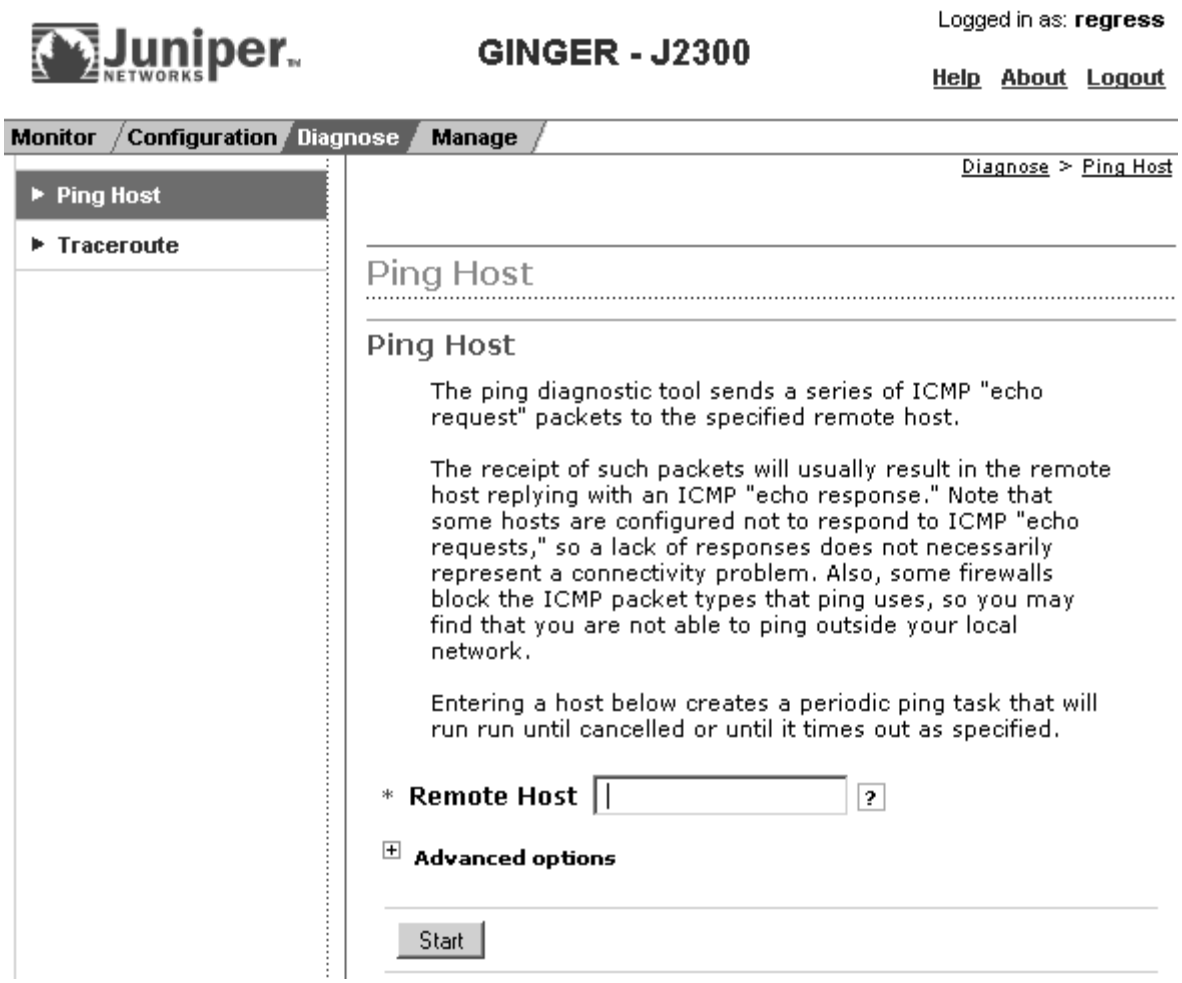

Table 81: J-Web Ping Host Summary

Field	Function	Your Action
Remote Host	Identifies the host to ping.	Type the hostname or IP address of the host to ping.
Advanced Options		
Don't Resolve Addresses	Determines whether to display hostnames of the hops along the path.	■ To suppress the display of the hop hostnames, select the check box. ■ To display the hop hostnames, clear the check box.
Interval	Specifies the interval, in seconds, between the transmission of each ping request.	From the drop-down list, select the interval.

Field	Function	Your Action
Packet Size	Specifies the size of the ping request packet.	Type the size, in bytes, of the packet. The size can be from 0 through 65468. The router adds 8 bytes of ICMP header to the size.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address.
Time-to-Live	Specifies the time-to-live (TTL) hop count for the ping request packet.	From the drop-down list, select the TTL.
Bypass Routing	Determines whether ping requests are routed by means of the routing table. If the routing table is not used, ping requests are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, ping responses are not sent.	■ To bypass the routing table and send the ping requests to hosts on the specified interface only, select the check box. ■ To route the ping requests using the routing table, clear the check box.
Interface	Specifies the interface on which the ping requests are sent.	From the drop-down list, select the interface on which ping requests are sent. If you select **any**, the ping requests are sent on all interfaces.
Count	Specifies the number of ping requests to send.	From the drop-down list, select the number of ping requests to send.
Don't Fragment	Specifies the Don't Fragment (DF) bit in the IP header of the ping request packet.	■ To set the DF bit, select the check box. ■ To clear the DF bit, clear the check box.
Record Route	Sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed in the main pane.	■ To record and display the path of the packet, select the check box. ■ To suppress the recording and display of the path of the packet, clear the check box.
Type-of-Service	Specifies the type-of-service (TOS) value in the IP header of the ping request packet.	From the drop-down list, select the decimal value of the TOS field.

Figure 52: Ping Host Results Page

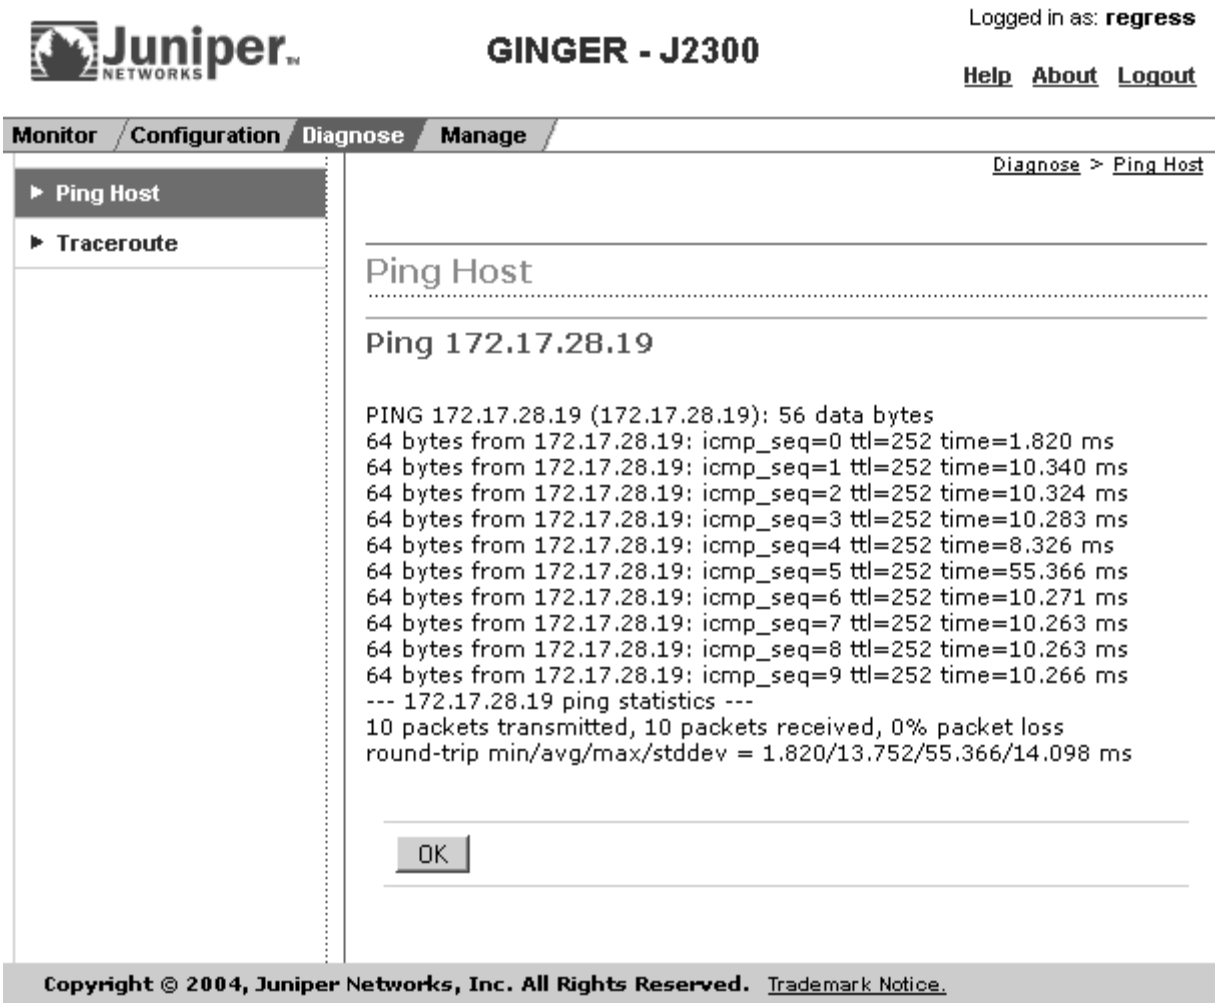

Table 82: J-Web Ping Host Results Summary

Field	Description
bytes bytes from *ip-address*	■ *bytes* —Size of ping response packet, which is equal to the value you entered in the Packet Size box, plus 8. ■ *ip-address* —IP address of destination host that sent the ping response packet.
icmp_seq=*number*	*number* —Sequence Number field of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
ttl=*number*	*number* —Time-to-live hop-count value of the ping response packet.
time=*time*	*time* —Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also known as *round-trip time* .
number packets transmitted	*number* —Number of ping requests (probes) sent to host.

Field	Description
number packets received	*number* —Number of ping responses received from host.
percentage packet loss	*percentage* —Number of ping responses divided by the number of ping requests, specified as a percentage.
round-trip min/avg/max/stddev = *min-time* / *avg-time* / *max-time* / *std-dev* ms	■ *min-time* —Minimum round-trip time (see time=*time* field in this table). ■ *avg-time* —Average round-trip time. ■ *max-time* —Maximum round-trip time. ■ *std-dev* —Standard deviation of the round-trip times.

If the Services Router does not receive ping responses from the destination host (the output shows a packet loss of 100 percent), one of the following might apply:

- The host is not operational.
- There are network connectivity problems between the Services Router and the host.
- The host might be configured to ignore ICMP echo requests.
- The host might be configured with a firewall filter that blocks ICMP echo requests or ICMP echo responses.
- The size of the ICMP echo request packet exceeds the MTU of a host along the path.
- The value you selected in the Time-to-Live box was less than the number of hops in the path to the host, in which case the host might reply with an ICMP error message.

For more information about ICMP, see RFC 792, *Internet Control Message Protocol*.

Using the J-Web Traceroute Tool

You can use the traceroute diagnostic tool to display a list of routers between the Services Router and a specified destination host. The output is useful for diagnosing a point of failure in the path from the Services Router to the destination host, and addressing network traffic latency and throughput problems.

The Services Router generates the list of routers by sending a series of ICMP traceroute packets in which the time-to-live (TTL) value in the messages sent to each successive router is incremented by 1. (The TTL value of the first traceroute packet is set to 1.) In this manner, each router along the path to the destination host replies with a Time Exceeded packet from which the source IP address can be obtained.

Alternatively, you can use the CLI traceroute command to generate the list. (See “Using the traceroute Command” on page 228.)

To use the traceroute tool:

1. Select **Diagnose > Traceroute**.
2. Next to Advanced options, click the expand icon (see Figure 53).
3. Enter information into the Traceroute page, as described in Table 83.

   The Remote Host field is the only required field.
4. Click **Start**.

   The results of the traceroute operation are displayed in the main pane. If no options are specified, each line of the traceroute display is in the following format:

   ```
   hop-number  host  ( ip-address ) [ as-number ] time1  time2  time3
   ```

   The Services Router sends a total of three traceroute packets to each router along the path and displays the round-trip time for each traceroute operation. If the Services Router times out before receiving a Time Exceeded message, an asterisk (*) is displayed for that round-trip time.

   Table 84 summarizes the output fields of the display.
5. To stop the traceroute operation before it is complete, click **OK** while the results of the traceroute operation are being displayed.

Figure 53: Traceroute Page

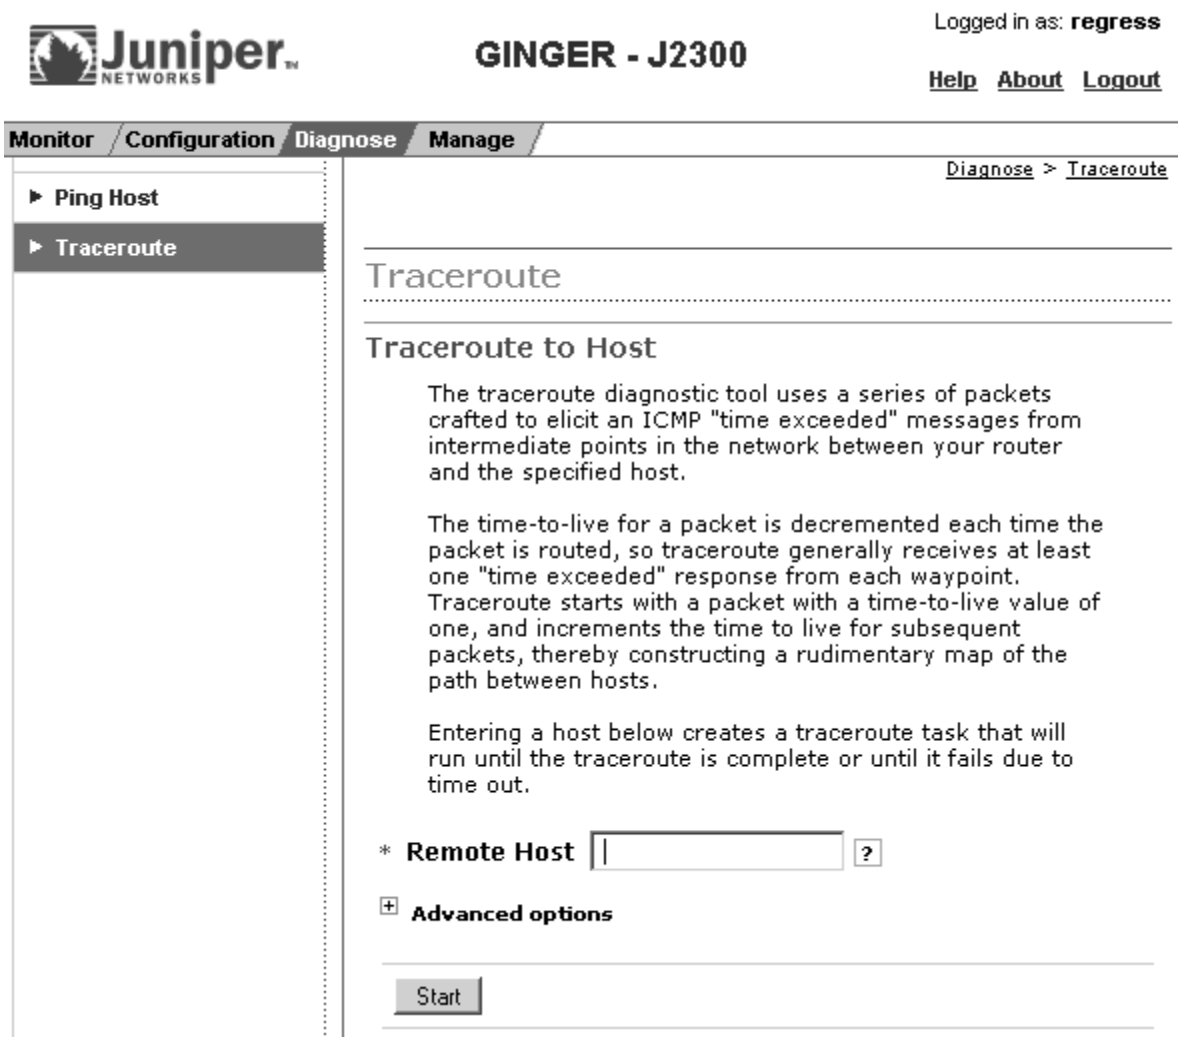

Table 83: Traceroute Summary

Field	Function	Your Action
Remote Host	Identifies the destination host of the traceroute.	Type the hostname or IP address of the destination host.
Advanced Options		
Don't Resolve Addresses	Determines whether hostnames of the hops along the path are displayed, in addition to IP addresses.	■ To suppress the display of the hop hostnames, select the check box. ■ To display the hop hostnames, clear the check box.
Gateway	Specifies the IP address of the gateway to route through.	Type the gateway IP address.

Field	Function	Your Action
Source Address	Specifies the source address of the outgoing traceroute packets.	Type the source IP address.
Bypass Routing	Determines whether traceroute packets are routed by means of the routing table. If the routing table is not used, traceroute packets are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, traceroute responses are not sent.	■ To bypass the routing table and send the traceroute packets to hosts on the specified interface only, select the check box. ■ To route the traceroute packets by means of the routing table, clear the check box.
Interface	Specifies the interface on which the traceroute packets are sent.	From the drop-down list, select the interface on which traceroute packets are sent. If you select **any**, the traceroute requests are sent on all interfaces.
Time-to-Live	Specifies the maximum time-to-live (TTL) hop count for the traceroute request packet.	From the drop-down list, select the TTL.
Type-of-Service	Specifies the type-of-service (TOS) value to include in the IP header of the traceroute request packet.	From the drop-down list, select the decimal value of the TOS field.
Resolve AS Numbers	Determines whether the autonomous system (AS) number of each intermediate hop between the router and the destination host is displayed.	■ To display the AS numbers, select the check box. ■ To suppress the display of the AS numbers, clear the check box.

Table 84: J-Web Traceroute Results Summary

Field	Description
hop-number	Number of the hop (router) along the path.
host	Hostname, if available, or IP address of the router. If the Don't Resolve Addresses check box is selected, the hostname is not displayed.
ip-address	IP address of the router.
as-number	AS number of the router.
time1	Round-trip time between the sending of the first traceroute packet and the receiving of the corresponding Time Exceeded packet from that particular router.
time2	Round-trip time between the sending of the second traceroute packet and the receiving of the corresponding Time Exceeded packet from that particular router.
time3	Round-trip time between the sending of the third traceroute packet and the receiving of the corresponding Time Exceeded packet from that particular router.

If the Services Router does not display the complete path to the destination host, one of the following might apply:

- The host is not operational.
- There are network connectivity problems between the Services Router and the host.
- The host, or a router along the path, might be configured to ignore ICMP traceroute messages.
- The host, or a router along the path, might be configured with a firewall filter that blocks ICMP traceroute requests or ICMP time exceeded responses.
- The value you selected in the Time Exceeded box was less than the number of hops in the path to the host. In this case, the host might reply with an ICMP error message.

For more information about ICMP, see RFC 792, *Internet Control Message Protocol*.

Using CLI Diagnostic Commands

This section describes how to use the CLI diagnostic tools. Because the CLI is a superset of the J-Web interface, you can perform certain tasks only through the CLI. For an overview of the CLI operational mode commands, along with instructions for filtering command output, see "CLI Diagnostic Commands Overview" on page 201.

This section contains the following topics:

Using the ping Command

Use the CLI ping command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The Services Router sends a series of ICMP echo (ping) requests to a specified host and expects to receive ICMP echo responses.

Alternatively, you can use the J-Web interface. (See "Using the J-Web Ping Host Tool" on page 218.)

Enter the ping command with the following syntax. Table 85 describes the ping command options.

```
user@host> ping host <interface source-interface> <bypass-routing>
<count number> <do-not-fragment> <inet> <interval seconds>
<loose-source [ hosts ]> <no-resolve> <pattern string> <rapid>
<record-route> <routing-instance routing-instance-name> <size bytes>
<source address> <strict> <strict-source [ hosts ]> <tos number>
<ttl number> <verbose> <wait seconds> <detail>
```

To quit the ping command, press Ctrl-C.

Table 85: CLI ping Command Options

Option	Description
host	Pings the hostname or IP address you specify.
interface *source-interface*	Sends the ping requests on the interface you specify. If you do not include this option, ping requests are sent on all interfaces.
bypass-routing	Bypasses the routing tables and send the ping requests only to hosts on directly attached interfaces. If the host is not on a directly attached interface, an error message is returned.
count *number*	Limits the number of ping requests to send. Specify a count from 0 through 1,000,000. If you do not specify a count, ping requests are continuously sent until you press Ctrl-C.
do-not-fragment	Sets the Don't Fragment (DF) bit in the IP header of the ping request packet.
inet	Forces the ping requests to an IPv4 destination.
interval *seconds*	Sets the interval between ping requests, in seconds. Specify an interval from 0.1 through 10,000. The default value is 1 second.
loose-source [*hosts*]	Sets the loose source routing option in the IP header of the ping request packet.
no-resolve	Suppresses the display of the hostnames of the hops along the path.
pattern *string*	Includes the hexadecimal string you specify, in the ping request packet.
rapid	Sends ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the count option.
record-route	Sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed on the screen.
routing-instance *routing-instance-name*	Uses the routing instance you specify for the ping request.
size *bytes*	Sets the size of the ping request packet. Specify a size from 0 through 65,468. The default value is 56 bytes, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.
source *address*	Uses the source address that you specify, in the ping request packet.
strict	Sets the strict source routing option in the IP header of the ping request packet.
strict-source [*hosts*]	Sets the strict source routing option in the IP header of the ping request packet, and uses the list of hosts you specify for routing the packet.
tos *number*	Sets the type-of-service (TOS) value in the IP header of the ping request packet. Specify a value from 0 through 255.
ttl *number*	Sets the time-to-live (TTL) value for the ping request packet. Specify a value from 0 through 255.

Option	Description
verbose	Displays detailed output.
wait *seconds*	Sets the maximum time to wait after sending the last ping request packet.
detail	Displays the interface on which the ping response was received.

Following is sample output from a ping command:

```
    user@host> ping host3 count 4

PING host3.site.net (176.26.232.111): 56 data bytes
64 bytes from 176.26.232.111: icmp_seq=0 ttl=122 time=0.661 ms
64 bytes from 176.26.232.111: icmp_seq=1 ttl=122 time=0.619 ms
64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms
64 bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms

--- host3.site.net ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

The fields in the display are the same as those displayed by the J-Web ping host diagnostic tool. Table 82 summarizes these output fields.

Using the traceroute Command

Use the CLI traceroute command to display a list of routers between the Services Router and a specified destination host. This command is useful for diagnosing a point of failure in the path from the Services Router to the destination host, and addressing network traffic latency and throughput problems.

The Services Router generates the list of routers by sending a series of ICMP traceroute packets in which the time-to-live (TTL) value in the messages sent to each successive router is incremented by 1. (The TTL value of the first traceroute packet is set to 1.) In this manner, each router along the path to the destination host replies with a Time Exceeded packet from which the source IP address can be obtained.

Alternatively, you can use the J-Web interface. (See "Using the J-Web Traceroute Tool" on page 222.)

Enter the traceroute command with the following syntax. Table 86 describes the traceroute command options.

```
    user@host> traceroute host <interface source-interface>
    <as-number-lookup> <bypass-routing> <gateway address>
    <inet> <logical-router logical-router-name> <no-resolve>
    <routing-instance routing-instance-name> <source address>
    <tos number> <ttl number> <wait seconds>
```

To quit the traceroute command, press Ctrl-C.

Table 86: CLI traceroute Command Options

Option	Description
host	Sends traceroute packets to the hostname or IP address you specify.
interface *source-interface*	Sends the traceroute packets on the interface you specify. If you do not include this option, traceroute packets are sent on all interfaces.
as-number-lookup	Displays the autonomous system (AS) number of each intermediate hop between the router and the destination host.
bypass-routing	Bypasses the routing tables and send the traceroute packets only to hosts on directly attached interfaces. If the host is not on a directly attached interface, an error message is returned.
gateway *address*	Uses the gateway you specify to route through.
logical-router *logical-router-name*	Sends traceroute packets to this logical router.
inet	Forces the traceroute packets to an IPv4 destination.
no-resolve	Suppresses the display of the hostnames of the hops along the path.
routing-instance *routing-instance-name*	Uses the routing instance you specify for the traceroute.
source *address*	Uses the source address you specify in the traceroute packet.
tos *number*	Sets the type-of-service (TOS) value in the IP header of the traceroute packet. Specify a value from 0 through 255.
ttl *number*	Sets the time-to-live (TTL) value for the traceroute packet. Specify a hop count from 0 through 255.
wait *seconds*	Sets the maximum time to wait for a response.

Following is sample output from a traceroute command:

```
user@host> traceroute host2

traceroute to 173.24.232.66 (172.24.230.41), 30 hops max, 40 byte packets
 1  173.18.42.253 (173.18.42.253)  0.482 ms  0.346 ms  0.318 ms
 2  host4.site1.net (173.18.253.5)  0.401 ms  0.435 ms  0.359 ms
 3  host5.site1.net (173.18.253.5)  0.401 ms  0.360 ms  0.357 ms
 4  173.24.232.65 (173.24.232.65)  0.420 ms  0.456 ms  0.378 ms
 5  173.24.232.66 (173.24.232.66)  0.830 ms  0.779 ms  0.834 ms
```

The fields in the display are the same as those displayed by the J-Web traceroute diagnostic tool. Table 84 summarizes these output fields.

Using the monitor interface Command

Use the CLI monitor interface command to display real-time traffic, error, alarm, and filter statistics about a physical or logical interface. Enter the command with the following syntax:

```
user@host> monitor interface (interface-name | traffic)
```

Replace *interface-name* with the name of a physical or logical interface. If you specify the traffic option, statistics for all active interfaces are displayed.

The real-time statistics are updated every second. The Current delta and Delta columns display the amount the statistics counters have changed since the monitor interface command was entered or since you cleared the delta counters. Table 87 and Table 88 list the keys you use to control the display using the *interface-name* and traffic options. (The keys are not case sensitive.)

Table 87: CLI monitor interface Output Control Keys

Key	Action
c	Clears (returns to 0) the delta counters in the Current delta column. The statistics counters are not cleared.
f	Freezes the display, halting the update of the statistics and delta counters.
i	Displays information about a different interface. You are prompted for the name of a specific interface.
n	Displays information about the next interface. The Services Router scrolls through the physical and logical interfaces in the same order in which they are displayed by the show interfaces terse command.
q or ESC	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

Table 88: CLI monitor interface Traffic Output Control Keys

Key	Action
b	Displays the statistics in units of bytes and bytes per second (bps).
c	Clears (returns to 0) the delta counters in the Delta column. The statistics counters are not cleared.
d	Displays the Delta column instead of the rate column—in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or ESC	Quits the command and returns to the command prompt.
r	Displays the rate column—in bps and pps—instead of the Delta column.

Following are sample displays from the monitor interface command:

```
user@host> monitor interface fe-0/0/0

host1                              Seconds: 11                 Time: 16:47:49
                                                               Delay: 0/0/0
Interface: fe-0/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:                                          Current delta
  Input bytes:                  381588589                          [11583]
  Output bytes:                   9707279                           [6542]
  Input packets:                  4064553                            [145]
  Output packets:                   66683                             [25]
Error statistics:
```

```
  Input errors:                    0                              [0]
  Input drops:                     0                              [0]
  Input framing errors:            0                              [0]
  Carrier transitions:             0                              [0]
  Output errors:                   0                              [0]
  Output drops:                    0                              [0]
```

NOTE: The output fields displayed when you enter the monitor interface *interface-name* command are determined by the interface you specify.

```
user@host> monitor interface traffic

Interface    Link  Input packets        (pps)     Output packets        (pps)
 fe-0/0/0      Up          42334          (5)              23306          (3)
 fe-0/0/1      Up      587525876      (12252)          589621478      (12891)
```

Using the monitor traffic Command

Use the CLI monitor traffic command to display packet headers transmitted through network interfaces.

Enter the monitor traffic command with the following syntax. Table 89 describes the monitor traffic command options.

```
user@host> monitor traffic <absolute-sequence> <count number>
<interface interface-name> <layer2-headers> <matching expression>
<no-domain-names> <no-promiscuous> <no-resolve> <no-timestamp>
<print-ascii> <print-hex> <size bytes> <brief | detail | extensive>
```

To quit the monitor traffic command and return to the command prompt, press Ctrl-C.

NOTE: Using the monitor traffic command can degrade Services Router performance. We recommend that you use filtering options—such as count and matching—to minimize the impact to packet throughput on the Services Router.

Table 89: CLI monitor traffic Command Options

Option	Description
absolute-sequence	Displays the absolute TCP sequence numbers.
count *number*	Displays the specified number of packet headers. Specify a value from 0 through 100,000. The command quits and exits to the command prompt after this number is reached.
interface *interface-name*	Displays packet headers for traffic on the specified interface. If an interface is not specified, the lowest numbered interface is monitored.
layer2-headers	Displays the link-layer packet header on each line.

Option	Description
matching *expression*	Displays packet headers that match an expression. Table 90 through Table 92 list match conditions, logical operators, and arithmetic, binary, and relational operators you can use in the expression.
no-domain-names	Suppresses the display of the domain name portion of the hostname.
no-promiscuous	Specifies *not* to place the monitored interface in promiscuous mode. In promiscuous mode, the interface reads every packet that reaches it. In non-promiscuous mode, the interface reads only the packets addressed to it.
no-resolve	Suppresses the display of hostnames.
no-timestamp	Suppresses the display of packet header timestamps.
print-ascii	Displays each packet header in ASCII format.
print-hex	Displays each packet header, except link-layer headers, in hexadecimal format.
size *bytes*	Displays the number of bytes for each packet that you specify. If a packet header exceeds this size, the displayed packet header is truncated. The default value is 96.
brief	Displays minimum packet header information. This is the default.
detail	Displays packet header information in moderate detail. For some protocols, you must also use the size option to see detailed information.
extensive	Displays the most extensive level of packet header information. For some protocols, you must also use the size option to see extensive information.

To limit the packet header information displayed by the monitor traffic command, include the matching *expression* option. An expression consists of one or more match conditions listed in Table 90, enclosed in quotation marks (“ “). You can combine match conditions by using the logical operators listed in Table 91 (shown in order of highest to lowest precedence).

For example, to display TCP or UDP packet headers, enter the following command:

```
user@host> monitor traffic matching “tcp || udp”
```

To compare the following types of expressions, use the relational operators listed in Table 92 (listed from highest to lowest precedence):

- Arithmetic—Expressions that use the arithmetic operators listed in Table 92.
- Binary—Expressions that use the binary operators listed in Table 92.
- Packet data accessor—Expressions that use the following syntax:

```
protocol [byte-offset <size>]
```

Replace *protocol* with any protocol in Table 90. Replace *byte-offset* with the byte offset, from the beginning of the packet header, to use for the comparison. The optional *size* parameter represents the number of bytes examined in the packet header—1, 2, or 4 bytes.

For example, the following command displays all multicast traffic:

```
user@host> monitor traffic matching “ether[0] & 1 !=0”
```

Table 90: CLI monitor traffic Match Conditions

Match Condition	Description
Entity Type	
host [*address* | *hostname*]	Matches packet headers that contain the specified address or hostname. You can preprend any of the following protocol match conditions, followed by a space, to host: arp, ip, rarp, or any of the Directional match conditions.
network *address*	Matches packet headers with source or destination addresses containing the specified network address.
network *address* mask *mask*	Matches packet headers containing the specified network address and subnet mask.
port [*port-number* | *port-name*]	Matches packet headers containing the specified source or destination TCP or UDP port number or port name.
Directional	Directional match conditions can be prepended to any Entity Type match conditions, followed by a space.
destination	Matches packet headers containing the specified destination.
source	Matches packet headers containing the specified source.
source and destination	Matches packet headers containing the specified source *and* destination.
source or destination	Matches packet headers containing the specified source *or* destination.
Packet Length	
less *bytes*	Matches packets with lengths less than or equal to the specified value, in bytes.
greater *bytes*	Matches packets with lengths greater than or equal to the specified value, in bytes.
Protocol	
arp	Matches all ARP packets.
ether	Matches all Ethernet frames.
ether [broadcast | multicast]	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with source or destination.
ether protocol [*address* | (\arp | \ip | \rarp)	Matches Ethernet frames with the specified address or protocol type. The arguments arp, ip, and rarp are also independent match conditions, so they must be preceded with a backslash (\) when used in the ether protocol match condition.
icmp	Matches all ICMP packets.
ip	Matches all IP packets.
ip [broadcast | multicast]	Matches broadcast or multicast IP packets.

Match Condition	Description
ip protocol [*address* | (\icmp | igrp | \tcp | \udp)]	Matches IP packets with the specified address or protocol type. The arguments icmp, tcp, and udp are also independent match conditions, so they must be preceded with a backslash (\) when used in the ip protocol match condition.
isis	Matches all IS-IS routing messages.
rarp	Matches all RARP packets.
tcp	Matches all TCP packets.
udp	Matches all UDP packets.

Table 91: CLI monitor traffic Logical Operators

Logical Operator	Description
!	Logical NOT.
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.
||	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
()	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

Table 92: CLI monitor traffic Arithmetic, Binary, and Relational Operators

Operator	Description
Arithmetic Operator	
+	Addition operator.
–	Subtraction operator.
/	Division operator.
Binary Operator	
&	Bitwise AND.
*	Bitwise exclusive OR.
|	Bitwise inclusive OR.
Relational Operator	
<=	A match occurs if the first expression is less than or equal to the second.
>=	A match occurs if the first expression is greater than or equal to the second.
<	A match occurs if the first expression is less than the second.
>	A match occurs if the first expression is greater than the second.

Operator	Description
=	A match occurs if the first expression is equal to the second.
!=	A match occurs if the first expression is not equal to the second.

Following is sample output from the monitor traffic command:

```
user@host> monitor traffic count 4 matching “arp” detail

Listening on fe-0/0/0, capture size 96 bytes

15:04:16.276780  In arp who-has 193.1.1.1 tell host1.site2.net
15:04:16.376848  In arp who-has host2.site2.net tell host1.site2.net
15:04:16.376887  In arp who-has 193.1.1.2 tell host1.site2.net
15:04:16.601923  In arp who-has 193.1.1.3 tell host1.site2.net
```

Using the monitor file Command

You can enter the monitor file command to display real-time additions to files such as system logs and trace files:

```
user@host> monitor start filename
```

When the Services Router adds a record to the file specified by *filename*, the record is displayed on the screen. For example, if you have configured a system log file named system-log (by including the syslog statement at the [edit system] hierarchy level), you can enter the monitor start system-log command to display the records added to the system log.

To display a list of files that are being monitored, enter the monitor list command. To stop the display of records for a specified file, enter the monitor stop *filename* command.

Using mtrace Commands

You can use CLI mtrace commands to trace information about multicast paths. This section covers the following mtrace commands:

- mtrace from-source—Displays information about a multicast path from a source to a receiver. See “Using the mtrace from-source Command” on page 236.
- mtrace monitor—Monitors and displays multicast trace operations. See “Using the mtrace monitor Command” on page 238.

For more information about the mtrace commands, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Using the mtrace from-source Command

To display information about a multicast path from a source to a receiver, enter the mtrace from-source command with the following syntax. Table 93 describes the mtrace from-source command options.

```
user@host> mtrace from-source source host <<extra-hops number>
| <group address> | <interval seconds> | <max-hops number>
| <max-queries number> | <response host> | <ttl number> |
<wait-time seconds>> <loop> <multicast-response | unicast-response>
<no-resolve> <no-router-alert> <brief | detail>
```

Table 93: CLI mtrace from-source Command Options

Option	Description
source *host*	Traces the path to the specified hostname or IP address.
extra-hops *number*	Sets the number of extra hops to trace past nonresponsive routers. Specify a value from 0 through 255.
group *address*	Traces the path for the specified group address. The default value is 0.0.0.0.
interval *seconds*	Sets the interval between statistics gathering. The default value is 10.
max-hops *number*	Sets the maximum number of hops to trace toward the source. Specify a value from 0 through 255. The default value is 32.
max-queries *number*	Sets the maximum number of queries for any hop. Specify a value from 1 through 32. The default value is 3.
response *host*	Sends the response packets to the specified hostname or IP address. By default, the response packets are sent to the router that sent the requests.
ttl *number*	Sets the time-to-live (TTL) value in the IP header of the query packets. Specify a hop count from 0 through 255. The default value for local queries to the *all routers* multicast group is 1. Otherwise, the default value is 127.
wait-time *seconds*	Sets the time to wait for a response packet. The default value is 3 seconds.
loop	Loops indefinitely, displaying rate and loss statistics. To quit the mtrace command, press Ctrl-C.
multicast-response	Forces the responses to use multicast.
unicast-response	Forces the response packets to use unicast.
no-resolve	Does not display hostnames.
no-router-alert	Does not use the router alert IP option in the IP header.
brief	Does not display packet rates and losses.
detail	Displays packet rates and losses if a group address is specified.

Following is sample output from the mtrace from-source command:

```
user@host> mtrace from-source source 192.1.4.1 group 224.1.1.1

Mtrace from 192.1.4.1 to 192.1.30.2 via group 224.1.1.1
Querying full reverse path... * *
  0  ? (192.1.30.2)
```

```
 -1  ? (192.1.30.1)  PIM  thresh^ 1
 -2  routerC.mycompany.net (192.1.40.2)  PIM  thresh^ 1
 -3  hostA.mycompany.net (192.1.4.1)
Round trip time 22 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

  Source        Response Dest    Overall     Packet Statistics For Traffic From
192.1.4.1  192.1.30.2       Packet      192.1.4.1 To 224.1.1.1
     v       __/  rtt   16 ms     Rate       Lost/Sent = Pct  Rate
192.168.195.37
192.1.40.2      routerC.mycompany.net
     v     ^      ttl    2                      0/0    = --    0 pps
192.1.40.1
192.1.30.1      ?
     v      \__   ttl    3                      ?/0            0 pps
192.1.30.2      192.1.30.2
  Receiver      Query Source
```

Each line of the trace display is usually in the following format (depending on the options selected and the responses from the routers along the path):

hop-number host (*ip-address*) *protocol ttl*

Table 94 summarizes the output fields of the display.

NOTE: The packet statistics gathered from Juniper Networks routers and routing nodes are always displayed as 0.

Table 94: CLI mtrace from-source Command Display Summary

Field	Description
hop-number	Number of the hop (router) along the path.
host	Hostname, if available, or IP address of the router. If the no-resolve option was entered in the command, the hostname is not displayed.
ip-address	IP address of the router.
protocol	Protocol used.
ttl	TTL threshold.
Round trip time *milliseconds* ms	Total time between the sending of the query packet and the receiving of the response packet.
total ttl of *number* required	Total number of hops required to reach the source.
Source	Source IP address of the response packet.
Response Dest	Response destination IP address.
Overall	Average packet rate for all traffic at each hop.
Packet Statistics For Traffic From	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.

Field	Description
Receiver	IP address receiving the multicast packets.
Query Source	IP address of the host sending the query packets.

Using the mtrace monitor Command

To monitor and display multicast trace operations, enter the mtrace monitor command:

```
user@host> mtrace monitor

Mtrace query at Apr 21 16:00:54 by 192.1.30.2, resp to 224.0.1.32, qid 2a83aa
packet from 192.1.30.2 to 224.0.0.2
from 192.1.30.2 to 192.1.4.1 via group 224.1.1.1 (mxhop=60)

Mtrace query at Apr 21 16:00:57 by 192.1.30.2, resp to 224.0.1.32, qid 25dc17
packet from 192.1.30.2 to 224.0.0.2
from 192.1.30.2 to 192.1.4.1 via group 224.1.1.1 (mxhop=60)

Mtrace query at Apr 21 16:01:00 by 192.1.30.2, resp to same, qid 20e046
packet from 192.1.30.2 to 224.0.0.2
from 192.1.30.2 to 192.1.4.1 via group 224.1.1.1 (mxhop=60)

Mtrace query at Apr 21 16:01:10 by 192.1.30.2, resp to same, qid 1d25ad
packet from 192.1.30.2 to 224.0.0.2
from 192.1.30.2 to 192.1.4.1 via group 224.1.1.1 (mxhop=60)
```

This example displays only mtrace queries. When the Services Router captures an mtrace response, the display is similar, but the complete mtrace response is also displayed—exactly as it is displayed in mtrace from-source command output.

Table 95 summarizes the output fields of the display.

Table 95: CLI mtrace monitor Command Display Summary

Field	Description
Mtrace *operation-type* at *time-of-day*	■ *operation-type*—Type of multicast trace operation: query or response. ■ *time-of-day*—Date and time the multicast trace query or response was captured.
by	IP address of the host issuing the query.
resp to *address*	*address*—Response destination address.
qid *qid*	*qid*—Query ID number.
packet from *source* to *destination*	■ *source*—IP address of the source of the query or response. ■ *destination*—IP address of the destination of the query or response.
from *source* to *destination*	■ *source*—IP address of the multicast source. ■ *destination*—IP address of the multicast destination.

Field	Description
via group *address*	*address*—Group address being traced.
mxhop=*number*	*number*—Maximum hop setting.

Chapter 11

Configuring SNMP for Network Management

The Simple Network Management Protocol (SNMP) is a client/server standard that helps you diagnose and monitor network health and statistics.

You can use either J-Web Quick Configuration or a configuration editor to configure SNMP.

This chapter contains the following topics. For more information about SNMP, see the *JUNOS Network Management Configuration Guide*.

Network Management Overview

A network is a complex organization of nodes and processes that must operate reliably and efficiently. Having a single node or link failure in a network can undermine the network's performance and result in a loss of service. Therefore, determining where and when a network failure is occurring is a necessity.

Additionally, gathering statistics about how a network is performing can help you diagnose the overall health of the network and pinpoint bottlenecks so that you can address network growth appropriately.

By querying individual network nodes and receiving triggered updates, SNMP clients are able to provide valuable feedback about the state of a network.

Managers and Agents

Because SNMP is a client/server protocol, SNMP nodes can be classified as either clients (SNMP managers) or servers (SNMP agents).

SNMP managers, also known as network management systems (NMSs), occupy central points in the network and they actively query and collect messages from SNMP agents in the network. SNMP agents are individual processes running on network nodes that gather information for a particular node and transfer the information to SNMP managers as queries are processed. Because SNMP agents are individual SNMP processes running on a host, multiple agents can be active on a single network node at any given time.

SMI, MIBs, and OIDs

Agents store information in a hierarchical database called the Structure of Management Information (SMI). The SMI resembles a file system; information is stored in individual files that are hierarchically arranged in the database. The individual files that store the information are known as Management Information Bases (MIBs). Each MIB contains nodes of information that are stored in a tree structure. Information branches down from a root node to individual leaves in the tree, and the individual leaves comprise the information that is queried by managers for a given MIB. The nodes of information are identified by an object ID (OID). The OID is a dotted integer identifier (1.3.6.1.2.1.2, for instance) or a subtree name (such as interfaces) that corresponds to an indivisible piece of information in the MIB.

Standard and Enterprise MIBs

A set of MIBs has been defined by the IETF and documented in various RFCs. These MIBs are common across many platforms. Additionally, individual enterprises can create their own set of enterprise-specific MIBs, provided they share the same structure as the standard MIBs. This structure is enforced through the Abstract Syntax Notation (ASN), which is a definition language used to store information.

SNMP Requests

Information is stored in MIBs, and MIBs are queried by SNMP managers. Managers send SNMP requests to process the information. SNMP requests come in two primary forms: get requests and set requests. These requests are processed by one or more agents on a particular node, and information is retrieved or modified on the MIB. When the agent has processed the request, it generates an SNMP response that either returns retrieved information from the MIB or acknowledges that information has been modified on the MIB.

SNMP Communities

To help ensure that only specific SNMP managers can access a particular SNMP agent, SNMP access is granted through communities. To control access, you first create an SNMP community. The community is assigned a name that is unique on the host. All SNMP requests that are sent to the agent must be configured with the same community name.

When multiple agents are configured on a particular host, the community name process ensures that SNMP requests are sorted to only those agents configured to handle the requests.

Additionally, communities allow you to specify one or more addresses or address prefixes to which you want to either allow or deny access. By specifying a list of clients, you can control exactly which SNMP managers have access to a particular agent.

SNMP Traps

The get and set commands that SNMP uses are useful for querying hosts within a network. However, the commands do not provide a means by which events can trigger a notification. For instance, if a link fails, the health of the link is unknown until an SNMP manager next queries that agent.

SNMP has traps, which are unsolicited notifications that are triggered by events on the host. When you configure a trap, you specify the types of events that can trigger trap messages, and you configure a set of targets to receive the generated messages.

Before You Begin

Before you begin configuring SNMP, complete the following tasks:

- Establish basic connectivity. See “Establishing Basic Connectivity” on page 47.
- Configure network interfaces. See “Configuring Network Interfaces” on page 79.

Configuring SNMP with Quick Configuration

J-Web Quick Configuration allows you to define system identification information, create SNMP communities, and create SNMP trap groups. Figure 54 shows the Quick Configuration page for SNMP.

Figure 54: Quick Configuration Page for SNMP

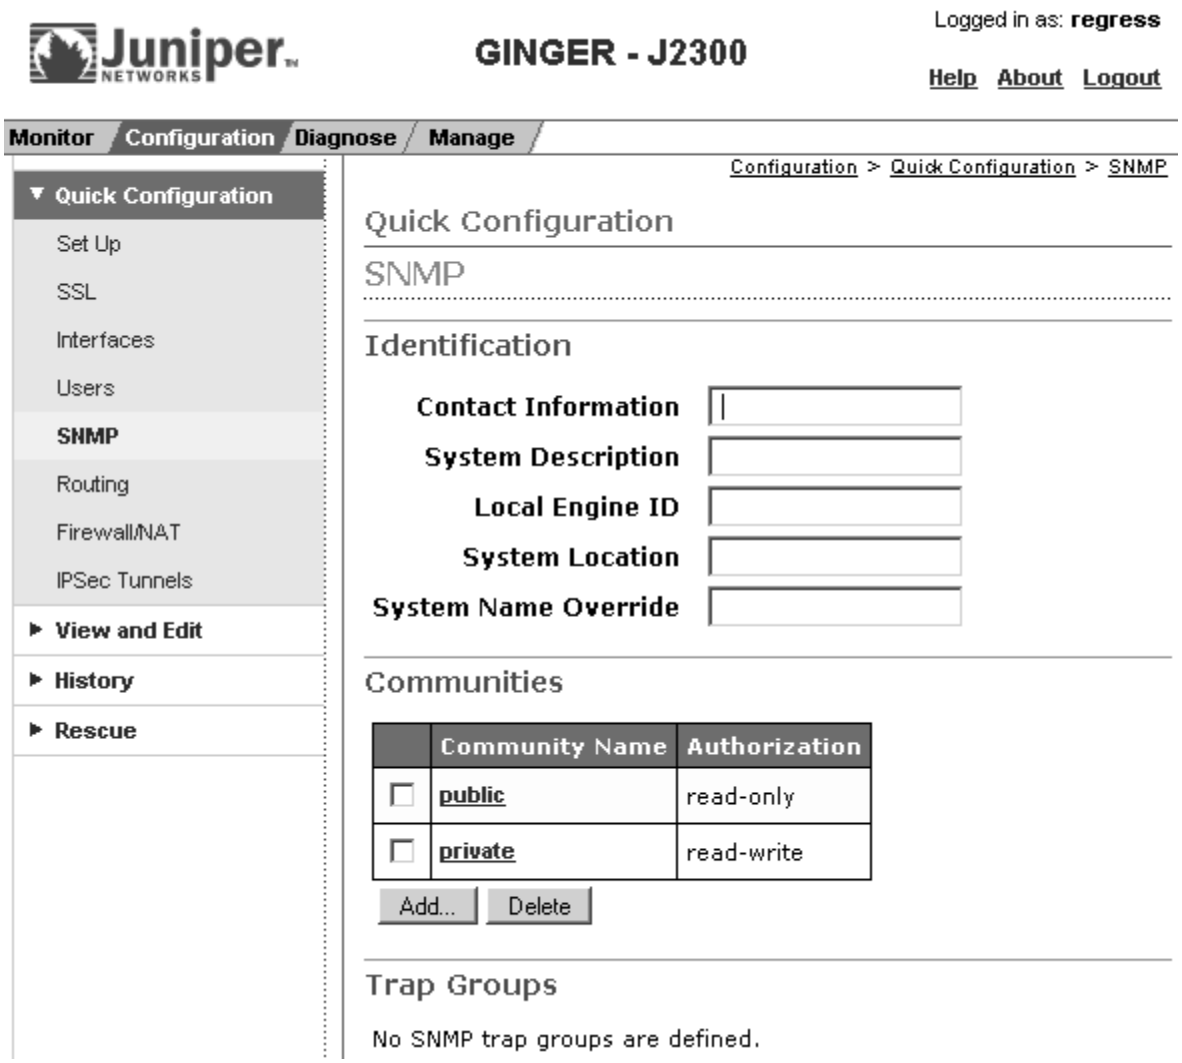

To configure SNMP features with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > SNMP**.
2. Enter information into the Quick Configuration page for SNMP, as described in Table 96.
3. From the SNMP Quick Configuration page, click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration page for SNMP, click **Apply**.

- To apply the configuration and return to the Quick Configuration SNMP page, click **OK**.
- To cancel your entries and return to the Quick Configuration for SNMP page, click **Cancel**.

4. To check the configuration, see "Verifying the SNMP Configuration" on page 251.

Table 96: SNMP Quick Configuration Summary

Field	Function	Your Action
Identification		
Contact Information	Free-form text string that specifies an administrative contact for the system.	Type any contact information for the administrator of the system (such as name and phone number).
System Description	Free-form text string that specifies a description for the system.	Type any system information that describes the system (*J4300 with 4 PIMs*, for example).
Local Engine ID	Provides an administratively unique identifier of an SNMPv3 engine for system identification. The local engine ID contains a prefix and a suffix. The prefix is formatted according to specifications defined in RFC 3411. The suffix is defined by the local engine ID. Generally, the local engine ID suffix is the MAC address of fe-0/0/0.	Type the MAC address of the fe-0/0/0 interface.
System Location	Free-form text string that specifies the location of the system.	Type any location information for the system (lab name or rack name, for example).
System Name Override	Free-form text string that overrides the system hostname defined in "Establishing Basic Connectivity" on page 47.	Type the name of the system.
Communities		Click **Add**.
Community Name	Specifies the name of the SNMP community.	Type the name of the community being added.
Authorization	Specifies the type of authorization (either read-only or read-write) for the SNMP community being configured.	Select the desired authorization (either read-only or read-write) from the drop-down menu.
Traps		Click **Add**.
Trap Group Name	Specifies the name of the SNMP trap group being configured.	Type the name of the SNMP trap group being configured.

Field	Function	Your Action
Categories	Specifies which trap categories are added to the trap group being configured.	■ To generate traps for authentication failures, select **Authentication**. ■ To generate traps for chassis and environment notifications, select **Chassis**. ■ To generate traps for configuration changes, select **Configuration**. ■ To generate traps for link-related notifications (up-down transitions), select **Link**. ■ To generate traps for remote operation notifications, select **Remote operations**. ■ To generate traps for remote network monitoring (RMON), select **RMON alarm**. ■ To generate traps for routing protocol notifications, select **Routing**. ■ To generate traps on system warm and cold starts, select **Startup**. ■ To generate traps on Virtual Router Redundancy Protocol (VRRP) events (such as new-master or authentication failures), select **VRRP events**.
Targets	One or more hostnames or IP addresses that specify the systems to receive SNMP traps generated by the trap group being configured.	1. Enter the hostname or IP address, in dotted decimal notation, of the target system to receive the SNMP traps. 2. Click **Add**.

Configuring SNMP with a Configuration Editor

To configure SNMP on a Services Router, you must perform the following tasks marked *(Required):*

- (Required) “Defining System Identification Information” on page 247
- (Required) “Configuring SNMP Agents and Communities” on page 248
- (Required) “Managing SNMP Trap Groups” on page 249
- (Optional) “Controlling Access to MIBs” on page 250

For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 127.

Defining System Identification Information

Basic system identification information for a Services Router can be configured with SNMP and stored in various MIBs. This information can be accessed through SNMP requests and either queried or reset. Table 97 identifies types of basic system identification and the MIB into which it is stored.

Table 97: System Identification Information and Corresponding MIBs

System Information	MIB
Contact	sysContact
System location	sysLocation
System description	sysDescription
System name override	sysName

To configure basic system identification for SNMP:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure basic system information using SNMP, perform the configuration tasks described in Table 98.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying the SNMP Configuration” on page 251.

Table 98: Configuring Basic System Identification

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **SNMP** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Snmp**.	From the top of the configuration hierarchy, enter edit snmp
Configure the system contact information (such as a name and phone number).	In the Contact box, type the contact information as a free-form text string.	Set the contact information: set contact *" contact-information "*
Configure the system location information (such as a lab name and a rack name).	In the Location box, type the location information as a free-form text string.	Set the location information: set location *" location-information "*
Configure the system description (*J4300 with 4 PIMs*, for example).	In the Description box, type the description information as a free-form text string.	Set the description information: set description *" description-information "*
Configure a system name to override the system hostname defined in "Establishing Basic Connectivity" on page 47.	In the System Name box, type the system name as a free-form text string.	Set the system name: set name *name*
Configure the local engine ID to use the MAC address of fe-0/0/0 as the engine ID suffix.	1. Select **Engine id**. 2. In the Engine id choice box, select **Use mac address** from the drop-down menu. 3. Click **OK**.	Set the engine ID to use the MAC address: set engine-id use-mac-address

Configuring SNMP Agents and Communities

To configure the SNMP agent, you must enable and authorize the network management system access to the Services Router, by configuring one or more communities. Each community has a community name, an authorization, which determines the kind of access the network management system has to the router, and, when applicable, a list of valid clients that can access the router.

To configure SNMP communities:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure SNMP communities, perform the configuration tasks described in Table 99.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see "Verifying the SNMP Configuration" on page 251.

Table 99: Configuring SNMP Agents and Communities

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **SNMP** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Snmp**.	From the top of the configuration hierarchy, enter edit snmp
Create and name a community.	1. Next to Community, click **Add new entry**. 2. In the Community box, type the name of the community as a free-form text string.	Create a community: set community *community-name*
Grant read-write access to the community.	In the Authorization box, select **read-write** from the drop-down menu.	Set the authorization to read-write: set community *community-name* authorization read-write
Allow community access to a client at a particular IP address—for example, at IP address 10.10.10.10.	1. Next to Clients, click **Add new entry**. 2. In the Prefix box, type the IP address, in dotted decimal notation. 3. Click **OK**.	Configure client access for the IP address 10.10.10.10: set community *community-name* clients 10.10.10.10
Allow community access to a group of clients—for example, all addresses within the 10.10.10.0/24 prefix, except those within the 10.10.10.10/29 prefix.	1. Next to Clients, click **Add new entry**. 2. In the Prefix box, type the IP address prefix 10.10.10.0/24, and click **OK**. 3. Next to Clients, click **Add new entry**. 4. In the Prefix box, type the IP address prefix 10.10.10.10/29. 5. Select the **Restrict** check box. 6. Click **OK**.	1. Configure client access for the IP address 10.10.10.0/24: set community *community-name* clients 10.10.10.0/24 2. Configure client access to restrict the IP addresses 10.10.10.10/29: set community *community-name* clients 10.10.10.10/29 restrict

Managing SNMP Trap Groups

SNMP traps are unsolicited notifications that are generated by conditions on the Services Router. When events trigger a trap, a notification is sent to the configured clients for that particular trap group. To manage a trap group, you must create the group, specify the types of traps that are included in the group, and define one or more targets to receive the trap notifications.

To configure SNMP trap groups:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure SNMP trap groups, perform the configuration tasks described in Table 100.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see "Verifying the SNMP Configuration" on page 251.

Table 100: Configuring SNMP Trap Groups

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **SNMP** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Snmp**.	From the top of the configuration hierarchy, enter edit snmp
Create a trap group.	1. Next to Trap group, click **Add new entry**. 2. In the Group name box, type the name of the group as a free-form text string.	Create a community: set trap-group *trap-group-name*
Configure the trap group to send all trap notifications to a target IP address—for example, to the IP address 192.174.6.6.	1. Next to Targets, click **Add new entry**. 2. In the Target box, type the IP address 192.174.6.6, and click **OK**.	Set the trap-group target to 192.174.6.6: set trap-group *trap-group-name* target 192.174.6.6
Configure the trap group to generate SNMP notifications on authentication failures, environment alarms, and changes in link state for any of the interfaces.	1. Click **Categories**. 2. Select the **Authentication**, **Chassis**, and **Link** check boxes. 3. Click **OK**.	Configure the trap group categories: set trap-group *trap-group-name* categories authentication chassis link

Controlling Access to MIBs

By default, an SNMP community is granted access to all MIBs. To control the MIBs to which a particular community has access, configure SNMP views that include the MIBs you want to explicitly grant or deny access to.

To configure SNMP views:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure SNMP views, perform the configuration tasks described in Table 101.

3. If you are finished configuring the network, commit the configuration.

4. To check the configuration, see "Verifying the SNMP Configuration" on page 251.

Table 101: Configuring SNMP Views

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **SNMP** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Snmp**.	From the top of the configuration hierarchy, enter: edit snmp
Create a view.	1. Next to View, click **Add new entry**. 2. In the Name box, type the name of the view as a free-form text string.	Create a view: set view *view-name*
Configure the view to include a MIB—for example, pingMIB.	1. Next to Oid, click **Add new entry**. 2. In the Name box, type the OID of the pingMIB, in either dotted integer or subtree name format. 3. In the View action box, select **include** from the drop-down menu, and click **OK**.	Set the pingMIB OID value and mark it for inclusion: set view *view-name* oid 1.3.6.1.2.1.80 include
Configure the view to exclude a MIB—for example, jnxPingMIB.	1. Next to Oid, click **Add new entry**. 2. In the Name box, type the OID of the jnxPingMIB, in either dotted integer or subtree name format. 3. In the View action box, select **exclude** from the drop-down menu, and click **OK** twice.	Set the jnxPingMIB OID value and mark it for exclusion: set view *view-name* oid jnxPingMIB exclude
Associate the view with a community.	1. On the Snmp page, under Community, click the name of the community to which you want to apply the view. 2. In the View box, type the view name. 3. Click **OK**.	Set the community view: set community *community-name* view *view-name*

Verifying the SNMP Configuration

To verify the SNMP configuration, perform the following verification task.

Verifying SNMP Agent Configuration

Purpose Verify that SNMP is running and that requests and traps are being properly transmitted.

Action From the CLI, enter the show snmp statistics command.

Sample Output

```
user@host> show snmp statistics

SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too bigs: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0
```

What It Means The output shows a list of the SNMP statistics, including details about the number and types of packets transmitted. Verify the following information:

- The number of requests and traps is increasing as expected with the SNMP client configuration.
- Under Bad community names, the number of bad (invalid) communities is not increasing. A sharp increase in the number of invalid community names generally means that one or more community strings are configured incorrectly.

For more information about show snmp statistics, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Part 5

Configuring Routing Protocols

Chapter 12
Routing Overview

At its most fundamental level, routing is the process of delivering a message across a network or networks. This task is divided into two primary components: the exchange of routing information to accurately forward packets from source to destination and the packet-forwarding process.

To use the routing capabilities of a J-series Services Router, you must understand the fundamentals of IP routing and the routing protocols that are primarily responsible for the transmission of unicast traffic. To read this chapter, you need a basic understanding of IP addressing and TCP/IP.

This chapter includes the following topics. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Routing Terms

To understand routing, become familiar with the terms defined in Table 102 .

Table 102: Routing Terms

Term	Definition
adjacency	Portion of the local routing information that pertains to the reachability of a single neighbor over a single circuit or interface.
area	Administrative group of OSPF networks within an autonomous system (AS) that operates independently from other areas in the AS. Multiple areas within an AS reduce the amount of link-state advertisement (LSA) traffic on the network and the size of topology databases.
area border router (ABR)	In OSPF, a router having interfaces in multiple areas of an autonomous system (AS) so that it can link the areas to each other. An area border router maintains a separate topological database for each area it is connected to and shares topology information between areas.

Term	Definition
AS path	In BGP, the list of autonomous system (ASs) that a packet must traverse to reach a given set of destinations within a single AS.
autonomous system (AS)	Network or collection of routers under a single administrative authority.
backbone area	In OSPF, the central area in an autonomous system (AS) to which all other areas are connected by area border routers (ABRs). The backbone area always has the area ID 0.0.0.0.
bidirectional connectivity	Ability of directly connected devices to communicate with each other over the same link.
Border Gateway Protocol (BGP)	Exterior gateway protocol used to exchange routing information among routers in different autonomous systems.
broadcast	Operation of sending network traffic from one network node to all other network nodes.
cluster	In BGP, a set of routers that have been grouped together. A cluster consists of one system that acts as a route reflector, along with any number of client peers. The client peers receive their route information only from the route reflector system. Routers in a cluster do not need to be fully meshed.
confederation	In BGP, a group of autonomous systems (ASs) that appears to external ASs to be a single AS.
confederation sequence	Ordered set of autonomous systems (ASs) for a confederation. The closest AS in the path is first in the sequence.
convergence	After a topology change, the time all the routers in a network take to receive the information and update their routing tables.
cost	Unitless number assigned to a path between neighbors, based on throughput, round-trip time, and reliability. The sum of path costs between source and destination hosts determines the overall path cost. OSPF uses the lowest cost to determine the best path.
designated router (DR)	In OSPF, a node designated to process link-state advertisements (LSAs) and distribute topology updates for an autonomous system (AS).
distance vector	Number of hops to a routing destination.
dynamic routing	Routing method that enables the route of a message through a network to change as network conditions change. Compare *static routing*.
exterior gateway protocol (EGP)	Protocol that exchanges routing information between autonomous systems (ASs). BGP is an EGP. Compare *interior gateway protocol (IGP)*.
external BGP (EBGP)	BGP configuration in which sessions are established between routers in different autonomous systems (ASs).
external peer	In BGP, a peer that resides in a different autonomous system (AS) from the Services Router.
external route	Route to an area outside the network.
flooding	Technique by which a router forwards traffic to every node attached to the router, except the node from which the traffic arrived. Flooding is a simple but sometimes inefficient way to distribute routing information quickly to every node in a network. RIP and OSPF are flooding protocols, but BGP is not.
forwarding table	JUNOS software forwarding information base (FIB). The JUNOS routing protocol process installs active routes from its routing tables into the Routing Engine forwarding table. The kernel copies this forwarding table into the Packet Forwarding Engine, which is responsible for determining which interface transmits the packets.
full mesh	Network in which devices are organized in a mesh topology, with each node connected to every other network node.
gateway router	Node on a network that serves as an entrance to another network.
global AS	Global autonomous system (AS). An AS consisting of multiple subautonomous systems (sub-ASs).

Term	Definition
handshake	Process of exchanging signaling information between two communications devices to establish the method and transmission speed of a connection.
hello packet	In OSPF, a packet sent periodically by a router to first establish and then maintain network adjacency, and to discover neighbor routers.
hold time	Maximum number of seconds allowed to elapse between the time a BGP system receives successive keepalive or update messages from a peer.
hop	Trip a data packet takes from one router to another in the network. The number of routers through which a packet passes to get from its source to its destination is known as the hop count. In general, the best route is the one with the shortest hop count.
Intermediate System-to-Intermediate System (IS-IS)	Link-state, interior gateway routing protocol for IP networks that also uses the shortest-path-first (SPF) algorithm to determine routes.
interior gateway protocol (IGP)	Protocol that exchanges routing information within autonomous systems (ASs). IS-IS, OSPF, and RIP are IGPs. Compare *exterior gateway protocol (EGP)*.
Internal BGP (IBGP)	BGP configuration in which sessions are established between routers in the same autonomous systems (ASs).
internal peer	In BGP, a peer that resides in the same autonomous system (AS) as the Services Router.
keepalive message	Periodic message sent by one BGP peer to another to verify that the session between them is still active.
latency	Delay.
link-state advertisement (LSA)	Messages that announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces (neighbors). The exchange of LSAs establishes bidirectional connectivity between neighbors.
local preference	Optional BGP path attribute carried in internal BGP update packets that indicates the degree of preference for an external route.
mesh	Network topology in which devices are organized in a manageable, segmented manner with many, often redundant, interconnections between network nodes. See also *full mesh*.
metric	Numerical value that determines how quickly a packet can reach its destination. See also *cost*.
multiple exit discriminator (MED)	Optional BGP path attribute consisting of a metric value that is used to determine the exit point to a destination when all other factors in determining the exit point are equal.
neighbor	Adjacent router interface. A node can directly route packets to its neighbors only. See also *peer*.
network	Series of nodes interconnected by communication paths.
network diameter	Maximum hop count in a network.
network topology	Arrangement of nodes and connections in a network.
node	Connection point that operates as a redistribution point or an end point in a network, recognizing data transmissions and either forwarding or processing them.
notification message	Message sent between BGP peers to inform the receiving peer that the sending peer is terminating the session because an error occurred, and explaining the error.
not-so-stubby area (NSSA)	In OSPF, a type of stub area in which external route advertisements can be flooded.
open message	Message sent between BGP peers to establish communication.
Open Shortest Path First protocol (OSPF)	A link-state interior gateway protocol (IGP) that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).
origin	Value assigned to a BGP route to indicate whether the first router to advertise the route learned it from an external, internal, or unknown source.

Term	Definition
path-vector protocol	Protocol that uses the path between autonomous systems (ASs) to select the best route, rather than the shortest distance or the characteristics of the route (link state). BGP is a path-vector protocol. In contrast, RIP is a distance-vector protocol, and OSPF and IS-IS are link-state protocols.
peer	Immediately adjacent router with which a protocol relationship has been established. See also *neighbor*.
peering	The practice of exchanging Internet traffic with directly connected peers according to commercial and contractual agreements.
point of presence (POP)	Access point to the Internet, having a unique IP address, where telecommunications equipment is located. POPs usually belong to Internet service providers (ISPs) or telephone companies.
poison reverse	An efficiency technique in a RIP network. By setting the number of hops to an unavailable router to 16 hops or more, a router informs all the other routers in the network. Because RIP allows only up to 15 hops to another router, this technique reduces RIP updates and helps defeat large routing loops. See also *split horizon*.
propagation	Process of translating and forwarding route information discovered by one routing protocol in the update messages of another routing protocol. Route propagation is also known as route redistribution.
reachability	In BGP, the feasibility of a route.
round-robin	Scheduling algorithm in which items have the same priority and are handled in a fixed cyclic order.
route advertisement	Distribution of routing information at specified intervals throughout a network, to establish adjacencies with neighbors and communicate usable routes to active destinations. See also *link-state advertisement (LSA)*.
route aggregation	Combining groups of routes with common addresses into a single entry in the routing table, to decrease routing table size and the number of route advertisements sent by a router.
route reflection	In BGP, configuring a group of routers into a cluster and having one system act as a route reflector, redistributing routes from outside the cluster to all routers in the cluster. Routers in a cluster do not need to be fully meshed.
Routing Information Protocol (RIP)	Distance-vector routing protocol that keeps a database of routing information gathered from periodic broadcasts by each router in a network.
routing table	Table stored on a router that keeps track of all possible paths (routes) between sources and destinations in a network and, in some cases, metrics associated with the routes.
split horizon	An efficiency technique in a RIP network. A router reduces the number of RIP updates in the network by not retransmitting a route advertisement out the interface through which it was received. Split-horizon updates also help prevent routing loops. See also *poison reverse*.
static routing	Routing method in which routes are manually entered in the routing table and do not change unless you explicitly update them. Unlike dynamic routes, which must be imported into the routing table each time a host comes online, static routes are available immediately. Static routes are generally preferred over other types of routes. Compare *dynamic routing*.
stub area	In OSPF, an area through which or into which autonomous system (AS) external route advertisements are not flooded.
subautonomous system (sub-AS)	Autonomous system (AS) members of a BGP confederation.
subnetwork	Subdivision of a network, which functions exactly like a network except that it has a more specific address and subnet mask (destination prefix).
three-way handshake	Process by which two routers synchronize protocols and establish a bidirectional connection.

Term	Definition
topology database	Map of connections between the nodes in a network. The topology database is stored in each node.
triggered update	In a network that uses RIP, a routing update that is automatically sent whenever routing information changes.
virtual link	In OSPF, a link you create between two area border routers (ABRs) that have an interface to a common nonbackbone area, to connect a third area to the backbone area. One of the area border routers must be directly connected to the backbone area.

Routing Overview

Routing is the transmission of data packets from a source to a destination address. For packets to be correctly forwarded to the appropriate host address, the host must have a unique numeric identifier or IP address. The unique IP address of the destination host forms entries in the routing table. These entries are primarily responsible for determining the path that a packet traverses when transmitted from source to destination.

This overview includes these topics:

Networks and Subnetworks

Large groups of machines that are interconnected and can communicate with one another form networks. Typically, networks identify large systems of computers and devices that are owned or operated by a single entity. Traffic is routed between or through the networks as data is passed from host to host.

As networks grow large, the ability to maintain the network and effectively route traffic between hosts within the network becomes increasingly difficult. To accommodate growth, networks are divided into subnetworks. Fundamentally, subnetworks behave exactly like networks, except that they are identified by a more specific network address and subnet mask

(destination prefix). Subnetworks have routing gateways and share routing information in exactly the same way as large networks.

Autonomous Systems

A large network or collection of routers under a single administrative authority is termed an autonomous system (AS). Autonomous systems are identified by a unique numeric identifier that is assigned by the Internet Assigned Numbers Authority (IANA). Typically, the hosts within an AS are treated as internal peers, and hosts in a peer AS are treated as external peers. The status of the relationship between hosts—internal or external—governs the protocol used to exchange routing information.

Interior and Exterior Gateway Protocols

Routing information that is shared within an AS is transmitted by an interior gateway protocol (IGP). Of the different IGPs, the most common are RIP, OSPF, and IS-IS. IGPs are designed to be fast acting and light duty. They typically incorporate only a moderate security system, because trusted internal peers do not require the stringent security measures that untrusted peers require. As a result, you can usually begin routing within an AS by enabling the IGP on all internal interfaces and performing minimal additional configuration. You do not need to establish individual adjacencies.

Routing information that is shared with a peer AS is transmitted by an exterior gateway protocol (EGP). The primary EGP in use in almost all networks is the Border Gateway Protocol (BGP). BGP is designed to be very secure. Individual connections must be explicitly configured on each side of the link. As a result, although large numbers of connections are difficult to configure and maintain, each connection is secure.

Routing Tables

To route traffic from a source host to a destination host, the routers through which the traffic will pass must learn the path that the packet is to take. Once learned, the information is stored in routing tables. The routing table maintains a list of all the possible paths from point A to point B. Figure 55 shows a simple network of routers.

Figure 55: Simple Network Topology

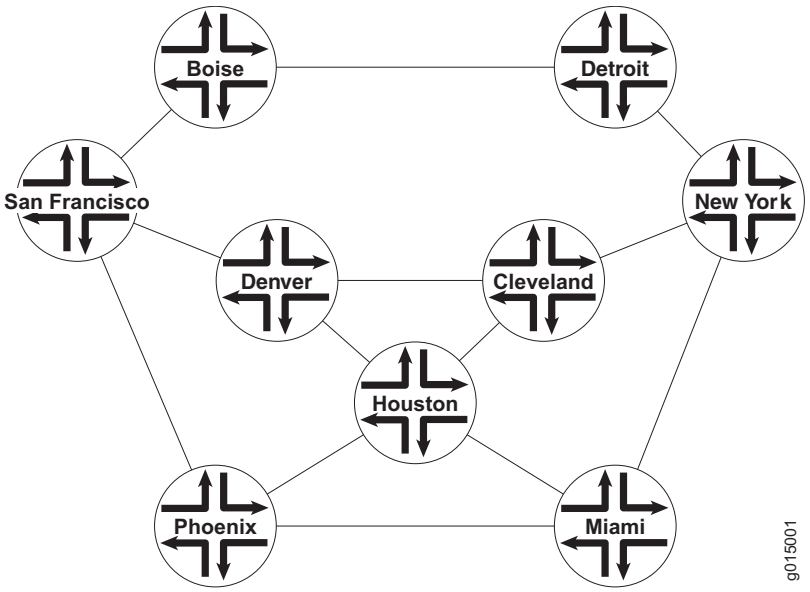

This simple network provides multiple ways to get from host San Francisco to host Miami. The packet can follow the path through Denver and Cleveland. Alternatively, the packet can be routed through Phoenix and directly to Miami. The routing table includes all the possible paths and combinations—an exhaustive list of all the ways to get from the source to the destination.

The routing table must include every possible path from a source to a destination. Routing tables for the network in Figure 55 must include entries for San Francisco-Denver, San Francisco-Cleveland, San Francisco-Miami, Denver-Cleveland, and so on. As the number of sources and destinations increases, the routing table quickly becomes large. The unwieldy size of routing tables in the primary reason for the division of networks into subnetworks.

Forwarding Tables

If the routing table is a list of all the possible paths a packet can take, the forwarding table is a list of only the best routes to a particular destination. The best path is determined according to the particular routing protocol being used, but generally the number of hops between the source and destination determines the best possible route.

In the network shown in Figure 55, because the path with the fewest number of hops from San Francisco to Miami is through Phoenix, the forwarding table distills all the possible San Francisco-Miami routes into the single route through Phoenix. All traffic with a destination address of Miami is sent directly to the next hop, Phoenix.

After it receives a packet, the Phoenix router performs another route lookup, using the same destination address. The Phoenix router then routes the packet

appropriately. Although it considers the entire path, the router at any individual hop along the way is responsible only for transmitting the packet to the next hop in the path. If the Phoenix router is managing its traffic in a particular way, it might send the packet through Houston on its route to Miami. This scenario is likely if specific customer traffic is treated as priority traffic and routed through a faster or more direct route, while all other traffic is treated as nonpriority traffic.

Dynamic and Static Routing

Entries are imported into a router's routing table from dynamic routing protocols or by manual inclusion as static routes. Dynamic routing protocols allow routers to learn the network topology from the network. The routers within the network send out routing information in the form of route advertisements. These advertisements establish and communicate active destinations, which are then shared with other routers in the network.

Although dynamic routing protocols are extremely useful, they have associated costs. Because they use the network to advertise routes, dynamic routing protocols consume bandwidth. Additionally, because they rely on the transmission and receipt of route advertisements to build a routing table, dynamic routing protocols create a delay (latency) between the time a router is powered on and the time during which routes are imported into the routing table. Some routes are therefore effectively unavailable until the routing table is completely updated, when the router first comes online or when routes change within the network (due to a host going offline, for example).

Static routing avoids the bandwidth cost and route import latency of dynamic routing. Static routes are manually included in the routing table, and never change unless you explicitly update them. Static routes are automatically imported into the routing table when a router first comes online. Additionally, all traffic destined for a static address is routed through the same router. This feature is particularly useful for networks with customers whose traffic must always flow through the same routers. Figure 56 shows a network that uses static routes.

Figure 56: Static Routing Example

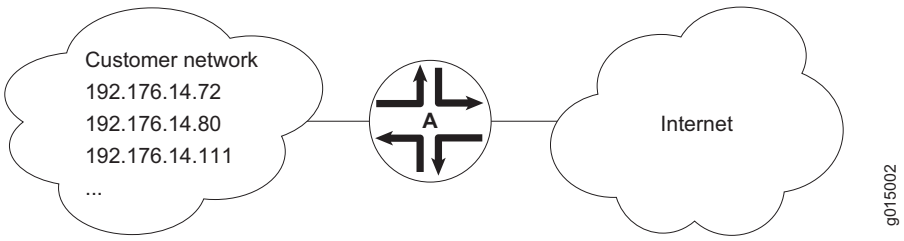

In Figure 56, the customer routes in the 192.176.14/24 subnetwork are static routes. These are hard links to specific customer hosts that never change. Because all traffic destined for any of these routes is forwarded through router A, these routes are included as static routes in router A's routing table. Router A then advertises these routes to other hosts so that traffic can be routed to and from them.

Route Advertisements

The routing table and forwarding table contain the routes for the routers within a network. These routes are learned through the exchange of route advertisements. Route advertisements are exchanged according to the particular protocol being employed within the network.

Generally, a router transmits hello packets out each of its interfaces. Neighboring routers detect these packets and establish adjacencies with the router. The adjacencies are then shared with other neighboring routers, which allows the routers to build up the entire network topology in a topology database, as shown in Figure 57.

Figure 57: Route Advertisement

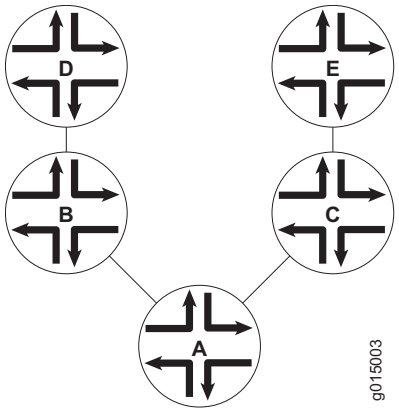

In Figure 57, router A sends out hello packets to each of its neighbors. Routers B and C detect these packets and establish an adjacent relationship with router A. Router B and C then share this information with their neighbors, routers D and E, respectively. By sharing information throughout the network, the routers create a network topology, which they use to determine the paths to all possible destinations within the network. The routes are then distilled into the forwarding table of best routes according to the route selection criteria of the protocol in use.

Route Aggregation

As the number of hosts in a network increases, the routing and forwarding tables must establish and maintain more routes. As these tables become larger, the time routers require to look up particular routes so that packets can be forwarded becomes prohibitive. The solution to the problem of growing routing tables is to group (aggregate) the routers by subnetwork, as shown in Figure 58.

Figure 58: Route Aggregation

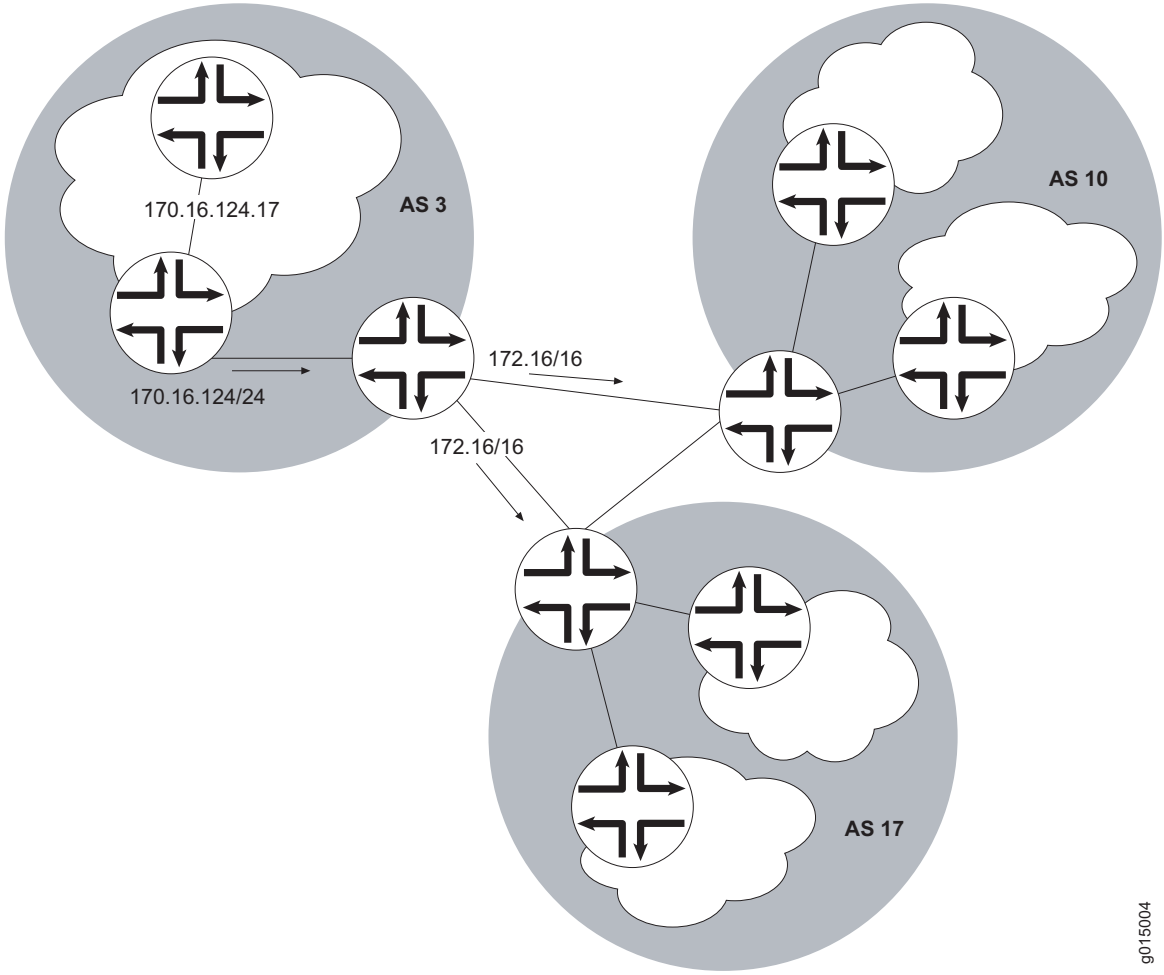

Figure 58 shows three different ASs. Each AS contains multiple subnetworks with thousands of host addresses. To allow traffic to be sent from any host to any host, the routing tables for each host must include a route for each destination. For the routing tables to include every combination of hosts, the flooding of route advertisements for each possible route becomes prohibitive. In a network of hosts numbering in the thousands or even millions, simple route advertisement is not only impractical but impossible.

By employing route aggregation, instead of advertising a route for each host in AS 3, the gateway router advertises only a single route that includes all the routes to all the hosts within the AS. For example, instead of advertising the particular route 170.16.124.17, the AS 3 gateway router advertises only 170.16/16. This single route advertisement encompasses all the hosts within the 170.16/16 subnetwork, which reduces the number of routes in the routing table from 2^{16} (one for every possible IP address within the subnetwork) to 1. Any traffic

destined for a host within the AS is forwarded to the gateway router, which is then responsible for forwarding the packet to the appropriate host.

Similarly, in this example, the gateway router is responsible for maintaining 2^{16} routes within the AS (in addition to any external routes). The division of this AS into subnetworks allows for further route aggregation to reduce this number. In the subnetwork in the example, the subnetwork gateway router advertises only a single route (170.16.124/24), which reduces the number of routes from 2^8 to 1.

RIP Overview

In a Routing Information Protocol (RIP) network, each router's forwarding table is distributed among the nodes through the flooding of routing table information. Because topology changes are flooded throughout the network, every node maintains the same list of destinations. Packets are then routed to these destinations based on path-cost calculations done at each node in the network.

This overview includes the following topics:

NOTE: The J-series Services Router supports both RIP version 1 and RIP version 2. In this guide, the term RIP refers to both versions of the protocol.

Distance-Vector Routing Protocols

Distance-vector routing protocols transmit routing information that includes a distance vector, typically expressed as the number of hops to the destination. This information is flooded out all protocol-enabled interfaces at regular intervals (every 30 seconds in the case of RIP) to create a network map that is stored in each node's local topology database.Figure 59 shows how distance-vector routing works.

Figure 59: Distance-Vector Protocol

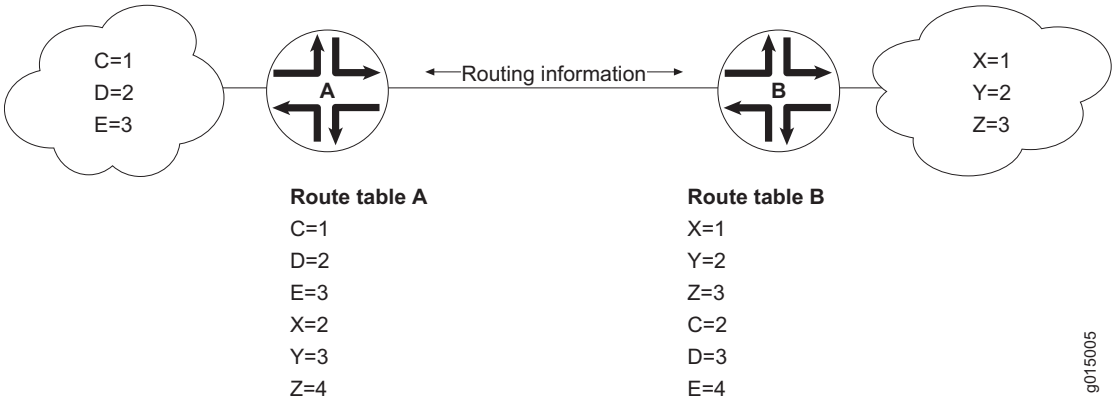

In Figure 59, routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors routers C, D, and E, which are 1, 2, and 3 hops away, respectively. Router B has known RIP neighbors routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each router floods its entire routing table information out all RIP-enabled interfaces. In this case, flooding exchanges routing table information across the RIP link.

When router A receives routing information from router B, it adds 1 to the hop count to determine the new hop count. For example, router X has a hop count of 1, but when router A imports the route to X, the new hop count is 2. The imported route also includes information about where the route was learned, so that the original route is imported as a route to router X through router B with a hop count of 2.

When multiple routes to the same host are received, RIP uses the distance-vector algorithm to determine which path to import into the forwarding table. The route with the smallest hop count is imported. If there are multiple routes with the same hop count, all are imported into the forwarding table, and traffic is sent along the paths in round-robin fashion.

Maximizing Hop Count

The successful routing of traffic across a RIP network requires that every node in the network maintain the same view of the topology. Topology information is broadcast between RIP neighbors every 30 seconds. If router A is many hops away from a new host, router B, the route to B might take significant time to propagate through the network and be imported into router A's routing table. If the two routers are 5 hops away from each other, router A cannot import the route to router B until 2.5 minutes after router B is online. For large numbers of hops, the delay becomes prohibitive. To help prevent this delay from growing arbitrarily large, RIP enforces a maximum hop count of 15 hops. Any prefix that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity. This maximum hop count is called the network diameter.

RIP Packets

Routing information is exchanged in a RIP network by RIP request and RIP response packets. A router that has just booted can broadcast a RIP request on all RIP-enabled interfaces. Any routers running RIP on those links receive the request and respond by sending a RIP response packet immediately to the router. The response packet contains the routing table information required to build the local copy of the network topology map.

In the absence of RIP request packets, all RIP routers broadcast a RIP response packet every 30 seconds on all RIP-enabled interfaces. The RIP broadcast is the primary way in which topology information is flooded throughout the network.

Once a router learns about a particular destination through RIP, it starts a timer. Every time it receives a new response packet with information about the destination, the router resets the timer to zero. However, if the router receives no updates about a particular destination for 180 seconds, it removes the destination from its RIP routing table.

In addition to the regular transmission of RIP packets every 30 seconds, if a router detects a new neighbor or detects that an interface is unavailable, it generates a triggered update. The new routing information is immediately broadcast out all RIP-enabled interfaces, and the change is reflected in all subsequent RIP response packets.

Split Horizon and Poison Reverse Efficiency Techniques

Because RIP functions by periodically flooding the entire routing table out to the network, it generates a lot of traffic. The split horizon and poison reverse techniques can help reduce the amount of network traffic originated by RIP hosts and make the transmission of routing information more efficient.

If a router receives a set of route advertisements on a particular interface, RIP determines that those advertisements do not need to be retransmitted out the same interface. This technique, known as split horizon, helps limit the amount of RIP routing traffic by eliminating information that other neighbors on that interface have already learned. Figure 60 shows an example of the split horizon technique.

Figure 60: Split Horizon Example

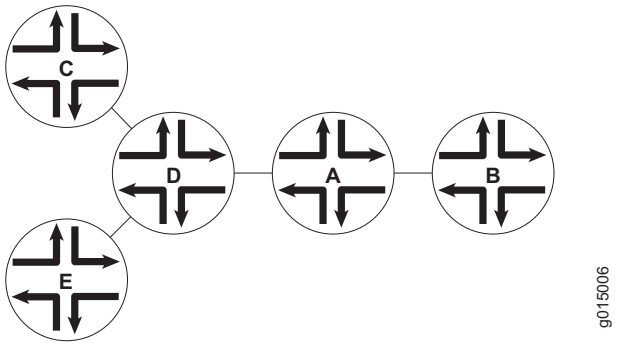

In Figure 60, router A advertises routes to routers C, D, and E to router B. In this example, router A can reach router C in 2 hops. When router A advertises the route to router B, B imports it as a route to router C through router A in 3 hops. If router B then readvertised this route to router A, A would import it as a route to router C through router B in 4 hops. However, the advertisement from router B to router A is unnecessary, because router A can already reach the route in 2 hops. The split horizon technique helps reduce extra traffic by eliminating this type of route advertisement.

Similarly, the poison reverse technique helps to optimize the transmission of routing information and improve the time to reach network convergence. If router A learns about unreachable routes through one of its interfaces, it advertises those routes as unreachable (hop count of 16) out the same interface. Figure 61 shows an example of the poison reverse technique.

Figure 61: Poison Reverse Example

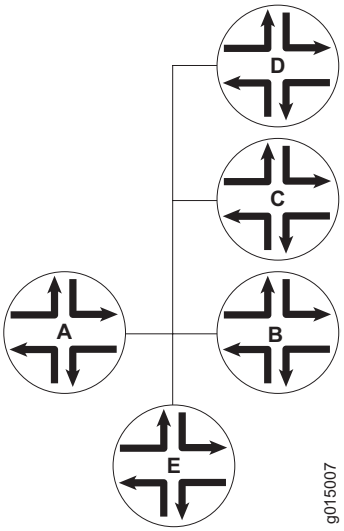

In Figure 61, router A learns through one of its interfaces that routes to routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs router B that hosts C, D, and E are definitely not reachable through router A.

Limitations of Unidirectional Connectivity

Because RIP processes routing information based solely on the receipt of routing table updates, it cannot ensure bidirectional connectivity. As Figure 62 shows, RIP networks are limited by their unidirectional connectivity.

Figure 62: Limitations of Unidirectional Connectivity

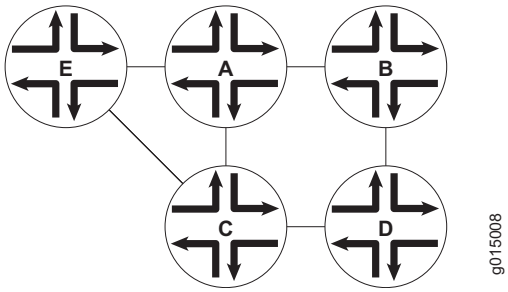

In Figure 62, routers A and D flood their routing table information to router B. Because the path to router E has the fewest hops when routed through router A, that route is imported into router B's forwarding table. However, suppose that router A can transmit traffic but is not receiving traffic from router B due to an unavailable link or invalid routing policy. If the only route to router E is through router A, any traffic destined for router A is lost, because bidirectional connectivity was never established.

OSPF establishes bidirectional connectivity with a three-way handshake. For more information, see "Link-State Advertisements" on page 270.

OSPF Overview

In an Open Shortest Path First (OSPF) network, the network topology is distributed among the nodes of the autonomous system (AS) and is regularly updated through the exchange of link-state advertisements (LSAs). As a result, OSPF is known as a link-state protocol. Because topology changes are flooded throughout the network, every node maintains the same copy of the network map in its local topological database. Packets are then routed based on the shared topology using the shortest path first (SPF) algorithm.

This overview includes the following topics:

NOTE: The J-series services gateway supports both OSPF version 2 and OSPF version 3. In this guide, the term OSPF refers to both versions of the protocol.

Link-State Advertisements

OSPF creates a topology map by flooding link-state advertisements (LSAs) across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in Figure 63.

Figure 63: OSPF Three-Way Handshake

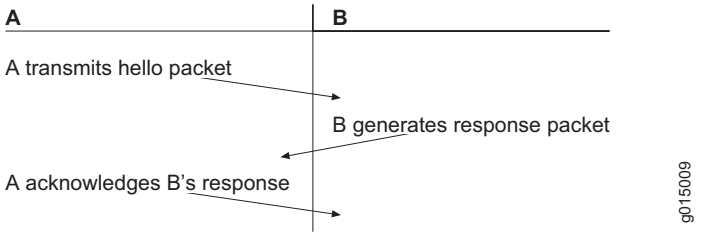

In Figure 63, router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that router B can receive traffic from router A. Router B generates a response to router A to acknowledge receipt of the hello packet. When router A receives the response, it establishes that router B can receive traffic from router A. Router A then generates a final response packet to inform router B that router A can receive traffic from router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

Role of the Designated Router

Large local area networks (LANs) that have many routers and therefore many OSPF adjacencies can produce heavy control-packet traffic as LSAs are flooded across the network. To alleviate the potential traffic problem, OSPF uses designated routers (DRs). Rather than broadcasting LSAs to all their OSPF neighbors, the routers send their LSAs to the designated router, which processes the LSAs, generates responses, and multicasts topology updates to all OSPF routers.

In LANs, the election of the designated router takes place when the OSPF network is initially established. When the first OSPF links are active, the router with the highest router identifier (defined by the router-id configuration value or the loopback address) is elected designated router. The router with the second highest router identifier is elected the backup designated router (BDR). If the designated router fails or loses connectivity, the BDR assumes its role and a new BDR election takes place between all the routers in the OSPF network.

Path Cost Metrics

Once the topology is shared across the network, OSPF uses it to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest path first (SPF) algorithm. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion.

OSPF allows you to manually assign a cost (or metric) to a particular path segment to control the flow of packets across the network.

Areas and Area Border Routers

The OSPF networks in an AS are administratively grouped into areas. Each area within an AS operates like an independent network and has a unique 32-bit area ID, which functions like a network address. Within an area, the topology database contains only information about the area, LSAs are flooded only to nodes within the area, and routes are computed only within the area. Subnetworks are divided into other areas, which are connected to form the whole of the main network.

The central area of an AS, called the backbone area, has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation, but they are not IP addresses. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by a router that has interfaces in more than one area. These connecting routers are called area border routers (ABRs). Figure 64 shows an OSPF topology of three areas connected by two area border routers.

Figure 64: Multiarea OSPF Topology

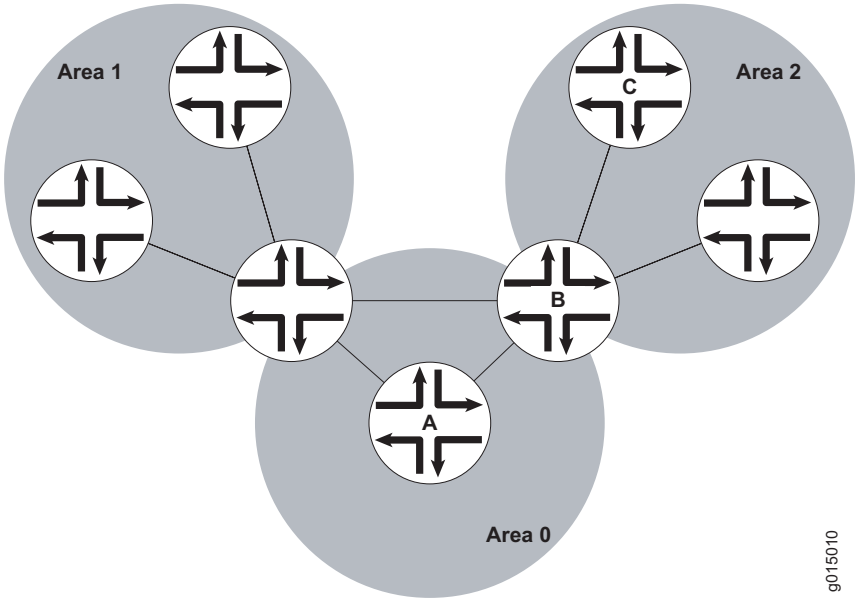

Area border routers are responsible for sharing topology information between areas. They summarize the link-state records of each area and advertise destination address summaries to neighboring areas. The advertisements contain the ID of the area in which each destination lies, so that packets are routed to the appropriate area border router. For example, in the OSPF areas shown in Figure 64, packets sent from router A to router C are automatically routed through area border router B.

Role of the Backbone Area

An OSPF restriction requires all areas to be directly connected to the backbone area so that packets can be properly routed. All packets are routed first to the backbone area by default. Packets that are destined for an area other than the backbone area are then routed to the appropriate area border router and on to the remote host within the destination area.

In large networks with many areas, in which direct connectivity between all areas and the backbone area is physically difficult or impossible, you can configure virtual links to connect noncontiguous areas. For example, Figure 65 shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

Figure 65: OSPF Topology with a Virtual Link

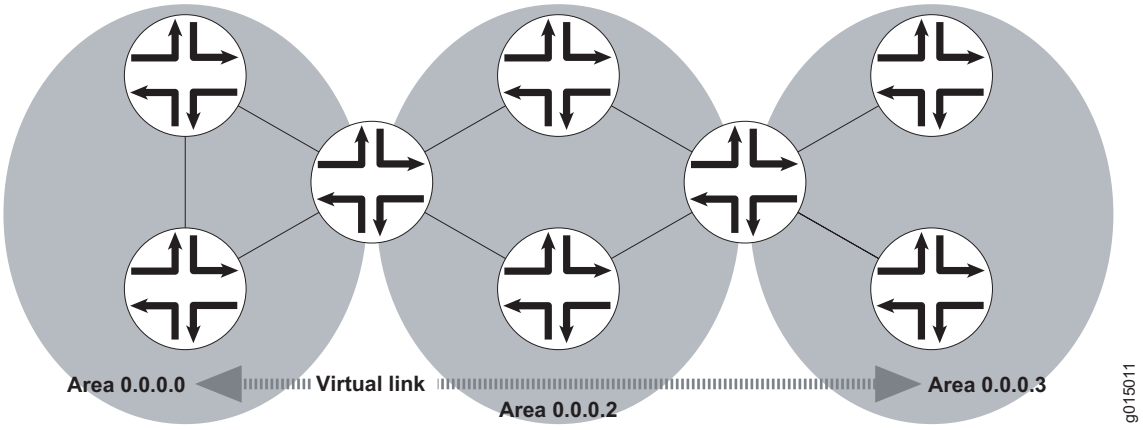

In the topology shown in Figure 65, a virtual link is established between area 0.0.0.3 and the backbone area through area 0.0.0.2. All outbound traffic destined for other areas is routed through area 0.0.0.2 to the backbone area and then to the appropriate area border router. All inbound traffic destined for area 0.0.0.3 is routed to the backbone area and then through area 0.0.0.2.

Stub Areas and Not-So-Stubby Areas

Figure 66 shows an AS across which many external routes are advertised. If external routes make up a significant portion of a topology database, you can suppress the advertisements in areas that do not have links outside the network. By doing so, you can reduce the amount of memory the nodes use to maintain the topology database and free it for other uses.

Figure 66: OSPF AS Network with Stub Areas and NSSAs

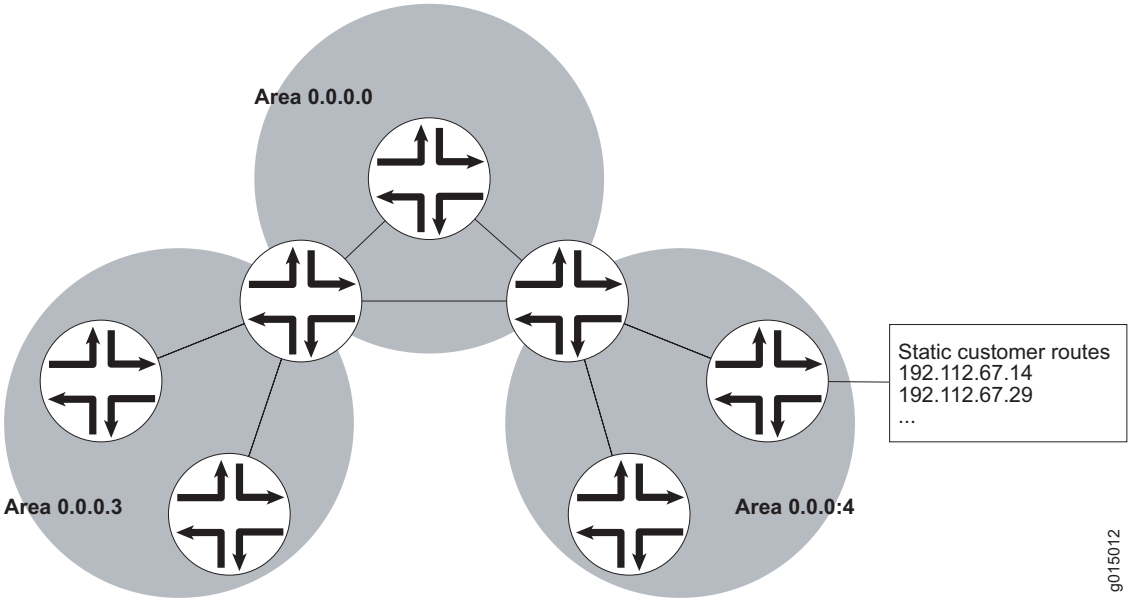

To control the advertisement of external routes into an area, OSPF uses stub areas. By designating an area border router interface to the area as a stub interface, you suppress external route advertisements through the area border router. Instead, the area border router automatically advertises a default route (through itself) in place of the external routes. Packets destined for external routes are automatically sent to the area border router, which acts as a gateway for outbound traffic and routes them appropriately.

For example, area 0.0.0.3 in Figure 66 is not directly connected to the outside network. All outbound traffic is routed through the area border router to the backbone and then to the destination addresses. By designating area 0.0.0.3 a stub area, you reduce the size of the topology database for that area by limiting the route entries to only those routes internal to the area.

Like area 0.0.0.3 in Figure 66, area 0.0.0.4 has no external connections. However, area 0.0.0.4 has static customer routes that are not internal OSPF routes. You can limit the external route advertisements to the area and advertise the static customer routes by designating it a not-so-stubby area (NSSA). External routes are flooded into the NSSA and then leaked to the other areas, but external routes from other areas are not advertised within the NSSA.

BGP Overview

The Border Gateway Protocol (BGP) is an exterior gateway protocol (EGP) used primarily to establish point-to-point connections and transmit data between peer ASs. Unlike the IGPs RIP and OSPF, BGP must explicitly advertise the routes between its peers. The route advertisements determine prefix reachability and the

way packets are routed between BGP neighbors. Because BGP uses the packet path to determine route selection, it is considered a path-vector protocol.

This overview includes the following topics:

Point-to-Point Connections

To establish point-to-point connections between peer ASs, you configure a BGP session on each interface of a point-to-point link. Figure 67 shows an example of a BGP peering session.

Figure 67: BGP Peering Session

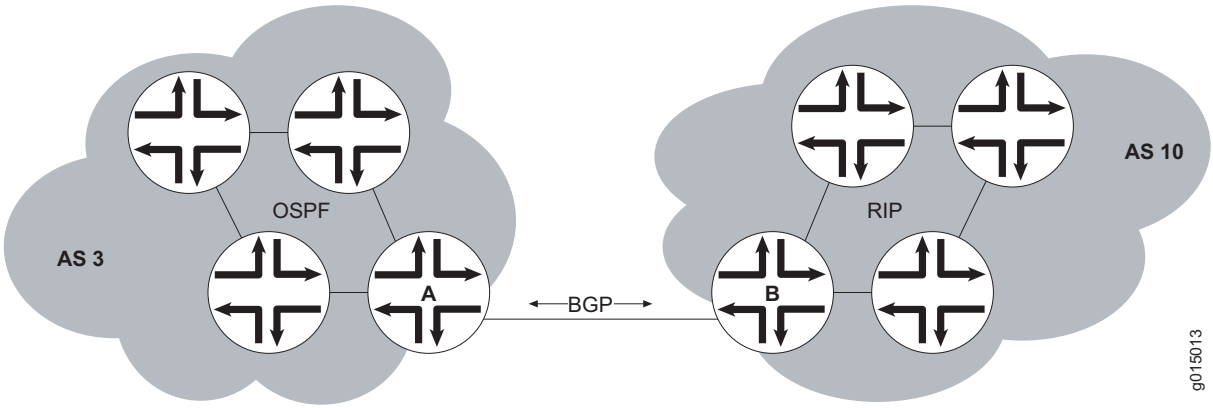

In Figure 67, router A is a gateway router for AS 3, and router B is a gateway router for AS 10. For traffic internal to either AS, an IGP (OSPF, for instance) is used. To route traffic between peer ASs, a BGP session is used.

BGP Messages for Session Establishment

When the routers on either end of a BGP session first boot, the session between them is in the Idle state. The BGP session remains idle until a start event is detected. Typically, the start event is the configuration of a new BGP session or the resetting of an existing BGP session. At boot time, the start event is generated by the router as the BGP session is initiated.

After it detects a start event, the BGP host sends TCP request packets to its configured BGP neighbors. These packets are directed only to neighboring interfaces that have been explicitly configured as BGP neighbors. Upon receipt of the TCP request packet, the neighboring host generates a TCP response to complete the three-way handshake and establish a TCP connection between the peers. While this handshake is taking place, the BGP state for the connection is Connect. If a TCP timeout occurs while the originating host is waiting for a TCP response packet, the BGP state for the connection is Active. The Active state indicates that the router is actively listening for a TCP response and the TCP retry timer has been initiated.

Once a TCP connection has been established between both ends of a BGP session, the BGP session state is OpenSent, indicating that the originating router has generated an open message. The open message is an initial BGP handshake that must occur before any route advertisement can take place. Upon receipt of the open message, the neighboring router generates a keepalive message. Receipt of the keepalive message establishes a point-to-point connection, and the BGP session state transitions to Established. While the originating host waits for the keepalive response packet, the BGP session state is OpenConfirm.

BGP Messages for Session Maintenance

Once a BGP session has been established, the BGP peers exchange route advertisements by means of update messages. Update messages contain a one or more route advertisements, and they can contain one or more prefixes that are to be removed from the BGP routing table. If the peers need to advertise multiple routes, they generate and send multiple update messages as they detect changes to the network. In the absence of changes to the routing table, no update messages are generated.

While a BGP session is active, each router on the BGP session generates keepalive messages periodically. The timing of these messages is determined by the hold time on the session. The hold time is a negotiated value specifying the number of seconds that can elapse without keepalive messages before BGP designates the link inactive. Three messages are sent during every hold time interval.

When a peer connection is closed (either by error or if the BGP session is closed), a notification message is generated and sent to the peer router that did not experience the error or did not terminate the BGP session.

IBGP and EBGP

BGP uses two primary modes of information exchange, internal BGP (IBGP) and external BGP (EBGP), to communicate with internal and external peers, respectively.

Peer ASs establish links through an external peer BGP session. As a result, all route advertisement between the external peers takes place by means of the EBGP mode of information exchange. To propagate the routes through the AS and advertise them to internal peers, BGP uses IBGP. To advertise the routes to a different peer AS, BGP again uses EBGP.

To avoid routing loops, IBGP does not advertise routes learned from an internal BGP peer to other internal BGP peers. For this reason, BGP cannot propagate routes throughout an AS by passing them from one router to another. Instead, BGP requires that all internal peers be fully meshed so that any route advertised by one router is advertised to all peers within the AS.

As a network grows, the full mesh requirement becomes increasingly difficult to manage. In a network with 1000 routers, the addition of a single router requires that all the routers in the network be modified to account for the new addition. To combat these scaling problems, BGP uses route reflection and BGP confederations.

For information about route reflection, see "Scaling BGP for Large Networks" on page 280. For information about routing confederations, see "Scaling BGP for Large Networks" on page 280.

Route Selection

A local BGP router uses the following primary criteria to select a route from the routing table for the forwarding table:

1. Next-hop accessible—If the next hop is inaccessible, the local router does not consider the route. The router must verify that it has a route to the BGP next-hop address. If a local route to the next hop does not exist, the local route does not include the router in its forwarding table. If such a route exists, route selection continues.

2. Highest local preference—The local router selects the route with the highest local preference value. If multiple routes have the same preference, route selection continues. (For more information, see "Local Preference" on page 278.)

3. Shortest AS path—The local router selects the route with the fewest entries in the AS path. If multiple routes have the same AS path length, route selection continues. (For more information, see "AS Path" on page 279.)

4. Lowest origin—The local router selects the route with the lowest origin value. If multiple routes have the same origin value, route selection continues. (For more information, see "Origin" on page 279.)

5. Lowest MED value—The local router selects the route with the lowest multiple exit discriminator (MED) value. If multiple routes have the same MED value, route selection continues. (For more information, see "Multiple Exit Discriminator" on page 280.)

If more than one route remains after all these criteria are evaluated, the local BGP router evaluates a set of secondary criteria to select the single route to a destination

for its forwarding table. The secondary criteria include whether the route was learned through an EBGP or IBGP, the IGP route metric, and the router ID.

Local Preference

The local preference is typically used to direct all outbound AS traffic to a certain peer. When you configure a local preference, all routes that are advertised through that peer are assigned the preference value. The preference is a numeric value, and higher values are preferred during BGP route selection. Figure 68 illustrates how to use local preference to determine BGP route selection.

Figure 68: Local Preference

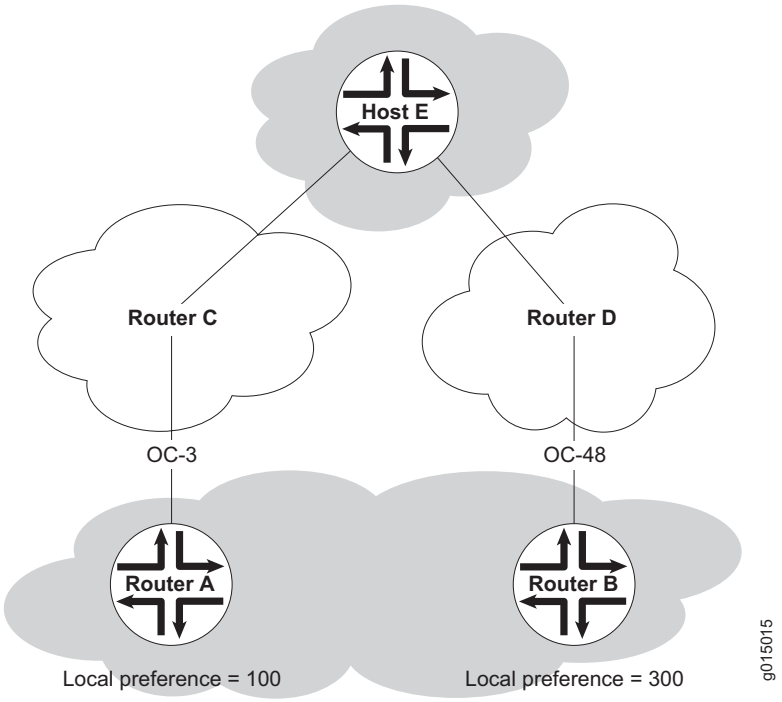

The network in Figure 68 shows two possible routes to the prefixes accessible through host E. The first route, through router A, uses an OC3 link to router C and is then forwarded to host E. The second route, through router B, uses an OC48 link to router D and is then forwarded to host E. Although the number of hops to host E is identical regardless of the route selected, the route through router B is more desirable because of the increased bandwidth. To force traffic through router B, you can set the local preference on router A to 100 and the local preference on router B to 300. During BGP route selection, the route with the higher local preference is selected.

NOTE: In contrast to almost every other metric associated with dynamic routing protocols, the local preference gives higher precedence to the larger value.

AS Path

Routes advertised by BGP maintain a list of the ASs through which the route travels. This information is stored in the route advertisement as the AS path, and it is one of the primary criteria that a local router uses to evaluate BGP routes for inclusion in its forwarding table. Figure 69 shows how BGP creates an AS path.

Figure 69: BGP AS Path

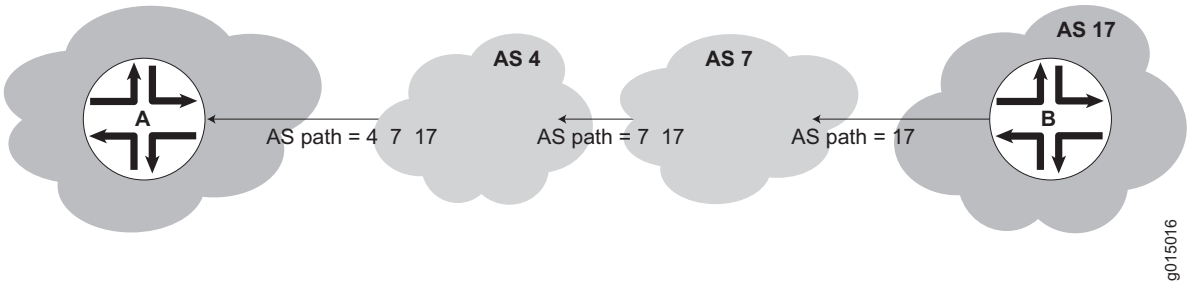

In the network shown in Figure 69, the route from host A to host B travels through two intermediate ASs. As the route advertisement is propagated through the BGP network, it accumulates an AS path number each time it exits one AS and enters another. Each AS number is prepended to the AS path, which is stored as part of the route advertisement. When the route advertisement first leaves host B's AS, the AS path is `17`. When the route is advertised between intermediate ASs, the AS number `7` is prepended to the AS path, which becomes `7 17`. When the route advertisement exits the third AS, the AS path becomes `4 7 17`. The route with the shortest AS path is preferred for inclusion into the BGP forwarding table.

Origin

The BGP router that first advertises a route assigns it of the following values to identify its origin. During route selection, the lowest origin value is preferred.

- `0`—The router originally learned the route through an IGP (OSPF, IS-IS, or a static route).
- `1`—The router originally learned the route through an EGP (most likely BGP).
- `2`—The route's origin is unknown.

Multiple Exit Discriminator

Because the AS path rather than the number of hops between hosts is the primary criterion for BGP route selection, an AS with multiple connections to a neighbor AS can have multiple equivalent AS paths. When the routing table contains two routes to the same host in a neighboring AS, a multiple exit discriminator (MED) metric assigned to each route can determine which to include in the forwarding table. The MED metric you assign can force traffic through a particular exit point in an AS. Figure 70 illustrates how to use an MED metric to determine route selection.

Figure 70: MED Example

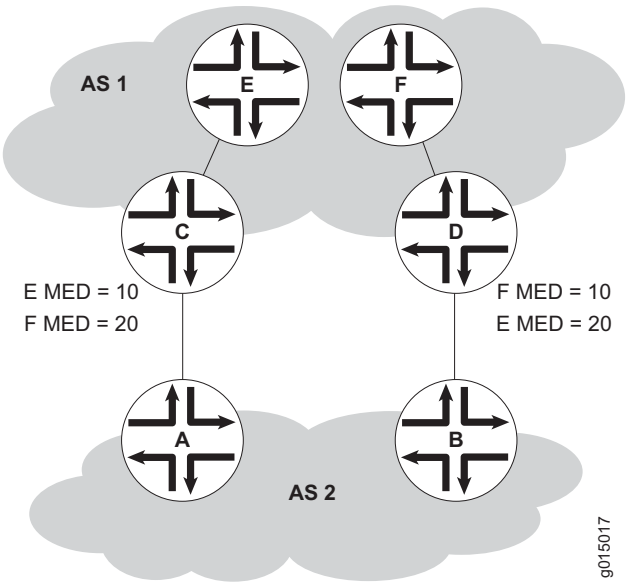

Figure 70 shows AS 1 and AS 2 connected by two separate BGP links to routers C and D. Host E in AS 1 is located nearer router C. Host F also in AS 1, and is located nearer router D. Because the AS paths are equivalent, two routes exist for each host, one through router C and one through router D. To force all traffic destined for host E through router C, network administrator for AS 2 assigns an MED metric for each router to host E at its exit point. An MED metric of 10 is assigned to the route to host E through router C, and an MED metric of 20 is assigned to the route to host E through router D. BGP routers in AS 2 then select the route with the lower MED metric for the forwarding table.

Scaling BGP for Large Networks

BGP is not a flooding protocol like RIP or OSPF. Instead, it is a peering protocol that exchanges routes with fully meshed peers only. However, in large networks, the full mesh requirement causes scaling problems. BGP combats scaling problems with the following methods:

- Route Reflectors—for Added Hierarchy on page 281

- Confederations—for Subdivision on page 283

Route Reflectors—for Added Hierarchy

To use route reflection in an AS, you designate one or more routers as a route reflector—typically, one per point of presence (POP) . Route reflectors have the special BGP ability to readvertise routes learned from an internal peer to other internal peers. So rather than requiring all internal peers to be fully meshed with each other, route reflection requires only that the route reflector be fully meshed with all internal peers. The route reflector and all its internal peers form a cluster, as shown in Figure 71.

Figure 71: Simple Route Reflector Topology (One Cluster)

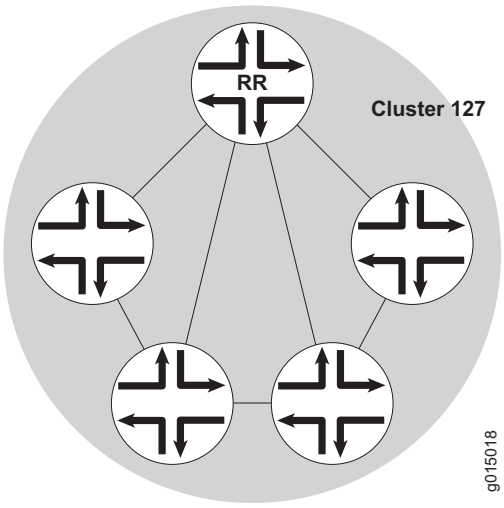

Figure 71 shows router RR configured as the route reflector for cluster 127. The other routers are designated internal peers within the cluster. BGP routes are advertised to router RR by any of the internal peers. RR then readvertises those routes to all other peers within the cluster.

You can configure multiple clusters and link them by configuring a full mesh of route reflectors (see Figure 72).

Figure 72: Basic Route Reflection (Multiple Clusters)

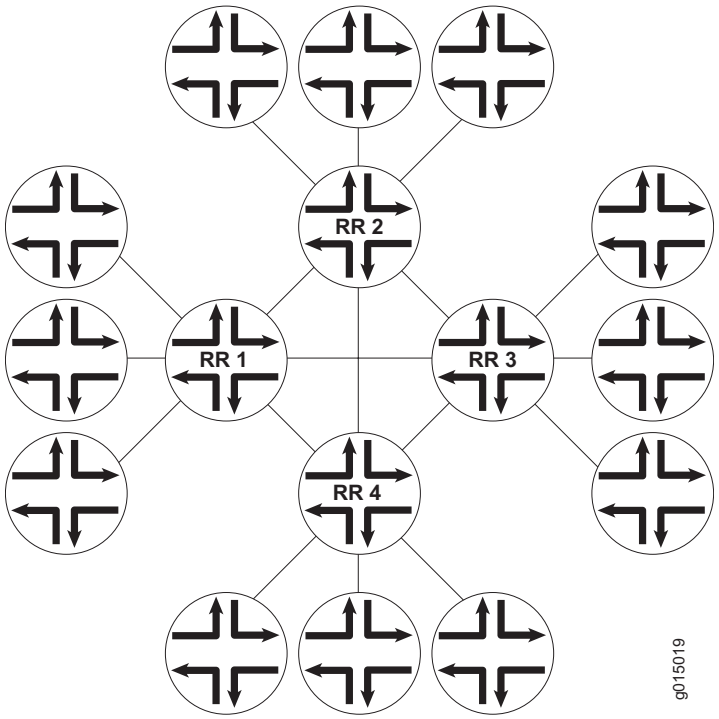

Figure 72 shows route reflectors RR1, RR2, RR3, and RR4 as fully meshed internal peers. When a router advertises a route to reflector RR1, RR1 readvertises the route to the other route reflectors, which, in turn, readvertise the route to the remaining routers within the AS. Route reflection allows the route to be propagated throughout the AS without the scaling problems created by the full mesh requirement.

However, as clusters become large, a full mesh with a route reflector becomes difficult to scale, as does a full mesh between route reflectors. To help offset this problem, you can group clusters of routers together into clusters of clusters for hierarchical route reflection (see Figure 73).

Figure 73: Hierarchical Route Reflection (Clusters of Clusters)

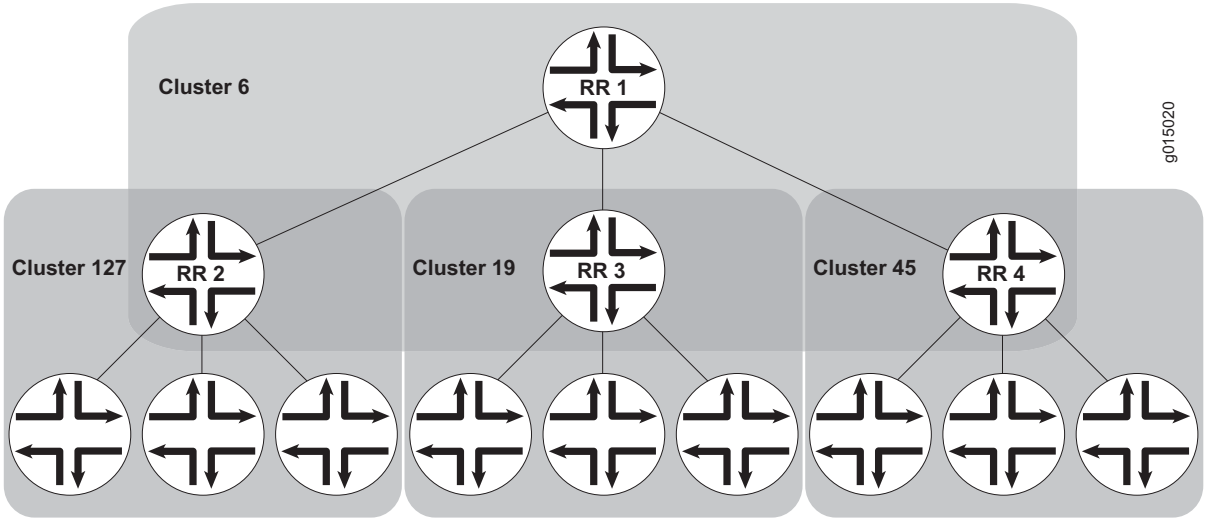

Figure 73 shows RR2, RR3, and RR4 as the route reflectors for clusters 127, 19, and 45, respectively. Rather than fully mesh those route reflectors, the network administrator has configured them as part of another cluster (cluster 6) for which RR1 is the route reflector. When a router advertises a route to RR2, RR2 readvertises the route to all the routers within its own cluster, and then readvertises the route to RR1. RR1 readvertises the route to the routers in its cluster, and those routers propagate the route down through their clusters.

Confederations—for Subdivision

BGP confederations are another way to solve the scaling problems created by the BGP full mesh requirement. BGP confederations effectively break up a large AS into subautonomous systems (sub-ASs). Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64512 and 65535.

Within a sub-AS, the same IBGP full mesh requirement exists. Connections to other confederations are made with standard EBGP, and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

The confederation AS appears whole to other confederation ASs. The AS path received by other ASs shows only the globally assigned AS number. It does not include the confederation sequence or the privately assigned sub-AS numbers. The sub-AS numbers are removed when the route is advertised out of the confederation AS. Figure 74 shows an AS divided into four confederations.

Figure 74: BGP Confederations

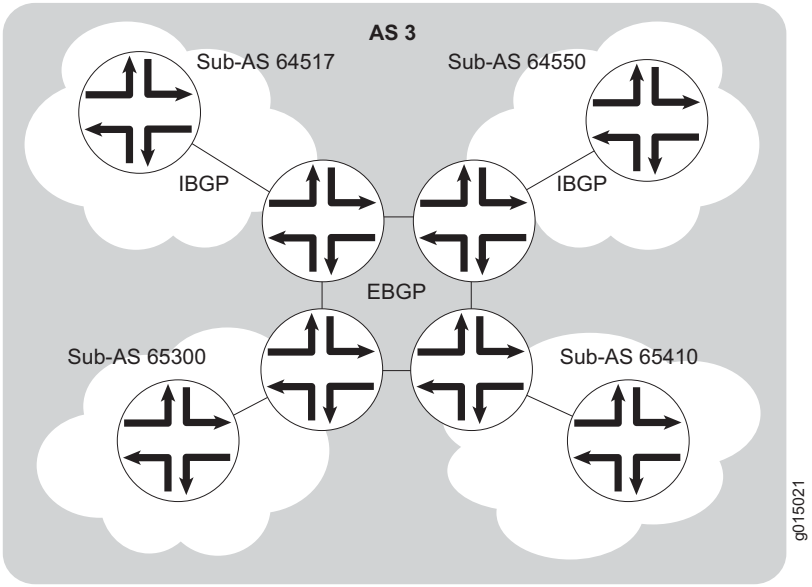

Figure 74 shows AS 3 divided into four sub-ASs, 64517, 64550, 65300, and 65410, which are linked through EBGP sessions. Because the confederations are connected by EBGP, they do not need to be fully meshed. EBGP routes are readvertised to other sub-ASs.

Chapter 13
Configuring Static Routes

Static routes are routes that you explicitly enter into the routing table as permanent additions. Traffic through static routes is always routed the same way.

You can use either J-Web Quick Configuration or a configuration editor to configure static routes.

This chapter contains the following topics. For more information about static routes, see the *JUNOS Routing Protocols Configuration Guide*.

Static Routing Overview

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

Static Route Preferences

A static route destination address can have multiple next hops associated with it. In this case, multiple routes are inserted into the routing table, and route selection must occur. Because the primary criterion for route selection is the route preference, you can control the routes that are used as the primary route for a particular destination by setting the route preference associated with a particular next hop. The routes with a higher preference are always used to route traffic. When you do not set a preferred route, traffic is alternated between routes in round-robin fashion.

Qualified Next Hops

In general, the default properties assigned to a static route apply to all the next-hop addresses configured for the static route. If, however, you want to configure two possible next-hop addresses for a particular route and have them treated differently, you can define one as a qualified next hop.

Qualified next hops allow you to associate one or more properties with a particular next-hop address. You can set an overall preference for a particular static route and then specify a different preference for the qualified next hop. For example, suppose two next-hop addresses (10.10.10.10 and 10.10.10.7) are associated with the static route 192.168.47.5/32. A general preference is assigned to the entire static route, and then a different preference is assigned to only the qualified next-hop address 10.10.10.7. For example:

```
route 192.168.47.5/32 {
next-hop 10.10.10.10;
qualified-next-hop 10.10.10.7 {
preference 2;
}
preference 6;
}
```

In this example, the qualified next hop 10.10.10.7 is assigned the preference 2, and the next-hop 10.10.10.10 is assigned the preference 6.

Control of Static Routes

You can control the importation of static routes into the routing and forwarding tables in a number of ways. Primary ways include assigning one or more of the following attributes to the route:

- retain—Keeps the route in the forwarding table after the routing process shuts down or the Services Router reboots. For more information, see “Route Retention” on page 286.
- no-readvertise—Prevents the route from being readvertised to other routing protocols. For more information, see “Readvertisement Prevention” on page 287.
- passive—Rejects traffic destined for the route. For more information, see “Forced Rejection of Passive Route Traffic” on page 287.

Route Retention

By default, static routes are not retained in the forwarding table when the routing process shuts down. When the routing process starts up again, any routes configured as static routes must be added 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. Retention ensures that the routes are always in the forwarding table, even immediately after a system reboot.

Readvertisement Prevention

Static routes are eligible for readvertisement by other routing protocols by default. In a stub area where you might not want to readvertise these static routes under any circumstances, you can flag the static routes as no-readvertise.

Forced Rejection of Passive Route Traffic

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 passive, 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 passive and its next-hop address is unreachable, the route is included in the routing table and all traffic destined for the route is rejected.

Default Properties

The basic configuration of static routes defines properties for a particular route. To define a set of properties to be used as defaults on all static routes, set those properties as default values. For example:

```
defaults {
retain;
no-readvertise;
passive;
}
route 0.0.0.0/0 next-hop 192.168.1.1;
route 192.168.47.5/32 {
next-hop 10.10.10.10;
qualified-next-hop 10.10.10.7 {
preference 6;
}
preference 2;
}
```

In this example, the retain, no-readvertise, and passive attributes are set as defaults for all static routes. If any local setting for a particular route conflicts with the default values, the local setting supersedes the default.

Before You Begin

Before you begin configuring static routes, complete the following tasks:

- Establish basic connectivity. See "Establishing Basic Connectivity" on page 47.
- Configure network interfaces. See "Configuring Network Interfaces" on page 79.

Configuring Static Routes with Quick Configuration

J-Web Quick Configuration allows you to configure static routes. Figure 75 shows the Quick Configuration Routing page for static routing.

Figure 75: Quick Configuration Routing Page for Static Routing

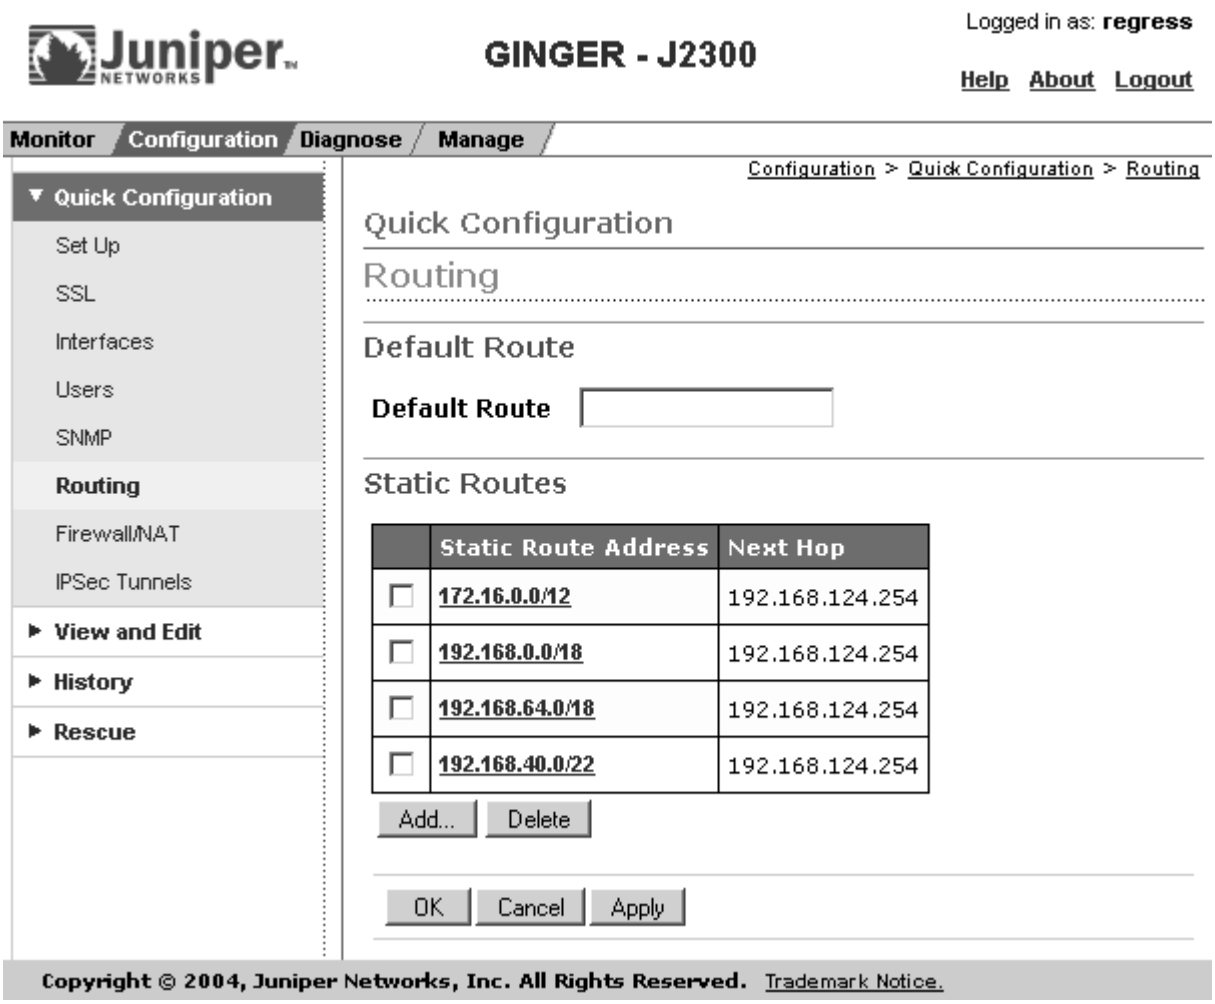

To configure static routes with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > Static Routing**.
2. Enter information into the Static Routing Quick Configuration, as described in Table 103.
3. From the main static routing Quick Configuration page, click one of the following buttons:

- To apply the configuration and stay on the Quick Configuration Routing page for static routing, click **Apply**.
- To apply the configuration and return to the Quick Configuration Routing page, click **OK**.
- To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.

4. To check the configuration, see "Verifying the Static Route Configuration" on page 295.

Table 103: Static Routing Quick Configuration Summary

Field	Function	Your Action
Default Route		
Default Route	Specifies the default gateway for the router.	Type the 32-bit IP address of the Services Gateway's default route in dotted decimal notation.
Static Routes		
Static Route Address (required)	Specifies the static route to add to the routing table.	1. On the main static routing Quick Configuration page, click **Add**. 2. In the Static Route Address box, type the 32-bit IP address of the static route in dotted decimal notation.
Next-Hop Addresses	Specifies the next-hop address or addresses to be used when routing traffic to the static route.	1. In the Add box, type the 32-bit IP address of the next-hop host. 2. Click **Add**. 3. Add more next-hop addresses as necessary. **NOTE:** If a route has multiple next-hop addresses, traffic is routed across each address in round-robin fashion. 4. When you have finished adding next-hop addresses, click **OK**.

Configuring Static Routes with a Configuration Editor

To configure static routes on the Services Router, you must perform the following tasks marked *(Required)*.

- (Required) “Configuring a Basic Set of Static Routes” on page 290
- (Optional)“Controlling Static Route Selection” on page 291
- (Optional)“Controlling Static Routes in the Routing and Forwarding Tables” on page 293
- (Optional)“Defining Default Behavior for All Static Routes” on page 294

For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 127.

Configuring a Basic Set of Static Routes

Customer routes that are connected to stub networks are often configured as static routes. Figure 76 shows a sample network.

Figure 76: Customer Routes Connected to a Stub Network

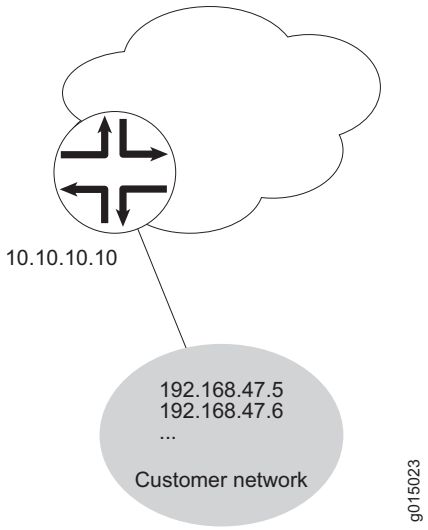

To configure customer routes as static routes, like the ones in Figure 76, follow these steps on the Services Router to which the customer routes are connected:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 104.

3. If you are finished configuring static routes, commit the configuration.

4. Go on to one of the following procedures:

   - To manually control static route selection, see "Controlling Static Route Selection" on page 291.
   - To determine how static routes are imported into the routing and forwarding tables, see "Controlling Static Routes in the Routing and Forwarding Tables" on page 293.
   - To define default properties for static routes, see "Defining Default Behavior for All Static Routes" on page 294.
   - To check the configuration, see "Verifying the Static Route Configuration" on page 295.

Table 104: Configuring Basic Static Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Static** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing options > Static**.	From the top of the configuration hierarchy, enter edit routing-options static
Add the static route 192.168.47.5/32, and define the next-hop address 10.10.10.10.	1. In the Route field, click **Add new entry**. 2. In the Destination box, enter **192.168.47.5/32**. 3. From the Next hop list, select **Next hop**. 4. In the Next hop field, click **Add new entry**. 5. In the Value box, enter **10.10.10.10**. 6. Click **OK**.	Define the static route and set the next-hop address: set route 192.168.47.5 next-hop 10.10.10.10

Controlling Static Route Selection

When multiple next hops exist for a single static route (see Figure 77), you can specify how traffic is to be routed to the destination.

Figure 77: Controlling Static Routes in the Routing and Forwarding Tables

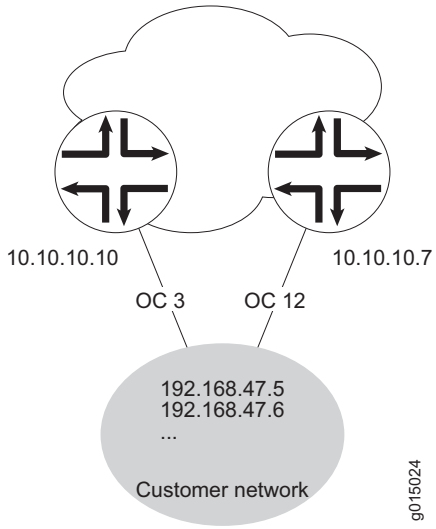

In this example, the static route 192.168.47.5/32 has two possible next hops. Because of the links between those next-hop hosts, host 10.10.10.7 is the preferred path. To configure the static route 192.168.47.5/32 with two next hops and give preference to host 10.10.10.7, follow these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 105.
3. If you are finished configuring static routes, commit the configuration.
4. Go on to one of the following procedures:
   - To determine how static routes are imported into the routing and forwarding tables, see “Controlling Static Routes in the Routing and Forwarding Tables” on page 293.
   - To define default properties for static routes, see “Defining Default Behavior for All Static Routes” on page 294.
   - To check the configuration, see “Verifying the Static Route Configuration” on page 295.

Table 105: Controlling Static Route Selection

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Static** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Static**.	From the top of the configuration hierarchy, enter edit routing-options static
Add the static route 192.168.47.5/32, and define the next-hop address 10.10.10.10.	1. In the Route field, click **Add new entry**. 2. In the Destination box, enter **192.168.47.5/32**. 3. From the Next hop list, select **Next hop**. 4. In the Next hop field, click **Add new entry**. 5. In the Value box, enter **10.10.10.10**. 6. Click **OK**.	Define the static route and set the next-hop address: set route 192.168.47.5 next-hop 10.10.10.10
Set the preference for the 10.10.10.10 next hop to 7.	1. Under Preference, in the Metric value box, enter **7**. 2. Click **OK**.	Set the preference to 7: set route 192.168.47.5 next-hop 10.10.10.10 preference 7
Define the qualified next-hop address 10.10.10.7.	1. In the Qualified next hop field, click **Add new entry**. 2. In the Nexthop field, enter **10.10.10.7**. 3. Click **OK**.	Set the qualified-next-hop address: set route 192.168.47.5 qualified-next-hop 10.10.10.7
Set the preference for the 10.10.10.7 qualified next hop to 6.	1. Under Preference, in the Metric value box, enter **6**. 2. Click **OK**.	Set the preference to 6: set route 192.168.47.5 qualified-next-hop 10.10.10.7 preference 6

Controlling Static Routes in the Routing and Forwarding Tables

Static routes have a number of attributes that define how they are inserted and maintained in the routing and forwarding tables. To customize this behavior for the static route 192.168.47.5/32, perform these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 106.
3. If you are finished configuring static routes, commit the configuration.
4. Go on to one of the following procedures:

- To define default properties for static routes, see “Defining Default Behavior for All Static Routes” on page 294.
- To check the configuration, see “Verifying the Static Route Configuration” on page 295.

Table 106: Controlling Static Routes in the Routing and Forwarding Tables

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **192.168.47.5/32** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing options > Static**, then click **192.168.47.5/32** in the Destination field.	From the top of the configuration hierarchy, enter edit routing-options static route 192.168.47.5/32
Specify that the route is to be retained in the forwarding table after the routing process shuts down. By default, static routes are not retained.	1. Next to Retain, select the **Yes** check box. 2. Click **OK**.	Set the retain attribute: set retain
Specify that the static route is not to be readvertised. By default, static routes are eligible to be readvertised.	1. Next to Readvertise, select the **No** check box. 2. Click **OK**.	Set the no-readvertise attribute: set no-readvertise
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	1. From the Passive flag list, select **Passive**. 2. Click **OK**.	Set the passive attribute: set passive

Defining Default Behavior for All Static Routes

Attributes that define static route behavior can be configured either at the individual route level or as a default behavior that applies to all static routes. In the case of conflicting configuration, the configuration at the individual route level overrides static route defaults. To configure static route defaults, perform these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 107.
3. If you are finished configuring static routes, commit the configuration.
4. To check the configuration, see “Verifying the Static Route Configuration” on page 295.

Table 107: Defining Static Route Defaults

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Defaults** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Static**, and then click **Configure** next to Defaults.	From the top of the configuration hierarchy, enter edit routing-options static defaults
Specify that the route is to be retained in the forwarding table after the routing process shuts down. By default, static routes are not retained.	1. Next to Retain, select the **Yes** check box. 2. Click **OK**.	Set the retain attribute: set retain
Specify that the static route is not to be readvertised. By default, static routes are eligible to be readvertised.	1. Next to Readvertise, select the **No** check box. 2. Click **OK**.	Set the no-readvertise attribute: set no-readvertise
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	1. From the Passive flag list, select **Passive**. 2. Click **OK**.	Set the passive attribute: set passive

Verifying the Static Route Configuration

Verify that the static routes are in the routing table and that those routes are active.

Displaying the Routing Table

Purpose Verify static route configuration as follows by displaying the routing table and checking its contents.

Action From the CLI, enter the show route terse command.

Sample Output

```
user@host> show route terse

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 192.168.47.5/32        S   5                       Reject
* 172.16.0.0/12      S   5                      >192.168.71.254
* 192.168.0.0/18     S   5                      >192.168.71.254
* 192.168.40.0/22    S   5                      >192.168.71.254
* 192.168.64.0/18    S   5                      >192.168.71.254
* 192.168.64.0/21    D   0                      >fxp0.0
* 192.168.71.246/32  L   0                       Local
* 192.168.220.4/30   D   0                      >fe-0/0/1.0
* 192.168.220.5/32   L   0                       Local
* 192.168.220.8/30   D   0                      >fe-0/0/2.0
* 192.168.220.9/32   L   0                       Local
* 192.168.220.12/30  D   0                      >fe-0/0/3.0
```

```
* 192.168.220.13/32  L   0                     Local
* 192.168.220.17/32  L   0                     Reject
* 192.168.220.21/32  L   0                     Reject
* 192.168.220.24/30  D   0                    >at-1/0/0.0
* 192.168.220.25/32  L   0                     Local
* 192.168.220.28/30  D   0                    >at-1/0/1.0
* 192.168.220.29/32  L   0                     Local
* 224.0.0.9/32       R 100          1          MultiRecv
```

What It Means The output shows a list of the routes that are currently in the inet.0 routing table. Verify the following information:

- Each configured static route is present. Routes are listed in ascending order by IP address. Static routes are identified with an S in the protocol (P) column of the output.
- Each static route is active. Routes that are active show the next-hop IP address in the Next hop column. If a route's next-hop address is unreachable, the next-hop address is identified as Reject. These routes are not active routes, but they appear in the routing table because the passive attribute is set.
- The preference for each static route is correct. The preference for a particular route is listed in the Prf column of the output.

Chapter 14
Configuring a RIP Network

The Routing Information Protocol (RIP) is an interior gateway protocol that routes packets within a single autonomous system (AS). To use RIP, you must understand the basic components of a RIP network and configure the J-series Services Router to act as a node in the network.

NOTE: The J-series Services Router supports only RIP version 1 and RIP version 2. Unless otherwise specified, the term *RIP* in this chapter refers to these versions of the protocol.

You can use either J-Web Quick Configuration or a configuration editor to configure a RIP network.

This chapter contains the following topics. For more information about RIP, see the *JUNOS Routing Protocols Configuration Guide*.

RIP Overview

To achieve basic connectivity between all RIP hosts in a RIP network, you need only enable RIP on every interface that is expected to transmit and receive RIP traffic. To do so, you define RIP groups, which are logical groupings of interfaces, and add interfaces to those groups. No additional configuration is required to pass traffic on a RIP network.

RIP Traffic Control with Metrics

To tune a RIP network and control traffic flowing through the network, you modify the incoming and outgoing metric attributes, which are set to 1 by default. These attributes manually specify the metric on any route that is advertised through that

host. By increasing or decreasing these metrics—and thus the cost—of links throughout the network, you can control packet transmission across the network.

Authentication

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. These authentication keys can be specified in either plain-text or MD5 form. Authentication provides an additional layer of security on the network beyond the other security features.

This type of authentication is not supported on RIPv1 networks.

Before You Begin

Before you begin configuring a RIP network, complete the following tasks:

- Establish basic connectivity. See “Establishing Basic Connectivity” on page 47.
- Configure network interfaces. See “Configuring Network Interfaces” on page 79.

Configuring a RIP Network with Quick Configuration

J-Web Quick Configuration allows you to create RIP networks. Figure 78 shows the Quick Configuration Routing page for RIP.

Figure 78: Quick Configuration Routing Page for RIP

To configure a RIP network with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > RIP Routing**.
2. Enter information into the Quick Configuration page for RIP, as described in Table 108.
3. From the main RIP routing Quick Configuration page, click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration Routing page for RIP, click **Apply**.
   - To apply the configuration and return to the Quick Configuration Routing page, click **OK**.
   - To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.
4. To check the configuration, see “Verifying the RIP Configuration” on page 307.

Table 108: RIP Routing Quick Configuration Summary

Field	Function	Your Action
RIP		
Enable RIP	Enables or disables RIP.	■ To enable RIP, select the check box. ■ To disable RIP, clear the check box.
Advertise Default Route	Advertises the default route using RIPv2.	■ To advertise the default route using RIPv2, select the check box. ■ To disable the default route advertisement, clear the check box.
RIP-Enabled Interfaces	Designates one or more Services Router interfaces on which RIP is enabled.	The first time you configure RIP, the Logical Interfaces box displays a list of all the logical interfaces configured on the Services Router. Do any of the following: ■ To enable RIP on an interface, click the interface name to highlight it, and click the left arrow to add the interface to the RIP interfaces list. ■ To enable RIP on multiple interfaces at once, press Ctrl while you click multiple interface names to highlight them. Then click the left arrow to add the interfaces to the RIP interfaces list. ■ To enable RIP on all logical interfaces except the special fxp0 management interface, select **All Interfaces** in the Logical Interfaces list and click the left arrow. ■ To enable RIP on all the interfaces displayed in the Logical Interfaces list, click **All** to highlight every interface. Then click the left arrow to add the interfaces to the RIP interfaces list. ■ To disable RIP on one or more interfaces, highlight the interface(s) in the RIP interfaces box and click the right arrow to move them back to the Logical Interfaces list.

Configuring a RIP Network with a Configuration Editor

To configure the Services Router as a node in a RIP network, you must perform the following task marked *(Required)*.

- (Required) “Configuring a Basic RIP Network” on page 301
- (Optional) “Controlling Traffic in a RIP Network” on page 302
- (Optional) “Enabling Authentication for RIP Exchanges” on page 305

For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 127.

Configuring a Basic RIP Network

To use RIP on the Services Router, you must configure RIP on all the RIP interfaces within a network like the one shown in Figure 79.

Figure 79: Typical RIP Network Topology

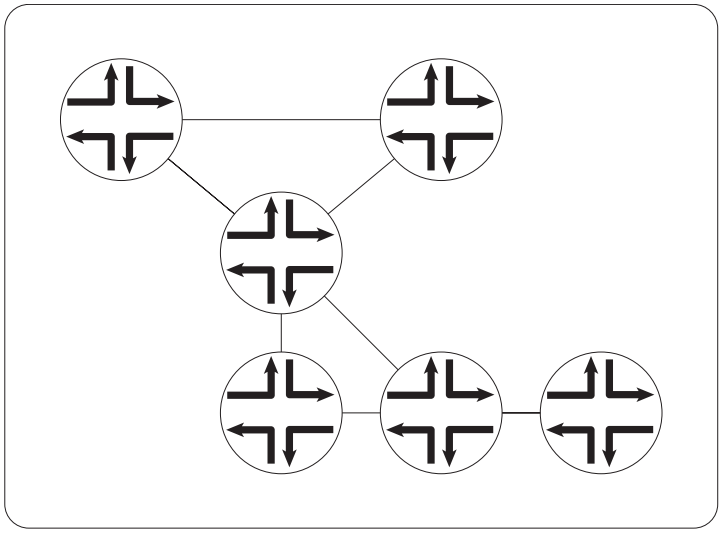

To configure a RIP network like the one in Figure 79, perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 109.
3. If you are finished configuring the network, commit the configuration.

After you add the appropriate interfaces to the RIP group, RIP begins sending routing information. No additional configuration is required to enable RIP traffic on the network.

4. Go on to one of the following procedures:

   - To control RIP traffic on the network, see "Controlling Traffic in a RIP Network" on page 302.
   - To authenticate RIP exchanges, see "Enabling Authentication for RIP Exchanges" on page 305.
   - To check the configuration, see "Verifying the RIP Configuration" on page 307.

Table 109: Configuring a RIP Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Rip** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Rip**.	From the top of the configuration hierarchy, enter `edit protocols rip`
Create the RIP group alpha1.	1. In the Group field, click **Add new entry**. 2. In the Group name box, type alpha1.	1. Create the RIP group alpha1, and add an interface: `set group alpha1 neighbor fe-0/0/0.0`
Add interfaces to the RIP group alpha1.	1. In the Neighbor field, click **Add new entry**. 2. In the Neighbor name box, type the name of an interface on the Services Router—for example, fe-0/0/0.0—and click **OK**. 3. Repeat Step 2 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.	2. Repeat Step 1 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.

Controlling Traffic in a RIP Network

There are two primary means for controlling traffic in a RIP network: the incoming metric and the outgoing metric. To modify these attributes, see the following sections:

- Controlling Traffic with the Incoming Metric on page 303
- Controlling Traffic with the Outgoing Metric on page 304

Controlling Traffic with the Incoming Metric

Depending on the RIP network topology and the links between nodes in the network, you might want to control traffic flow through the network to maximize flow across higher-bandwidth links. Figure 80 shows a network with alternate routes between routers A and D.

Figure 80: Controlling Traffic in a RIP Network with the Incoming Metric

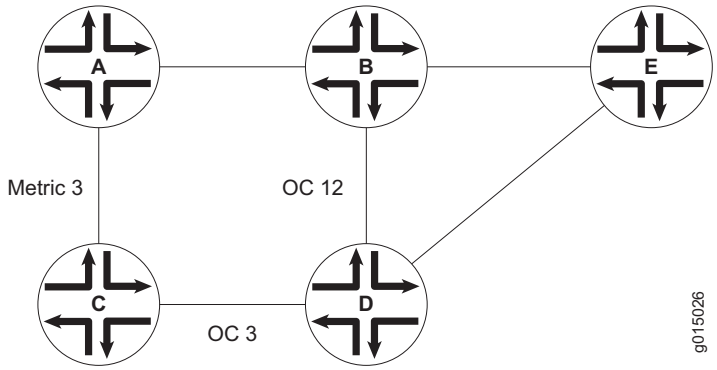

In this example, each of the two routes from router A to router D has two hops. However, because the link from router B to router D has a higher bandwidth than the link from router C to D, you want traffic to flow from router A through B to D. To force this flow, you can increase the incoming metric on router C from `1` (the default) to `3` to make this route less preferable.

To modify the incoming metric on router C and force traffic through router D:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 110.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges” on page 305.
   - To check the configuration, see “Verifying the RIP Configuration” on page 307.

Table 110: Modifying the Incoming Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
In the configuration hierarchy, navigate to the level of an interface in the alpha1 RIP group.	1. In the configuration editor hierarchy, select **Protocols > Rip**, and click **alpha1** in the Group name field. 2. Click the interface name—for example, **fe-0/0/0.0**—in the Neighbor name field.	From the top of the configuration hierarchy, enter edit protocols rip group alpha1 neighbor fe-0/0/0
Increase the incoming metric to 3.	In the Metric in box, type 3, and click **OK**.	Set the incoming metric to 3: set metric-in 3

Controlling Traffic with the Outgoing Metric

If a route being exported was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the outgoing metric. Figure 81 shows a network with alternate routes between routers A and D.

Figure 81: Controlling Traffic in a RIP Network with the Outgoing Metric

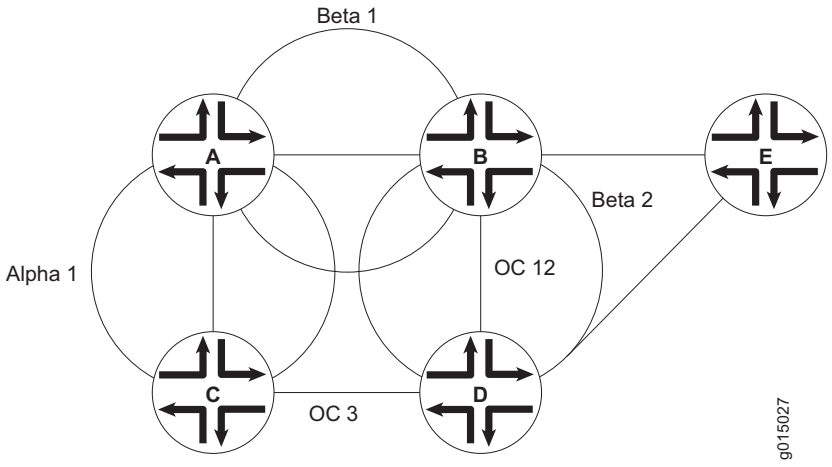

In this example, each of the two routes from router A to router D has two hops. However, because the link from router B to router D has a higher bandwidth than the link from router C to D, you want traffic to flow from router A through B to D. In this case, the A-to-B link is in a different RIP group from the B-to-D

link. As a result, the incoming metric is not sufficient to control traffic flow. To force traffic through router B, you can increase the outgoing metric on router C to make the route through C less preferable.

To modify the outgoing metric on router C and force traffic through router D:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 111.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To authenticate RIP exchanges, see "Enabling Authentication for RIP Exchanges" on page 305.
   - To check the configuration, see "Verifying the RIP Configuration" on page 307.

Table 111: Modifying the Outgoing Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **alpha1** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Rip**, and then click **alpha1** in the Group name field.	From the top of the configuration hierarchy, enter `edit protocols rip group alpha1`
Increase the outgoing metric to 5.	In the Metric out box, type 5, and click **OK**.	Set the outgoing metric to 5: `set metric-out 5`

Enabling Authentication for RIP Exchanges

All RIPv2 protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, this authentication is disabled. Authentication is performed at the area level, requiring all routers within the area to have the same authentication and corresponding key configured.

You can enable RIP authentication exchanges by either of the following methods:

- Enabling Authentication with Plain-Text Passwords on page 306
- Enabling Authentication with MD5 Authentication on page 306

Enabling Authentication with Plain-Text Passwords

To configure authentication that requires a plain-text password to be included in the transmitted packet, enable simple authentication by performing these steps on all RIP Services Routers in the area:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 112.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see "Verifying the RIP Configuration" on page 307.

Table 112: Configuring Simple RIP Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to **Rip** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Rip**.	From the top of the configuration hierarchy, enter edit protocols rip
Set the authentication type to simple.	From the Authentication type list, select **simple**.	Set the authentication type to simple: set authentication-type simple
Set the authentication key to a simple-text password. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.	In the Authentication key box, type a simple-text password, and click **OK**.	Set the authentication key to a simple-text password: set authentication-key *password*

Enabling Authentication with MD5 Authentication

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP Services Routers in the area:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 113.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see "Verifying the RIP Configuration" on page 307.

Table 113: Configuring MD5 RIP Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to **Rip** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Rip**.	From the top of the configuration hierarchy, enter edit protocols rip
Set the authentication type to MD5.	From the Authentication type list, select **md5**.	Set the authentication type to md5: set authentication-type md5
Set the MD5 authentication key (password). The key can be from 1 through 16 contiguous characters long and can include any ASCII strings.	In the Authentication key box, type an MD5 authentication key, and click **OK**.	Set the MD5 authentication key: set authentication-key *password*

Verifying the RIP Configuration

To verify the RIP configuration, perform these tasks:

- Verifying the RIP-Enabled Interfaces on page 307
- Verifying Reachability of All Hosts in the RIP Network on page 308

Verifying the RIP-Enabled Interfaces

Purpose Verify that all the RIP-enabled interfaces are available and active.

Action From the CLI, enter the show rip neighbor command.

Sample Output

```
user@host> show rip neighbor

Source           Destination     Send   Receive   In
Neighbor           State  Address         Address         Mode   Mode     Met
--------           -----  -------         -----------     ----   -------  ---
fe-0/0/0.0            Dn (null)          (null)          mcast  both       1
fe-0/0/1.0            Up 192.168.220.5   224.0.0.9       mcast  both       1
```

What It Means The output shows a list of the RIP neighbors that are configured on the Services Router. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the Destination State column. A state of Up indicates that the link is passing RIP traffic. A state of Dn indicates that the link is not passing RIP traffic. In a

point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

Verifying Reachability of All Hosts in the RIP Network

Purpose By using the traceroute tool on each loopback address in the network, verify that all hosts in the RIP network are reachable from each Services Router.

Action For each Services Router in the RIP network:

1. In the J-Web interface, select **Diagnose > Traceroute**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router.
3. Click **Start**. Output appears on a separate page.

Sample Output

```
1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 routera-fxp0.englab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms
```

What It Means Each numbered row in the output indicates a router (“hop”) in the path to the host. The three time increments indicate the round-trip time (RTT) between the Services Router and the hop, for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is likely not reachable.

For information about the traceroute command and its output, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Chapter 15
Configuring an OSPF Network

The Open Shortest Path First protocol (OSPF) is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). To use OSPF, you must understand the basic components of an OSPF network and configure the J-series Services Router to act as a node in the network.

NOTE: The J-series Services Router supports both OSPF version 2 and OSPF version 3. In this chapter, the term *OSPF* refers to both versions of the protocol.

You can use either J-Web Quick Configuration or a configuration editor to configure an OSPF network.

This chapter contains the following topics. For more information about OSPF, see the *JUNOS Routing Protocols Configuration Guide*.

OSPF Overview

In an OSPF network, the network topology is distributed among the nodes of the autonomous system (AS) and is regularly updated. Because topology changes are flooded throughout the network, every node maintains the same copy of the network map in its local topological database. Packets are then routed based on the shared topology.

Enabling OSPF

To activate OSPF on a network, you must enable the protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF on

one or more interfaces, you must configure one or more interfaces on the Services Router within an OSPF area. Once the interfaces are configured, OSPF link-state advertisements (LSAs) are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

OSPF Areas

OSPF is enabled on a per-interface basis. Those interfaces are configured as OSPF enabled, and are assigned to an area. In a simple, single-area network, the area has the numeric identifier 0.0.0.0, which designates it as the backbone area. As the network grows, it is divided into multiple subnetworks or areas that are identified by numeric identifiers unique to the AS.

In a multiarea network, all areas must be directly connected to the backbone area by area border routers (ABRs). Because all areas are adjacent to the backbone area, OSPF routers send all traffic not destined for their own area through the backbone area. The ABRs in the backbone area are then responsible for transmitting the traffic through the appropriate ABR to the destination area.

Path Cost Metrics

Once the topology is shared across the network, OSPF uses it to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest path first (SPF) algorithm. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion.

OSPF allows you to manually assign a cost (or metric) to a particular path segment to control the flow of packets across the network.

Before You Begin

Before you begin configuring an OSPF network, complete the following tasks:

- Establish basic connectivity. See “Establishing Basic Connectivity” on page 47.
- Configure network interfaces. See “Configuring Network Interfaces” on page 79.

Configuring an OSPF Network with Quick Configuration

J-Web Quick Configuration allows you to create single-area OSPF networks. Figure 82 shows the Quick Configuration Routing page for OSPF.

Figure 82: Quick Configuration Routing Page for OSPF

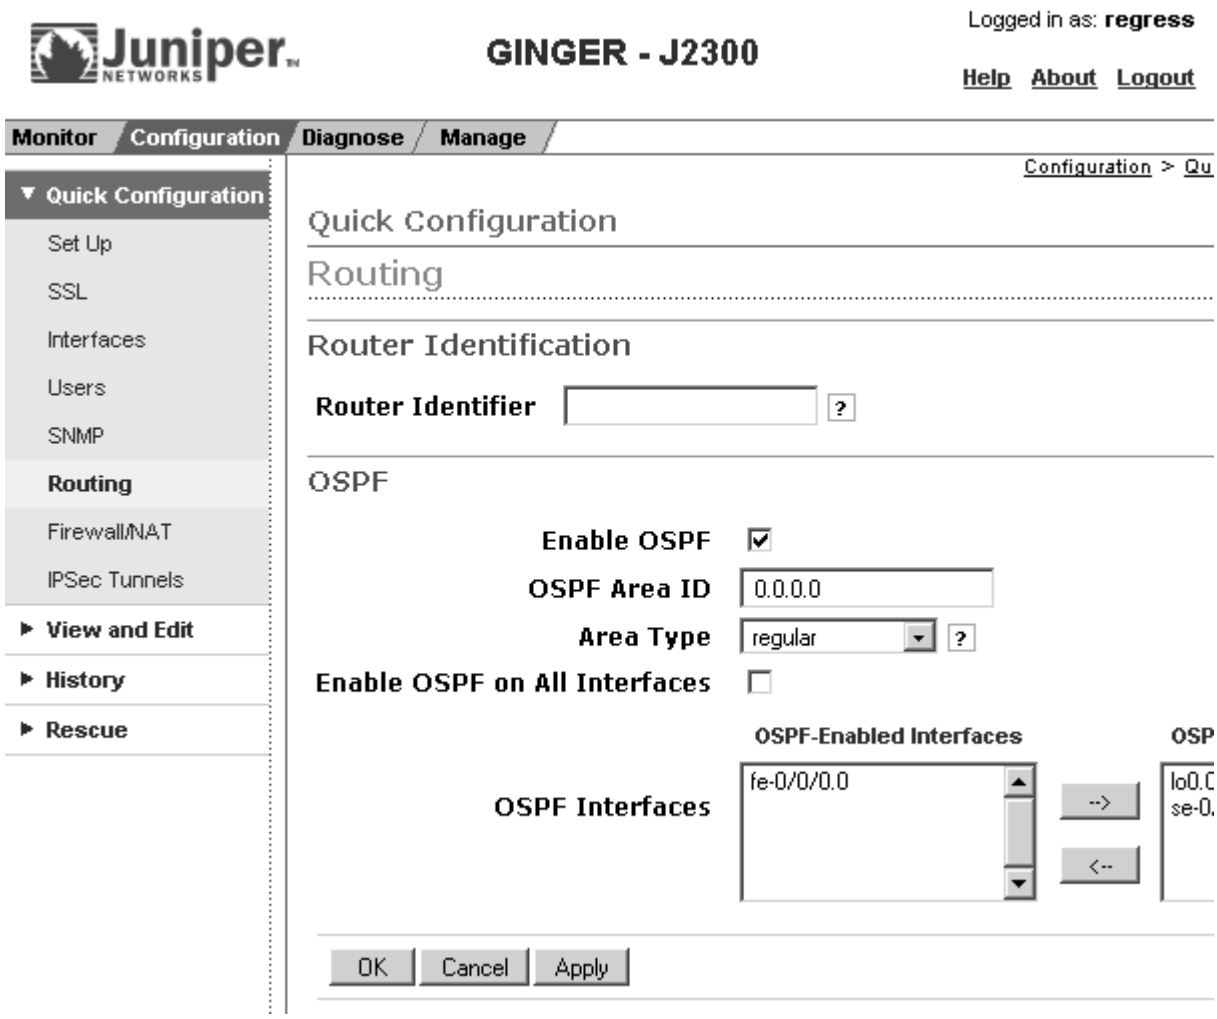

To configure a single-area OSPF network with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > OSPF Routing**.
2. Enter information into the Quick Configuration Routing page for OSPF, as described in Table 114.
3. Click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration Routing page for OSPF, click **Apply**.
   - To apply the configuration and return to the Quick Configuration Routing page, click **OK**.

- To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.

4. To check the configuration, see "Verifying an OSPF Configuration" on page 325.

Table 114: OSPF Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router.	Type the Services Router's 32-bit IP address, in dotted decimal notation.
OSPF		
Enable OSPF	Enables or disables OSPF.	■ To enable OSPF, select the check box. ■ To disable OSPF, clear the check box.
OSPF Area ID	Uniquely identifies the area within its AS.	Type a 32-bit numeric identifier for the area, or an integer. If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the area is 0.0.0.3.

Field	Function	Your Action
Area Type	Designates the type of OSPF area.	From the drop-down list, select the type of OSPF area you are creating: ■ **regular**—A regular OSPF area, including the backbone area ■ **stub**—A stub area ■ **nssa**—A not-so-stubby area (NSSA)
OSPF-Enabled Interfaces	Designates one or more Services Router interfaces on which OSPF is enabled.	The first time you configure OSPF, the Logical Interfaces box displays a list of all the logical interfaces configured on the Services Router. Do any of the following: ■ To enable OSPF on an interface, click the interface name to highlight it, and click the left arrow to add the interface to the OSPF interfaces list. ■ To enable OSPF on multiple interfaces at once, press Ctrl while you click multiple interface names to highlight them. Then click the left arrow to add the interfaces to the OSPF interfaces list. ■ To enable OSPF on all logical interfaces except the special fxp0 management interface, select **All Interfaces** in the Logical Interfaces list and click the left arrow. ■ To enable OSPF on all the interfaces displayed in the Logical Interfaces list, click **All** to highlight every interface. Then click the left arrow to add the interfaces to the OSPF interfaces list. ■ To disable OSPF on one or more interfaces, highlight the interface or interfaces in the OSPF interfaces box and click the right arrow to move them back to the Logical Interfaces list.

Configuring an OSPF Network with a Configuration Editor

To configure the Services Router as a node in an OSPF network, you must perform the following tasks marked *(Required)*.

- (Required) "Configuring the Router Identifier" on page 314Configuring the Router Identifier on page 9
- (Required) "Configuring a Single-Area OSPF Network" on page 315
- (Optional) "Configuring a Multiarea OSPF Network" on page 316
- (Optional) "Configuring Stub and Not-So-Stubby Areas" on page 319

For information about using the J-Web and CLI configuration editors, see "Using J-series Configuration Tools" on page 127.

Configuring the Router Identifier

The router identifier is the IP address that uniquely identifies the J-series Services Router.

OSPF uses the router identifier to elect a designated router, unless you manually specify a priority value. When the OSPF network first becomes active, by default, the router with the highest router identifier is elected the designated router.

To configure the router identifier for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 115.
3. Go on to "Configuring a Single-Area OSPF Network" on page 315.

Table 115: Configuring the Router Identifier

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Routing-options** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing-options**.	From the top of the configuration hierarchy, enter edit routing-options
Enter the router ID value.	In the Router Id box, type the IP address of the Services Router, in dotted decimal notation.	Set the router-id value to the IP address of the Services Router, in dotted decimal notation. For example: set router-id 177.162.4.24
Apply your configuration changes.	Click **OK** to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the set command.

Configuring a Single-Area OSPF Network

To use OSPF on the Services Router, you must configure at least one OSPF area, like the one shown in Figure 83.

Figure 83: Typical Single-Area OSPF Network Topology

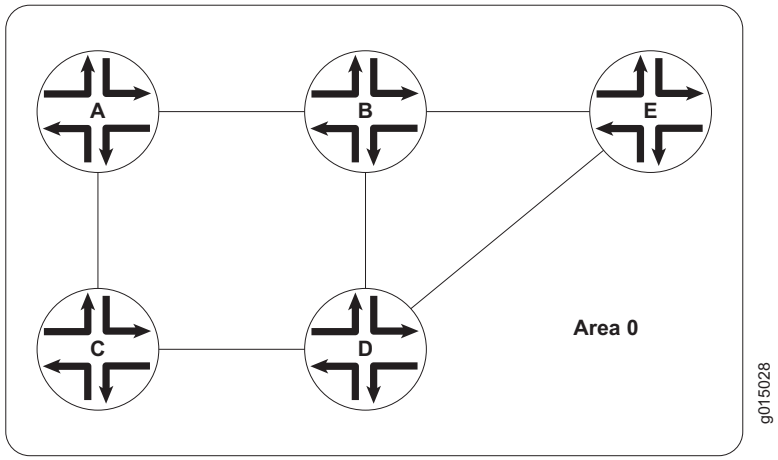

To configure a single-area OSPF network with a backbone area, like the one in Figure 83, perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 116.
3. If you are finished configuring the network, commit the configuration.

   After you create the backbone area and add the appropriate interfaces to the area, OSPF begins sending LSAs. No additional configuration is required to enable OSPF traffic on the network.
4. Go on to one of the following procedures:

   - To add more areas to the AS, see "Configuring a Multiarea OSPF Network" on page 316.
   - To control external route advertisement in the AS, see "Configuring Stub and Not-So-Stubby Areas" on page 319.
   - To improve network operation, see "Tuning an OSPF Network for Efficient Operation" on page 321.
   - To check the configuration, see "Verifying an OSPF Configuration" on page 325.

Table 116: Configuring a Single-Area OSPF Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Ospf** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf**.	From the top of the configuration hierarchy, enter edit protocols ospf
Create the backbone area with area ID 0.0.0.0.	1. In the Area box, click **Add new entry**. 2. In the Area ID box, type 0.0.0.0.	1. Set the backbone area ID to 0.0.0.0 and add an interface. For example: set area 0.0.0.0 interface fe-0/0/0
Add interfaces as needed to the OSPF area.	1. In the Interface box, click **Add new entry**. 2. In the Interface name box, type the name of an interface on the Services Router and click **OK**. 3. Repeat Step 1 and Step 2 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.	2. Repeat Step 1 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required. Changes in the CLI are applied automatically when you execute the set command.

Configuring a Multiarea OSPF Network

To reduce traffic and topology maintenance for the Services Routers in an OSPF autonomous system (AS), you can group them into multiple areas, as shown in Figure 84.

Figure 84: Typical Multiarea OSPF Network Topology

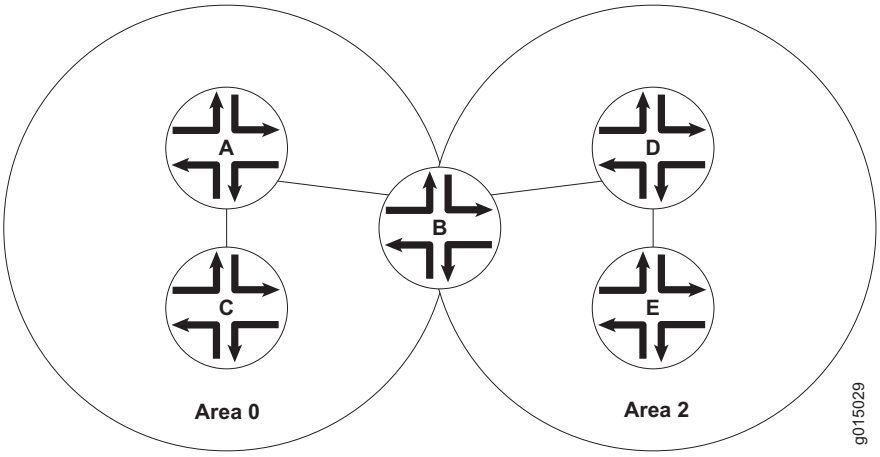

To configure a multiarea OSPF network shown in Figure 84, perform the following tasks on the appropriate Services Routers in the network. You must create a

backbone area. To link each additional area to the backbone area, you must configure one of the Services Routers as an area border router (ABR).

- "Creating the Backbone Area" on page 317
- "Creating Additional OSPF Areas" on page 317
- "Configuring Area Border Routers" on page 318

Creating the Backbone Area

On each Services Router that is to operate as an ABR in the network, create backbone area 0.0.0.0 with at least one interface enabled for OSPF.

For instruction, see "Configuring a Single-Area OSPF Network" on page 315.

Creating Additional OSPF Areas

To create additional OSPF areas:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 117.
3. If you are finished configuring the network, commit the configuration.

Table 117: Configuring a Multiarea OSPF Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Ospf** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf**.	From the top of the configuration hierarchy, enter edit protocols ospf
Create the additional area with a unique area ID, in dotted decimal notation.	1. In the Area box, click **Add new entry**. 2. In the Area ID box, type 0.0.0.2.	1. Set the area ID to 0.0.0.2 and add an interface. For example: set area 0.0.0.2 interface fe-0/0/0
Add interfaces as needed to the OSPF area.	1. In the Interface box, click **Add new entry**. 2. In the Interface name box, type the name of an interface on the Services Router and click **OK**. 3. Repeat Step 1 and Step 2 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.	2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required. Changes in the CLI are applied automatically when you execute the set command.

Configuring Area Border Routers

A Services Router operating as an area border router (ABR) has interfaces enabled for OSPF in the backbone area and in the area you are linking to the backbone. For example, Services Router B acts as the ABR in Figure 84 and has interfaces in both the backbone area and area 0.0.0.3.

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 118.
3. If you are finished configuring the network, commit the configuration.

   After you create the areas on the appropriate Services Routers and add and enable the appropriate interfaces to the areas, no additional configuration is required to enable OSPF traffic within or across the areas.
4. Go on to one of the following procedures:

   - To control external route advertisement in the AS, see "Configuring Stub and Not-So-Stubby Areas" on page 319.
   - To improve network operation, see "Tuning an OSPF Network for Efficient Operation" on page 321.
   - To check the configuration, see "Verifying an OSPF Configuration" on page 325.

Table 118: Configuring Area Border Routers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Ospf** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf**.	From the top of the configuration hierarchy, enter edit protocols ospf
Verify that the backbone area has at least one interface enabled for OSPF.	Click **0.0.0.0** to display the Area ID 0.0.0.0 page, and verify that the backbone area has at least one interface enabled for OSPF. For example, Services Router B in Figure 84 has the following interfaces enabled for OSPF in the backbone area: ■ Interface fe-0/0/0.0 ■ Interface fe-0/0/1.0 To enable an interface on the backbone area, see "Configuring a Single-Area OSPF Network" on page 315.	View the configuration using the show command: show For example, Services Router B in Figure 84 has the following interfaces enabled for OSPF in the backbone area: area 0.0.0.0 { interface fe-0/0/0.0; interface fe-0/0/1.0; } To enable an interface on the backbone area, see "Configuring a Single-Area OSPF Network" on page 315.

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the additional area with a unique area ID, in dotted decimal format.	1. In the Area box, click **Add new entry**. 2. In the Area ID box, type 0.0.0.2.	1. Set the area ID to 0.0.0.2 and add an interface. For example: set area 0.0.0.2 interface fe-0/0/0 2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required. Changes in the CLI are applied automatically when you execute the set command.
Add interfaces as needed to the OSPF area.	1. In the Interface box, click **Add new entry**. 2. In the Interface name box, type the name of an interface on the Services Router and click **OK**. 3. Repeat Step 1 and Step 2 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.	

Configuring Stub and Not-So-Stubby Areas

To control the advertisement of external routes into an area, you can create stub areas and not-so-stubby areas (NSSAs) in an OSPF network. In the network shown in Figure 85, area 0.0.0.7 has no external connections and can be configured as a stub area. Area 0.0.0.9 only has external connections to static routes and can be configured as an NSSA.

Figure 85: OSPF Network Topology with Stub Areas and NSSAs

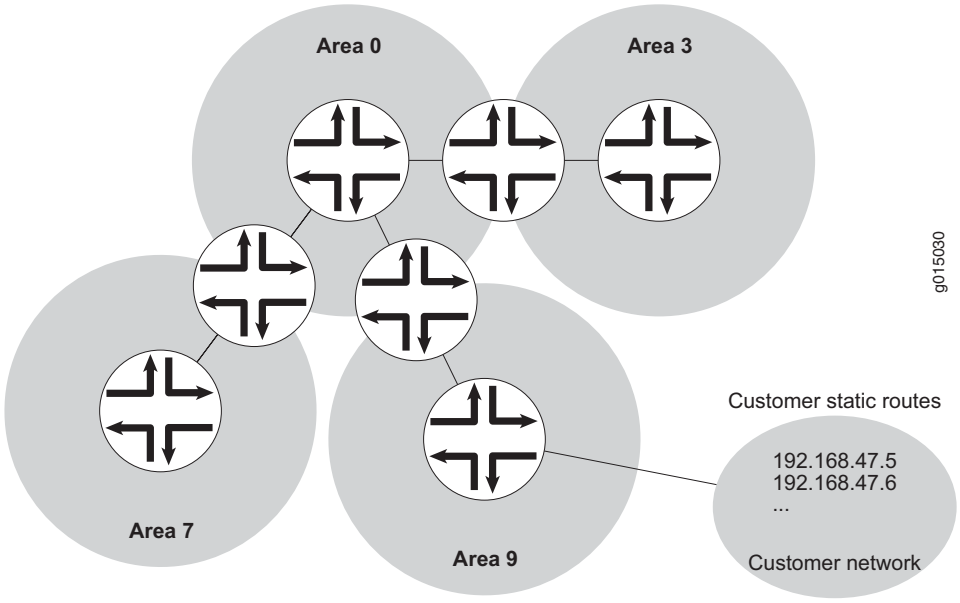

To configure stub areas and NSSAs in an OSPF network like the one shown in Figure 85:

1. Create the area and enable OSPF on the interfaces within that area.

   For instructions, see "Creating Additional OSPF Areas" on page 317.

2. Configure an area border router to bridge the areas.

   For instructions, see "Configuring Area Border Routers" on page 318.

3. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.

4. To configure each Services Router in area 0.0.0.7 as a stub area router, perform the configuration tasks described in Table 119.

5. If you are finished configuring the network, commit the configuration.

6. Go on to one of the following procedures:

   - To improve network operation, see "Tuning an OSPF Network for Efficient Operation" on page 321.
   - To check the configuration, see "Verifying an OSPF Configuration" on page 325.

Table 119: Configuring Stub Area and Not-So-Stubby Area Routers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **0.0.0.7** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf > Area id 0.0.0.7**.	From the top of the configuration hierarchy, enter `edit protocols ospf area 0.0.0.7`
Configure each Services Router in area 0.0.0.7 as a stub router.	1. In the Stub option list, select **Stub** and click **OK**. 2. Repeat Step 1 for every Services Router in the stub area to configure them with the stub parameter for the area.	1. Set the stub attribute: `set stub` 2. Repeat Step 1 for every Services Router in the stub area to configure them with the stub parameter for the area.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **0.0.0.9** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf > Area > 0.0.0.9**.	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.9
Configure each Services Router in area 0.0.0.9 as an NSSA router.	1. In the Stub option list, select **Nssa** and click **OK**. 2. Repeat Step 1 for every Services Router in the NSSA to configure them with the nssa parameter for the area.	1. Set the nssa attribute: set nssa 2. Repeat Step 1 for every Services Router in the NSSA to configure them with the nssa parameter for the area. Changes in the CLI are applied automatically when you execute the set command.

Tuning an OSPF Network for Efficient Operation

To make your OSPF network operate more efficiently, you can change some default settings on the Services Router by performing the following tasks:

- "Controlling Route Selection in the Forwarding Table" on page 321
- "Controlling the Cost of Individual Network Segments" on page 322
- "Enabling Authentication for OSPF Exchanges" on page 323
- "Controlling Designated Router Election" on page 324

Controlling Route Selection in the Forwarding Table

OSPF uses route preferences to select the route that is installed in the forwarding table when several routes have the same shortest path first (SFP) calculation. To evaluate a route, OSPF calculates the sum of the individual preferences of every router along the path and selects the route with the lowest total preference.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. Suppose all routers in your OSPF network use the default preference values. By setting the internal preference to 7 and the external preference to 130, you can ensure that the path through a particular Services Router is selected for the forwarding table any time multiple equal-cost paths to a destination exist.

To modify the default preferences on a Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 120.

Table 120: Controlling Route Selection in the Forwarding Table by Setting Preferences

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Ospf** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf**.	From the top of the configuration hierarchy, enter edit protocols ospf
Set the external and internal route preferences.	1. In the External preference box, type an external preference value—for example, 7. 2. In the Preference box, type an internal preference value—for example, 130. 3. Click **OK**.	1. Set the internal preference. For example: set preference 7 2. Set the external preference. For example: set external-preference 130 Changes in the CLI are applied automatically when you execute the set command.

Controlling the Cost of Individual Network Segments

When evaluating the cost of individual network segments, OSPF evaluates the reference bandwidth. For any link faster than 100 Mbps, the default cost metric is 1. When OSPF calculates the SPF algorithm, it sums the metrics of all interfaces along a path to determine the overall cost of the path. The path with the lowest metric is selected for the forwarding table.

To control the cost of the network segment, you can modify the metric value on an individual interface. Suppose all routers in the OSPF network use default metric values. If you increase the metric on an interface to 5, all paths through this interface have a calculated metric higher than the default and are *not* preferred.

To manually set the cost of a network segment on the stub area's Fast Ethernet interface by modifying the interface metric:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 121.

Table 121: Controlling the Cost of Individual Network Segments by Modifying the Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **fe-0/0/0.0** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf > Area id 0.0.0.0 > Interface name fe-0/0/0.0**.	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.0 interface fe-0/0/0.0
Set the interface metric and the external and route preference.	1. In the Metric box, type an interface metric value—for example, 5. 2. Click **OK**.	1. Set the interface metric. For example: set metric 5 2. Set the external preference. For example: set external-preference 130 Changes in the CLI are applied automatically when you execute the set command.

Enabling Authentication for OSPF Exchanges

All OSPFv2 protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, OSPF authentication is disabled.

NOTE: OSPFv3 does not support authentication.

You can enable either of two authentication types:

- Simple authentication—Authenticates by means of a plain-text password (key) included in the transmitted packet.
- MD5 authentication—Authenticates by means of an MD5 checksum included in the transmitted packet.

Because OSPF performs authentication at the area level, all routers within the area must have the same authentication and corresponding password (key) configured. For MD5 authentication to work, both the receiving and transmitting routers must have the same MD5 key.

To enable OSPF authentication on the stub area:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 122.

Table 122: Enabling OSPF Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **0.0.0.0** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf > Area id 0.0.0.0**.	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.0
Set the authentication type.	1. From the Authentication type list, select the type of authentication to enable on the stub area: **simple** **md5** 2. Click **OK**.	Set the authentication type to either simple or md5. For example: set authentication–type md5 Changes in the CLI are applied automatically when you execute the set command.
Navigate to the *interface-name* level in the configuration hierarchy.	In the configuration editor hierarchy under **Protocols > Ospf > Area > 0.0.0.0 > interface**, click an interface name.	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.0 interface *interface-name*
Set the authentication password (key) and, if applicable, the key identifier.	1. In the Key name box, type a password: For simple authentication, type from 1 through 8 ASCII characters. For MD5 authentication, type from 1 through 16 ASCII characters. 2. For MD5 authentication only, in the Key ID box, type any value between 0 (the default) and 255 to associate with the MD5 password. 3. Click **OK**. 4. Repeat Step 1 through Step 3 for each interface in the stub area for which you are enabling authentication.	1. Set the authentication password: For simple authentication, type from 1 through 8 ASCII characters. For MD5 authentication, type from 1 through 16 ASCII characters. 2. For MD5 authentication only, set the key identifier to associate with the MD5 password to any value between 0 (the default) and 255. For example: set authentication-key Chey3nne key-id 2 Changes in the CLI are applied automatically when you execute the command. 3. Repeat Step 1 and Step 2 for each interface in the stub area for which you are enabling authentication.

Controlling Designated Router Election

At designated router election, the router priorities are evaluated first, and the router with the highest priority is elected designated router.

By default, routers have a priority of 128. A priority of 0 marks the router as ineligible to become the designated router. To configure a router so it is always the designated router, set its priority to 255.

To change the priority of a Services Router to control designated router election:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 123.

Table 123: Controlling Designated Router Election

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the OSPF interface address for the Services Router. For example, navigate to the **fe-/0/0/1** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Ospf > area id 0.0.0.3 > Interface name fe-/0/0/1**.	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.3 interface fe-0/0/1
Set the Services Router priority.	1. In the Priority box, type a value between 0 and 255. The default value is 128. 2. Click **OK**.	Set the priority to a value between 0 and 255. The default value is 128. For example: set priority 200 Changes in the CLI are applied automatically when you execute the set command.

Verifying an OSPF Configuration

To verify an OSPF configuration, perform these tasks:

- "Verifying OSPF-Enabled Interfaces" on page 325
- "Verifying OSPF Neighbors" on page 326
- "Verifying the Number of OSPF Routes" on page 327
- "Verifying Reachability of All Hosts in an OSPF Network" on page 328

Verifying OSPF-Enabled Interfaces

Purpose Verify that OSPF is running on a particular interface and that the interface is in the desired area.

Action From the CLI, enter the show ospf interface command.

Sample Output

```
user@host> show ospf interface

Intf              State     Area            DR ID           BDR ID       Nbrs
```

```
at-5/1/0.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
ge-2/3/0.0          DR       0.0.0.0         192.168.4.16    192.168.4.15    1
lo0.0               DR       0.0.0.0         192.168.4.16    0.0.0.0         0
so-0/0/0.0          Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/1.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-6/0/2.0          Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/3.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
```

What It Means The output shows a list of the Services Router interfaces that are configured for OSPF. Verify the following information:

- Each interface on which OSPF is enabled is listed.
- Under Area, each interface shows the area for which it was configured.
- Under Intf and State, the Services Router loopback (lo0.0) interface and LAN interface that are linked to the OSPF network's designated router (DR) are identified.
- Under DR ID, the IP address of the OSPF network's designated router appears.
- Under State, each interface shows a state of PtToPt to indicate a point-to-point connection. If the state is Waiting, check the output again after several seconds. A state of Down indicates a problem.
- The designated router addresses always show a state of DR.

For more information about show ospf interface, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying OSPF Neighbors

Purpose OSPF neighbors are interfaces that have an immediate adjacency. On a point-to-point connection between the Services Router and another router running OSPF, verify that each router has a single OSPF neighbor.

Action From the CLI, enter the show ospf neighbor command.

Sample Output

```
user@host> show ospf neighbor

Address         Intf              State      ID              Pri  Dead
192.168.254.225  fxp3.0             2Way      10.250.240.32    128   36
192.168.254.230  fxp3.0             Full      10.250.240.8     128   38
192.168.254.229  fxp3.0             Full      10.250.240.35    128   33
10.1.1.129       fxp2.0             Full      10.250.240.12    128   37
10.1.1.131       fxp2.0             Full      10.250.240.11    128   38
10.1.2.1         fxp1.0             Full      10.250.240.9     128   32
10.1.2.81        fxp0.0             Full      10.250.240.10    128   33
```

What It Means The output shows a list of the Services Router's OSPF neighbors and their addresses, interfaces, states, router IDs, priorities, and number of seconds allowed for inactivity ("dead" time). Verify the following information:

- Each interface that is immediately adjacent to the Services Router is listed.
- The Services Router's own loopback address and the loopback addresses of any routers with which the Services Router has an immediate adjacency are listed.
- Under State, each neighbor shows a state of Full. Because full OSPF connectivity is established over a series of packet exchanges between clients, the OSPF link might take several seconds to establish. During that time, the state might be displayed as Attempt, Init, or 2way, depending on the stage of negotiation.

  If, after 30 seconds, the state is not Full, the OSPF configuration between the neighbors is not functioning correctly.

For more information about show ospf neighbor, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying the Number of OSPF Routes

Purpose Verify that the OSPF routing table has entries for the following:

- Each subnetwork reachable through an OSPF link
- Each loopback address reachable on the network

For example, Figure 86 shows a sample network with an OSPF topology.

Figure 86: Sample OSPF Network Topology

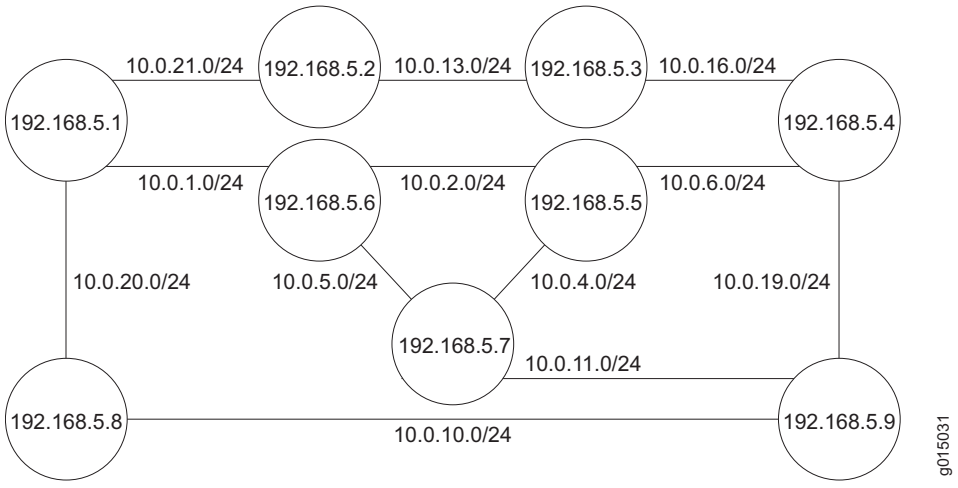

In this topology, OSPF is being run on all interfaces. Each segment in the network is identified by an address with a /24 prefix, with interfaces on either end of the segment being identified by unique IP addresses.

Action From the CLI, enter the show ospf route command.

Sample Output

```
user@host> show ospf route

Prefix              Path   Route       NH   Metric  NextHop       Nexthop
                    Type   Type        Type         Interface     addr/label
10.10.10.1/24       Intra  Network     IP   1       fe-0/0/2.0    10.0.21.1
10.10.10.2/24       Intra  Network     IP   1       fe-0/0/2.0    10.0.21.1
10.10.10.4/24       Intra  Network     IP   1       fe-0/0/1.0    10.0.13.1
10.10.10.5/24       Intra  Network     IP   1       fe-0/0/2.0    10.0.21.1
10.10.10.6/24       Intra  Network     IP   1       fe-0/0/1.0    10.0.13.1
10.10.10.10/24      Intra  Network     IP   1       fe-0/0/2.0    10.0.21.1
10.10.10.11/24      Intra  Network     IP   1       fe-0/0/1.0    10.0.13.1
10.10.10.13/24      Intra  Network     IP   1       fe-0/0/1.0
10.10.10.16/24      Intra  Network     IP   1       fe-0/0/1.0    10.0.13.1
10.10.10.19/24      Intra  Network     IP   1       fe-0/0/1.0    10.0.13.1
10.10.10.20/24      Intra  Network     IP   1       fe-0/0/2.0    10.0.21.1
10.10.10.21/24      Intra  Network     IP   1       fe-0/0/2.0
192.168.5.1         Intra  Router      IP   1       fe-0/0/2.0    10.0.21.1
192.168.5.2         Intra  Router      IP   1       lo0
192.168.5.3         Intra  Router      IP   1       fe-0/0/1.0    10.0.13.1
192.168.5.4         Intra  Router      IP   1       fe-0/0/1.0    10.0.13.1
192.168.5.5         Intra  Router      IP   1       fe-0/0/1.0    10.0.13.1
192.168.5.6         Intra  Router      IP   1       fe-0/0/2.0    10.0.21.1
192.168.5.7         Intra  Router      IP   1       fe-0/0/2.0    10.0.21.1
192.168.5.8         Intra  Router      IP   1       fe-0/0/2.0    10.0.21.1
192.168.5.9         Intra  Router      IP   1       fe-0/0/1.0    10.0.13.1
```

What It Means The output lists each route, sorted by IP address. Routes are shown with a route type of Network, and loopback addresses are shown with a route type of Router.

For the example shown in Figure 86, verify that the OSPF routing table has 21 entries, one for each network segment and one for each router's loopback address.

For more information about show ospf route, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying Reachability of All Hosts in an OSPF Network

Purpose By using the traceroute tool on each loopback address in the network, verify that all hosts in the network are reachable from each Services Router.

Action For each Services Router in the OSPF network:

1. In the J-Web interface, select **Diagnose > Traceroute**.
2. In the Host Name box, type the name of a host for which you want to verify reachability from the Services Router.

3. Click **Start**. Output appears on a separate page.

Sample Output

```
1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 routera-fxp0.englab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms
```

What It Means

Each numbered row in the output indicates a router (“hop”) in the path to the host. The three time increments indicate the round-trip time (RTT) between the Services Router and the hop, for each traceroute packet. To ensure that the OSPF network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is likely not reachable. In this case, verify the routes with the show ospf route command.

For information about ospf routeshow , see “Verifying the Number of OSPF Routes” on page 327.

For information about the traceroute command and its output, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Chapter 16
Configuring BGP Sessions

Connections between peering networks are typically made through an exterior gateway protocol, most commonly the Border Gateway Protocol (BGP).

You can use either J-Web Quick Configuration or a configuration editor to configure BGP sessions.

This chapter contains the following topics. For more information about BGP, see the *JUNOS Routing Protocols Configuration Guide*.

BGP Overview

BGP is a heavy-duty, secure protocol that must be configured on a per-peer basis. Once a peering session has been configured, BGP uses a TCP connection to establish a session. After a BGP session is established, traffic is passed along the BGP-enabled link.

Although BGP requires a full-mesh topology to share route information, you can use route reflectors and confederations in a large autonomous system (AS) to reduce scaling problems.

BGP Peering Sessions

Unlike RIP and OSPF links, BGP peering sessions must be explicitly configured at both ends. To establish a session between BGP peers, you must manually specify the interface address to which you are establishing a connection. Once this configuration is complete on both ends of a link, a TCP negotiation takes place and a BGP session is established.

The type of the BGP peering session depends on whether the peer is outside or inside the host's autonomous system (AS):

- Peering sessions established with hosts outside the local AS are external sessions . Traffic that passes along such links uses external BGP (EBGP) as its protocol.
- Peering sessions established with hosts within the local AS are internal sessions. Traffic that passes along such links uses internal BGP (IBGP) as its protocol.

IBGP Full Mesh Requirement

By default, BGP does not readvertise routes that are learned from BGP. To share route information throughout the network, BGP requires a full mesh of internal peering sessions within an AS. To achieve an IBGP full mesh, you configure a direct peering session every host to every other host within the network. These sessions are configured on every router within the network, as type internal.

Route Reflectors and Clusters

In larger networks, the overhead needed to implement the IBGP full-mesh requirement is prohibitive. Many networks use route reflectors to avoid having to configure an internal connection to each node for every new router.

A route reflector can readvertise routes learned through BGP to its BGP neighbors. If you define clusters of routers and configure a single router as a route reflector within each cluster, a full mesh is required only between the route reflectors and all their internal peers within the network. The route reflector is responsible for propagating BGP routes throughout the cluster.

For more information about route reflectors, see "Route Reflectors—for Added Hierarchy" on page 281

BGP Confederations

Large ASs can be divided into smaller sub-ASs, which are groups of routers known as confederations. You configure EBGP peering sessions between confederations, and IBGP peering sessions within confederations. Within a confederation, the IBGP full mesh is required. For more information about confederations, see "Confederations—for Subdivision" on page 283

Before You Begin

Before you begin configuring a BGP network, complete the following tasks:

- Establish basic connectivity. See "Establishing Basic Connectivity" on page 47.
- Configure network interfaces. See "Configuring Network Interfaces" on page 79.

Configuring a BGP Network with Quick Configuration

J-Web Quick Configuration allows you to create BGP peering sessions. Figure 87 shows the Quick Configuration Routing page for BGP.

Figure 87: Quick Configuration Routing Page for BGP

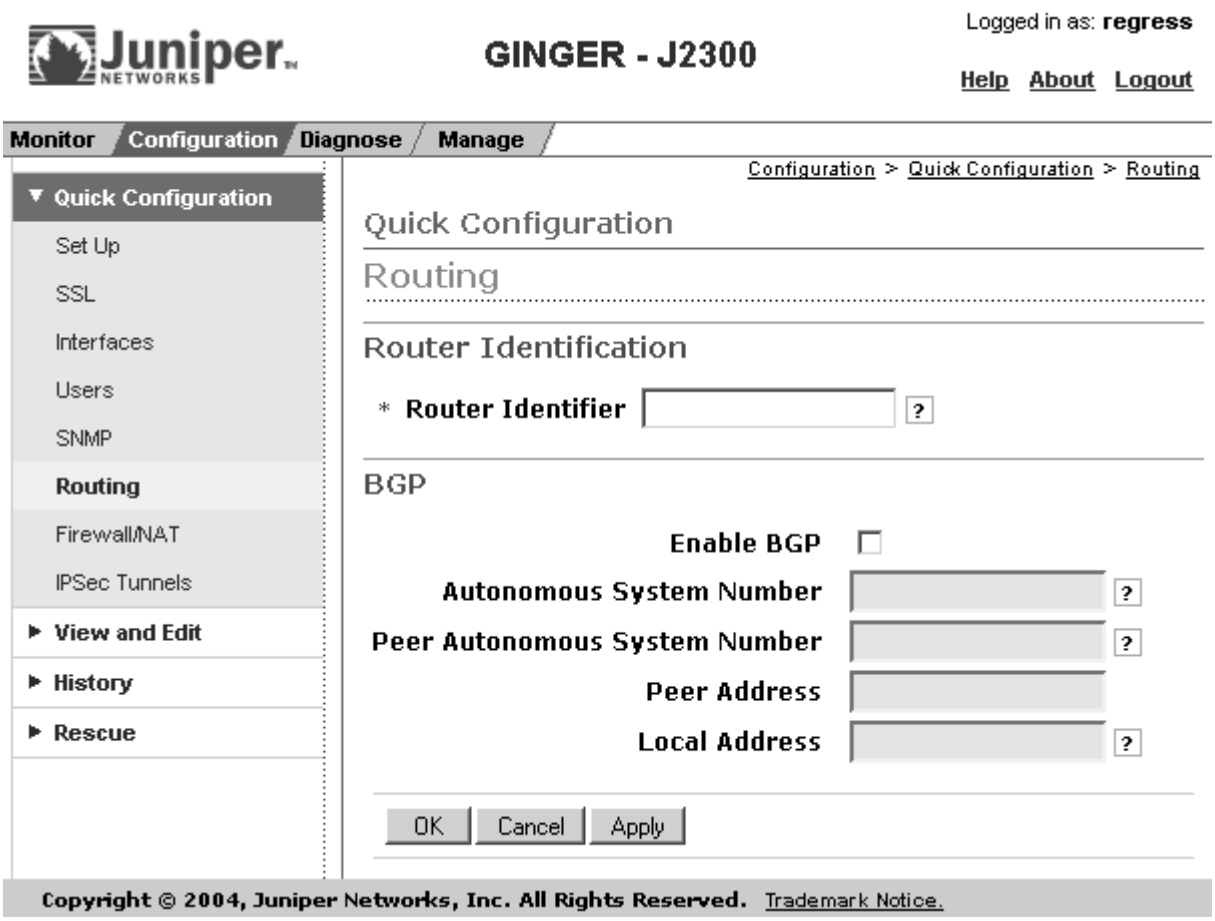

To configure a BGP peering session with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > BGP Routing**.
2. Enter information into the Quick Configuration page for BGP, as described in Table 124.
3. From the main BGP routing Quick Configuration page, click one of the following buttons:
   - To apply the configuration and stay on the Quick Configuration Routing page for BGP, click **Apply**.

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

4. To check the configuration, see "Verifying a BGP Configuration" on page 344.

Table 124: BGP Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router	Type the Services Router's 32-bit IP address, in dotted decimal notation.
BGP		
Enable BGP	Enables or disables BGP.	■ To enable BGP, select the check box. ■ To disable BGP, clear the check box.
Autonomous System Number	Sets the unique numeric identifier of the AS in which the services router is configured.	Type the Services Router's 32-bit AS number, in dotted decimal notation. If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the AS is 0.0.0.3.
Peer Autonomous System Number	Sets the unique numeric identifier of the AS in which the peer host resides.	Type the peer host's 32-bit AS number, in dotted decimal notation. If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the AS is 0.0.0.3.
Peer Address	Specifies the IP address of the peer host's interface to which the BGP session is being established.	Type the IP address of the peer host's adjacent interface, in dotted decimal notation.
Local Address	Specifies the IP address of the local host's interface from which the BGP session is being established.	Type the IP address of the local host's adjacent interface, in dotted decimal notation.

Configuring BGP Networks with a Configuration Editor

To configure the Services Router as a node in a BGP network, you must perform the following tasks marked *(Required)*.

- (Required) "Configuring a Point-to-Point Peering Session" on page 335
- (Required) "Configuring BGP Within a Network" on page 338
- (Optional) "Configuring a Route Reflector" on page 339
- (Optional) "Configuring BGP Confederations" on page 342

For information about using the J-Web and CLI configuration editors, see "Using J-series Configuration Tools" on page 127.

Configuring a Point-to-Point Peering Session

To enable BGP traffic across one or more links, you must configure a BGP peering session with the adjacent host. Generally, such sessions are made at network exit points with neighboring hosts outside the autonomous system. Figure 88 shows a network with BGP peering sessions.

In the sample network, a Services Router in AS 17 has BGP peering sessions to a group of peers called external-peers. Peers A, B, and C reside in AS 22 and have IP addresses 10.10.10.10, 10.10.10.11, and 10.10.10.12. Peer D resides in AS 79, at IP address 10.21.7.2.

Figure 88: Typical Network with BGP Peering Sessions

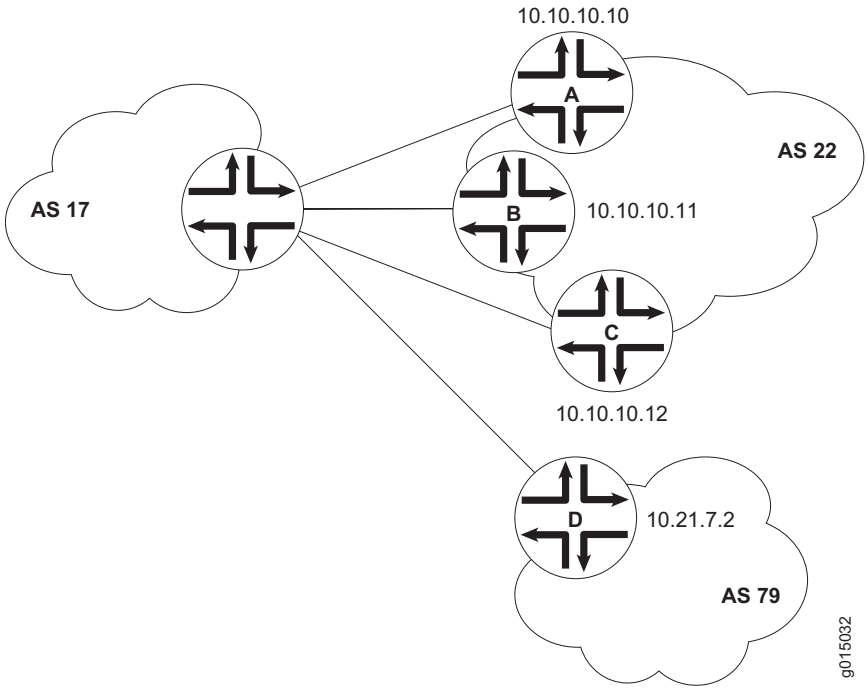

To configure the BGP peering sessions shown in Figure 88:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 125.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To configure IBGP sessions between peers, see “Configuring BGP Within a Network” on page 338.
   - To configure route reflector clusters, see “Configuring a Route Reflector” on page 339.
   - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations” on page 342.
   - To check the configuration, see “Verifying a BGP Configuration” on page 344.

Table 125: Configuring BGP Peering Sessions

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Routing-options** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing-options**.	From the top of the configuration hierarchy, enter edit routing-options
Set the network's AS number to 17.	1. In the AS Number box, enter 17. 2. Click **OK**.	Set the AS number to 17: set autonomous-system 17
Navigate to the **Bgp** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Bgp**.	From the top of the configuration hierarchy, enter edit protocols bgp
Create the BGP group external-peers, and add the external neighbor addresses to the group.	1. In the Group box, click **Add new entry**. 2. In the Group name box, type the name of the group of external BGP peers—external-peers in this case. 3. In the Neighbor box, click **Add new entry**. 4. In the Address box, type the IP address of an external BGP peer, in dotted decimal notation, and click **OK**. 5. Repeat Step 3 and Step 4 for each BGP neighbor within the external group that you are configuring.	1. Create the group external-peers, and add the address of an external neighbor: set group external-peers neighbor 10.10.10.10 2. Repeat Step 1 for each BGP neighbor within the external peer group that you are configuring.
At the group level, set the AS number for the group external-peers to 22. Because three of the peers in this group (peers A, B, and C) reside in one AS, you can set their AS number as a group.	1. In the Peer as box, type the number of the AS in which most peers in the external-peers group reside. 2. Click **OK**.	From the [edit protocols bgp] hierarchy level: set group external-peers peer-as 22
At the individual neighbor level, set the AS number for peer D to 79. Because peer D is a member of the group external-peers, it inherits the peer AS number configured at the group level. You must override this value at the individual neighbor level.	1. Under Neighbor, in the Address column, click the IP address of peer D—10.21.7.2 in this case. 2. In the Peer as box, type the AS number of the peer. 3. Click **OK**.	From the [edit protocols bgp group external-peers] hierarchy level: set neighbor 10.21.7.2 peer-as 79
Set the group type to external.	1. From the Type drop-down menu, select **external**. 2. Click **OK**.	From the [edit protocols bgp group external-peers] hierarchy level: set type external

Configuring BGP Within a Network

To configure BGP sessions between peering networks, you must configure point-to-point sessions between the external peers of the networks. Additionally, you must configure BGP internally to provide a means by which BGP route advertisements can be forwarded throughout the network. Because of the full mesh requirement of IBGP, you must configure individual peering sessions between all internal nodes of the network—unless you use route reflectors or confederations.

Figure 89 shows a typical network with external and internal peer sessions. In the sample network, the Services Router in AS 17 is fully meshed with its internal peers in the group internal-peers, which have IP addresses starting at 192.168.6.4.

Figure 89: Typical Network with EBGP External Sessions and IBGP Internal Sessions

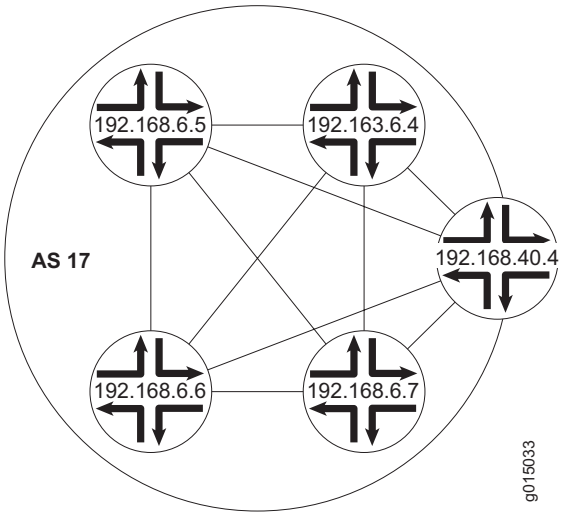

To configure IBGP in the network shown in Figure 89:

1. Configure all external peering sessions as described in “Configuring a Point-to-Point Peering Session” on page 335.
2. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
3. Perform the configuration tasks described in Table 126.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
   - To configure route reflector clusters, see “Configuring a Route Reflector” on page 339.
   - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations” on page 342.

- To check the configuration, see “Verifying a BGP Configuration” on page 344.

Table 126: Configuring IBGP Peering Sessions

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Bgp** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Bgp**.	From the top of the configuration hierarchy, enter edit protocols bgp
Create the BGP group internal-peers, and add the internal neighbor addresses to the group. You must configure a full IBGP mesh, which requires that each peer be configured with every other internal peer as a BGP neighbor.	1. In the Group box, click **Add new entry**. 2. In the Group name box, type the name of the group of internal BGP peers—internal-peers in this case. 3. In the Neighbor box, click **Add new entry**. 4. In the Address box, type the IP address of an internal BGP peer, in dotted decimal notation. 5. Click **OK**. 6. Repeat Step 3 and Step 4 for each internal BGP peer within the network.	1. Create the group internal-peers, and add the address of an internal neighbor: set group internal-peers neighbor 192.168.6.4 2. Repeat Step 1 for each internal BGP neighbor within the network.
Set the group type to internal.	1. From the Type drop-down menu, select **internal**. 2. Click **OK**.	From the [edit protocols bgp group internal-peers] hierarchy level: set type internal
Configure a routing policy to advertise BGP routes.	See “Injecting OSPF Routes into the BGP Routing Table” on page 380.	

Configuring a Route Reflector

Because of the IBGP full-mesh requirement, most networks use route reflectors to simplify configuration. Using a route reflector, you group routers into clusters, which are identified by numeric identifiers unique to the AS. Within the cluster, you must configure a BGP session from a single router (the route reflector) to each internal peer. With this configuration, the IBGP full-mesh requirement is met.

NOTE: You must have an Advanced BGP Feature license installed on each Services Router that uses a route reflector. For license details, see “Managing J-series Licenses” on page 69.

Figure 90 shows an IBGP network with a Services Router at IP address 192.168.40.4 acting as a route reflector. In the sample network, each router in cluster 2.3.4.5 has an internal client relationship to the route reflector. To configure the cluster:

- On the Services Router, create an internal group, configure an internal peer (neighbor) relationship to every other router in the cluster, and assign a cluster identifier.
- On each other router you are assigning to the cluster, create the cluster group and configure a client relationship to the route reflector.

Figure 90: Typical IBGP Network Using a Route Reflector

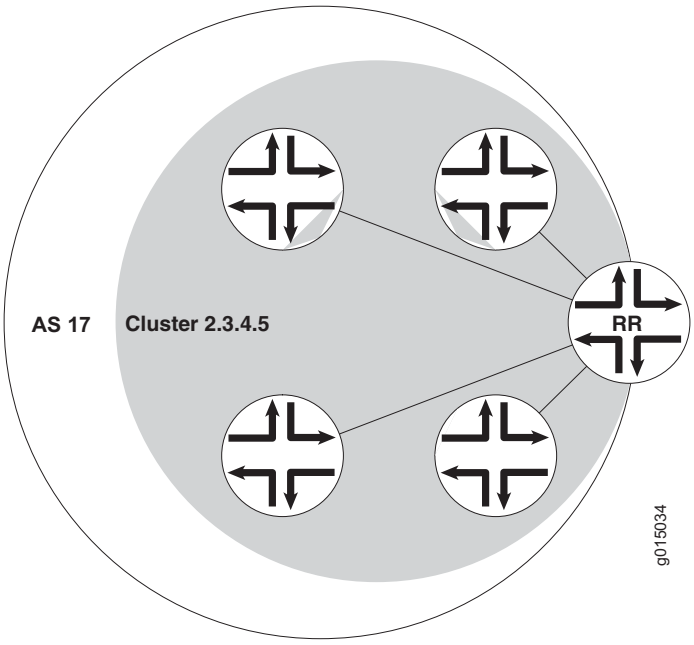

To configure IBGP in the network using the Services Router as a route reflector:

1. Configure all external peering sessions as described in “Configuring a Point-to-Point Peering Session” on page 335.
2. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
3. Perform the configuration tasks described in Table 127.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
   - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations” on page 342.

- To check the configuration, see “Verifying a BGP Configuration” on page 344.

Table 127: Configuring a Route Reflector

Task	J-Web Configuration Editor	CLI Configuration Editor
On the Services Router that you are using as a route reflector, navigate to the **Bgp** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Bgp**.	From the top of the configuration hierarchy, enter edit protocols bgp
On the Services Router that you are using as a route reflector, create the BGP group cluster-peers, and add to the group the IP addresses of the internal neighbors that you want in the cluster.	1. In the Group box, click **Add new entry**. 2. In the Group name box, type the name of the group in which the BGP peer is configured—cluster-peers in this case. 3. In the Neighbor box, click **Add new entry**. 4. In the Address box, type the IP address of a BGP peer, in dotted decimal notation. 5. Click **OK**. 6. Repeat Step 3 and Step 4 for each BGP neighbor within the cluster that you are configuring.	1. Create the group cluster-peers, and add the address of an internal neighbor: set group cluster-peers neighbor 192.168.6.4 2. Repeat Step 1 for each BGP neighbor within the cluster that you are configuring.
On the Services Router that you are using as a route reflector, set the group type to internal.	From the Type drop-down menu, select **internal**.	From the [edit protocols bgp group internal-peers] hierarchy level: set type internal
On the Services Router that you are using as a route reflector, configure the cluster identifier for the route reflector.	1. In the Cluster box, enter the unique numeric cluster identifier. 2. Click **OK**.	Set the cluster identifier: set cluster 2.3.4.5

Task	J-Web Configuration Editor	CLI Configuration Editor
On the other routers in the cluster, create the BGP group cluster-peers, and add the internal IP address of the route reflector. You do not need to include the neighbor addresses of the other internal peers, or configure the cluster identifier on these route reflector clients. They need only be configured as internal neighbors. **NOTE:** If the other routers in the network are Services Routers, follow the steps in this row. Otherwise, consult the router documentation for instructions.	On a client Services Router in the cluster: 1. In the configuration editor hierarchy, select **Protocols > Bgp**. 2. In the Group box, click **Add new entry**. 3. In the Group name box, type the name of the group in which the BGP peer is configured—cluster-peers in this case. 4. In the Neighbor box, click **Add new entry**. 5. In the Address box, type the IP address of the route reflector, in dotted decimal notation—in this case, 192.168.40.4. 6. Click **OK**.	On a client Services Router in the cluster: 1. From the top of the configuration hierarchy, enter edit protocols bgp 2. Create the group cluster-peers, and add only the route reflector address to the group: set group cluster-peers neighbor 192.168.40.4
Configure a routing policy to advertise BGP routes.	See “Injecting OSPF Routes into the BGP Routing Table” on page 380.	

Configuring BGP Confederations

To help solve BGP scaling problems caused by the IBGP full-mesh requirement, you can divide your AS into sub-ASs called confederations. As Figure 91 shows, the connections between the sub-ASs are made through EBGP sessions, and the internal connections are made through standard IBGP sessions.

In the sample network, AS 17 has two separate confederations (sub-AS 64512 and sub-AS 64513), each of which has multiple routers. Within a sub-AS, an IGP (OSPF, for example) is used to establish network connectivity with internal peers. Between sub-ASs, an external BGP peering session is established.

Figure 91: Typical Network Using BGP Confederations

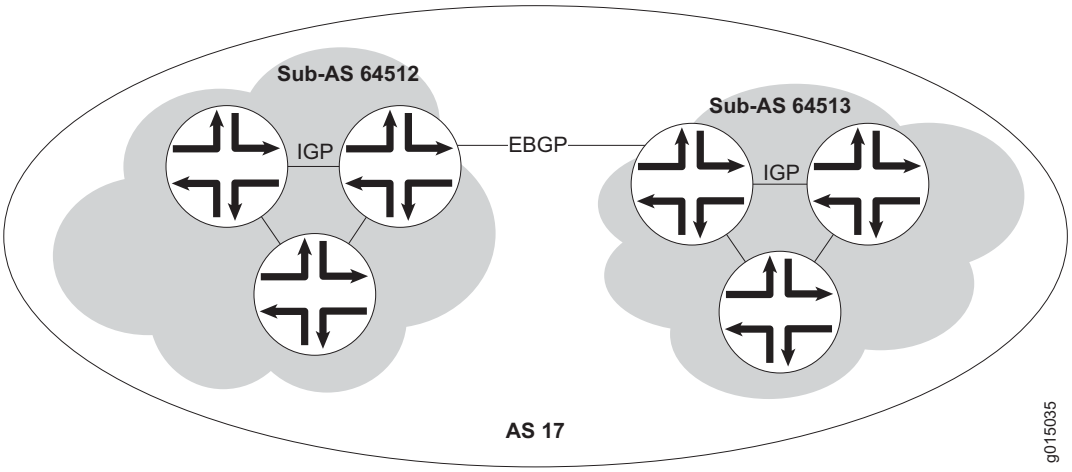

To configure the BGP confederations shown in Figure 91:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 128.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see "Verifying a BGP Configuration" on page 344.

Table 128: Configuring BGP Confederations

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Routing-options** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing-options**.	From the top of the configuration hierarchy, enter edit routing-options
Set the AS number to the sub-AS number 64512. The sub-AS number is a unique AS number that is usually taken from the pool of private AS numbers—64512 through 65535.	1. In the AS Number box, enter the sub-AS number. 2. Click **OK**.	Set the sub-AS number: set autonomous-system 64512
Navigate to the **Confederation** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing-options > Confederation**.	From the top of the configuration hierarchy, enter edit routing-options confederation
Set the confederation number to the AS number 17.	In the Confederation as box, enter 17.	Set the confederation AS number: set 17

Task	J-Web Configuration Editor	CLI Configuration Editor
Add the sub-ASs as members of the confederation. Every sub-AS within the AS must be added as a confederation member.	1. In the Members field, click **Add new entry**. 2. In the Value box, enter the sub-ASs that are members of this confederation. Separate multiple sub-ASs with a space.	Add members to the confederation: set 17 members 64512 64513
Using EBGP, configure the peering session between the confederations (from router A to router B in this example). When setting the peer AS number for these sessions, use the sub-AS number rather than the AS number.	See "Configuring a Point-to-Point Peering Session" on page 335.	
Using IBGP, configure internal sessions within a sub-AS. You can configure an IBGP full mesh, or you can configure a route reflector.	■ To configure an IBGP full mesh, see "Configuring BGP Within a Network" on page 338. ■ To configure a route reflector, see "Configuring a Route Reflector" on page 339.	

Verifying a BGP Configuration

To verify a BGP configuration, perform these tasks:

- "Verifying BGP Neighbors" on page 344
- "Verifying BGP Groups" on page 345
- "Verifying BGP Summary Information" on page 346
- "Verifying Reachability of All Peers in a BGP Network" on page 347

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From the CLI, enter the show bgp neighbor command.

Sample Output

```
user@host> show bgp neighbor

Peer: 10.255.245.12+179 AS 35  Local: 10.255.245.13+2884 AS 35
  Type: Internal    State: Established  (route reflector client)Flags: Sync
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh
  Address families configured: inet-vpn-unicast inet-labeled-unicast
```

```
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages:  Total 9      Updates 6       Refreshes 0     Octets 403
  Output messages: Total 7      Updates 3       Refreshes 0     Octets 365
  Output Queue[0]: 0
  Output Queue[1]: 0
  Trace options: detail packets
  Trace file: /var/log/bgpgr size 131072 files 10
```

What It Means The output shows a list of the BGP neighbors with detailed session information. Verify the following information:

- Each configured peering neighbor is listed.
- For State, each BGP session is Established.
- For Type, each peer is configured as the correct type (either internal or external).
- For AS, the AS number of the BGP neighbor is correct.

For more information about show bgp neighbor, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From the CLI, enter the show bgp group command.

Sample Output

```
user@host> show bgp group

Group Type: Internal    AS: 10045       Local AS: 10045
  Name: pe-to-asbr2                                 Flags: Export Eval
  Export: [ match-all ]
  Total peers: 1        Established: 1
  4.4.4.4+179
  bgp.l3vpn.0: 1/1/0
  vpn-green.inet.0: 1/1/0

Groups: 1   Peers: 1    External: 0    Internal: 1    Down peers: 0   Flaps: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0            1          1          0          0          0          0
```

What It Means The output shows a list of the BGP groups with detailed group information. Verify the following information:

- Each configured group is listed.
- For AS, each group's remote AS is configured correctly.
- For Local AS, each group's local AS is configured correctly.
- For Group Type, each group has the correct type (either internal or external).
- For Total peers, the expected number of peers within the group is shown.
- For Established, the expected number of peers within the group have BGP sessions in the Established state.
- The IP addresses of all the peers within the group are present.

For more information about show bgp group, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From the CLI, enter the show bgp summary command.

Sample Output

```
user@host> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/R
10.0.0.2        65002      88675      88652       0       2       42:38 2/4/0
10.0.0.3        65002      54528      54532       0       1     2w4d22h 0/0/0
10.0.0.4        65002      51597      51584       0       0     2w3d22h 2/2/0
```

What It Means The output shows a summary of BGP session information. Verify the following information:

- For Groups, the total number of configured groups is shown.
- For Peers, the total number of BGP peers is shown.
- For Down Peers, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under Peer, the IP address for each configured peer is shown.
- Under AS, the peer AS for each configured peer is correct.
- Under Up/Dwn State, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being damped (such as 0/0/0). If the field is Active, it indicates a problem in the establishment of the BGP session.

For more information about show bgp summary, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying Reachability of All Peers in a BGP Network

Purpose By using the ping tool on each peer address in the network, verify that all peers in the network are reachable from each Services Router.

Action For each Services Router in the BGP network:

1. In the J-Web interface, select **Diagnose > Ping Host**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router.
3. Click **Start**. Output appears on a separate page.

Sample Output

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

What It Means If a host is active, it generates an ICMP response. If this response is received, the round-trip time is listed in the time field. For more information about the ping output, see Table 82.

For more information about using the J-Web interface to ping a host, see "Using the J-Web Ping Host Tool" on page 218.

For information about the ping command, see "Using the ping Command" on page 226 or the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Part 6

Configuring Routing Policy, Firewall Filters, and Class of Service

Chapter 17

Policy, Firewall Filter, and Class-of-Service Overview

Several mechanisms can help you control the way routing information and data packets are handled by a router—routing policy, firewall filters, and class-of-service (CoS) rules. Routing policies control how information is imported to and exported from the routing tables, acting exclusively at the Routing Engine level. Firewall filters examine packets at the entry (ingress) and exit (egress) points of the Services Router, filtering traffic at the router level. CoS rules determine packet scheduling, buffering, and queueing within the router. These three mechanisms are at the core of managing how a router forwards traffic.

To manage the flow of information into and out of a Services Router, you must understand the fundamentals of routing policies, firewall filters, and CoS rules. To read this chapter, you need a basic understanding of IP routing protocols.

This chapter contains the following topics. For more information see the *JUNOS Policy Framework Configuration Guide* and the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Policy, Firewall Filter, and CoS Terms

Before configuring routing policies, firewall filters, or class of service (CoS) with Differentiated Services (DiffServ) on a Services Router, become familiar with the terms defined in Table 129.

Table 129: Policy, Firewall Filter, and CoS Terms

Term	Definition
assured forwarding (AF)	CoS packet forwarding class that provides a group of values you can define and includes four subclasses, AF1, AF2, AF3, and AF4, each with three drop probabilities, low, medium, and high.
behavior aggregate (BA) classifier	Feature that can be used to determine the forwarding treatment for each packet. The BA classifier maps a code point to a loss priority. The loss priority is used later in the work flow to select one of the two drop profiles used by random early detection (RED).
best-effort (BE)	CoS packet forwarding class that provides no service profile. For the BE forwarding class, loss priority is typically not carried in a code point, and random early detection (RED) drop profiles are more aggressive.
class of service (CoS)	Method of classifying traffic on a packet-by-packet basis, using information in the type-of-service (TOS) byte to assign traffic flows to different service levels.
Differentiated Services (DiffServ)	Services based on RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*. The DiffServ method of CoS uses the type-of-service (ToS) byte to identify different packet flows on a packet-by-packet basis. DiffServ adds a Class Selector code point (CSCP) and a DiffServ code point (DSCP).
DiffServ code point (DSCP)	Values for a 6-bit field defined in IP packet headers that can be used to enforce class-of-service (CoS) distinctions in a Services Router.
drop profile	Drop probabilities for different levels of buffer fullness that are used by random early detection (RED) to determine from which Services Router scheduling queue to drop packets.
expedited forwarding (EF)	CoS packet forwarding class that provides end-to-end service with low loss, low latency, low jitter, and assured bandwidth.
multifield (MF) classifier	Firewall filter that scans through a variety of packet fields to determine the forwarding class and loss priority for a packet and polices traffic to a specific bandwidth and burst size. Typically, a classifier performs matching operations on the selected fields against a configured value.
network address port translation (NAPT)	Method of concealing a set of host ports on a private network behind a pool of public addresses. It can be used as a security measure to protect the host ports from direct targeting in network attacks.
network address translation (NAT)	Method of concealing a set of host addresses on a private network behind a pool of public addresses. It can be used as a security measure to protect the host addresses from direct targeting in network attacks.
network control (NC)	CoS packet forwarding class that is typically high priority because it supports protocol control.
PLP bit	Packet loss priority bit. Used to identify packets that have experienced congestion or are from a transmission that exceeded a service provider's customer service license agreement. A Services Router can use the PLP bit as part of a congestion control strategy. The bit can be configured on an interface or in a filter.
policer	Feature that limits the amount of traffic passing into or out of an interface. It is an essential component of firewall filters that is designed to thwart denial-of-service (DoS) attacks. A policer applies rate limits on bandwidth and burst size for traffic on a particular Service Router interface.
policing	Applying rate and burst size limits to traffic on an interface.
random early detection (RED)	Gradual drop profile for a given class, used for congestion avoidance. RED attempts to anticipate congestion and reacts by dropping a small percentage of packets from the head of a queue to prevent congestion.

Term	Definition
rule	Guide that the Services Router follows when applying services. A rule consists of a match direction and one or more terms.
service set	Collection of services. Examples of services include stateful firewall filters and network address translation (NAT).
stateful firewall filter	Type of firewall filter that evaluates the context of connections, permits or denies traffic based on the context, and updates this information dynamically. Context includes IP source and destination addresses, TCP port numbers, TCP sequencing information, and TCP connection flags.
stateless firewall filter	Type of firewall filter that statically evaluates the contents of packets transiting the router, and packets originating from, or destined for, the router. Information about connection states is not maintained.
term	Firewall filters contain one or more terms that specify filter match conditions and actions.
trusted network	Network from which all originating traffic can be trusted—for example, an internal enterprise LAN. Stateful firewall filters allow traffic to flow from trusted to untrusted networks.
untrusted network	Network from which all originating traffic cannot be trusted—for example, a WAN. Unless configured otherwise, stateful firewall filters do not allow traffic to flow from untrusted to trusted networks.

Routing Policy Overview

Routing protocols send information about routes to a router's neighbors. This information is processed and used to create routing tables, which are then distilled into forwarding tables. Routing policies control the flow of information between the routing protocols and the routing tables and between the routing tables and the forwarding tables. Using policies, you can determine which routes are advertised, specify which routes are imported into the routing table, and modify routes to control which routes are added to the forwarding table.

This section contains the following topics:

Routing Policy Components

Routing policies are made up of one or more terms, which contain a set of match conditions and a set of actions. Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route. These actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

This section contains the following topics:

Routing Policy Terms

A term is a named structure in which match conditions and actions are defined. Each routing policy contains one or more terms,

Generally, a Services Router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of accept or reject is taken. If none of the terms in the policy match the route, the Services Router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

Routing Policy Match Conditions

A match condition defines the criteria that a route must match for an action to take place. Each term can have one or more match conditions. If a route matches all the match conditions for a particular term, the actions defined for that term are processed.

Each term can consist of two statements, to and from, that define match conditions:

- In the from statement, you define the criteria that an *incoming* route must match. You can specify one or more match conditions. If you specify more than one, all conditions must match the route for a match to occur.
- In the to statement, you define the criteria that an *outgoing* route must match. You can specify one or more match conditions. If you specify more than one, all conditions must match the route for a match to occur.

The order of match conditions in a term is not important, because a route must match all match conditions in a term for an action to be taken.

Table 130 summarizes the routing policy match conditions.

Table 130: Summary of Routing Policy Match Conditions

Match Condition	Description
aggregate-contributor	Matches routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.

Match Condition	Description
area *area-id*	Matches a route learned from the specified OSPF area during the exporting of OSPF routes into other protocols.
as-path *name*	Name of an AS path regular expression. BGP routes whose AS path matches the regular expression are processed.
color *preference*	Color value. You can specify preference values that are finer-grained than those specified in the *preference* match conditions. The color value can be a number from 0 through 4,294,967,295 9 ($2^{32} - 1$). A lower number indicates a more preferred route.
community	Name of one or more communities. If you list more than one name, only one name needs to match for a match to occur. (The matching is effectively a logical OR operation.)
external [type *metric-type*]	Matches external OSPF routes, including routes exported from one level to another. In this construct type is an optional keyword. The metric-type value can be either 1 or 2. When you do not specify type, this condition matches all external routes.
interface *interface-name*	Name or IP address of one or more router interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP). Depending on where the policy is applied, this match condition matches routes learned from or advertised through the specified interface.
internal	Matches a routing policy against the internal flag for simplified next-hop self policies.
level *level*	Matches the IS-IS level. Routes that are from the specified level or are being advertised to the specified level are processed.
local-preference *value*	BGP local preference attribute. The preference value can be from 0 through 4,294,967,295 9 ($2^{32} - 1$).
metric *metric* metric2 *metric*	Metric value. The metric value corresponds to the multiple exit discriminator (MED), and metric2 corresponds to the interior gateway protocol (IGP) metric if the BGP next hop runs back through another route.
neighbor *address*	Address of one or more neighbors (peers). For BGP export policies, the address can be a directly connected or indirectly connected peer. For all other protocols, the address is the neighbor from which the advertisement is received.
next-hop *address*	Next-hop address or addresses specified in the routing information for a particular route. For BGP routes, matches are performed against each protocol next hop.
origin *value*	BGP origin attribute, which is the origin of the AS path information. The value can be one of the following: ■ egp—Path information originated from another AS. ■ igp—Path information originated from within the local AS. ■ incomplete—Path information was learned by some other means.
policy [*policy-names*]	Name of one or more policies to evaluate as a subroutine.
preference *preference* preference2 *preference*	Preference value. You can specify a primary preference value (preference) and a secondary preference value (preference2). The preference value can be a number from 0 through 4,294,967,295 9 ($2^{32} - 1$). A lower number indicates a more preferred route.

Match Condition	Description
prefix-list *name*	Named list of IP addresses configured at the **Policy-options** level in the configuration hierarchy. This match condition can be used on import policies only.
protocol *protocol*	Name of the protocol from which the route was learned or to which the route is being advertised. It can be one of the following: aggregate, bgp, direct, dvmrp, isis, local, ospf, pim-dense, pim-sparse, rip, ripng, or static.
route-filter *destination-prefix match-type <actions>*	List of destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route or a less precise match using match types. You can configure either a common action that applies to the entire list or an action associated with each prefix. Route filters can be used on import policies only.
route-type *value*	Type of route. The value can be either external or internal.
source-address-filter *destination-prefix match-type <actions>*	List of multicast source addresses. When specifying a source address, you can specify an exact match with a specific route or a less precise match using match types. You can configure either a common action that applies to the entire list or an action associated with each prefix. Source-address filters can be used on import policies only.

Routing Policy Actions

An action defines what the Services Router does with the route when the route matches all the match conditions in the from and to statements for a particular term. If a term does not have from and to statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the then statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy
- Actions that manipulate route characteristics
- Trace action, which logs route matches

Table 131 summarizes the routing policy actions.

If you do not specify an action, one of the following results occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the accept or reject action specified by the default policy is executed.

Table 131: Summary of Key Routing Policy Actions

Action	Description
Flow Control Actions	These actions control the flow of routing information into and out of the routing table.
accept	Accepts the route and propagates it. After a route is accepted, no other terms in the routing policy and no other routing policies are evaluated.
reject	Rejects the route and does not propagate it. After a route is rejected, no other terms in the routing policy and no other routing policies are evaluated.
next term	Skips to and evaluates the next term in the same routing policy. Any accept or reject action specified in the then statement is ignored. Any actions specified in the then statement that manipulate route characteristics are applied to the route.
next policy	Skips to and evaluates the next routing policy. Any accept or reject action specified in the then statement is ignored. Any actions specified in the then statement that manipulate route characteristics are applied to the route.
Route Manipulation Actions	These actions manipulate the route characteristics.
as-path-prepend *as-path*	Appends one or more autonomous system (AS) numbers at the beginning of the AS path. If you are specifying more than one AS number, include the numbers in quotation marks. The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.
as-path-expand last-as count *n*	Extracts the last AS number in the existing AS path and appends that AS number to the beginning of the AS path *n* times. Replace *n* with a number from 1 through 32. The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.
class *class-name*	Applies the specified class-of-service (CoS) parameters to routes installed into the routing table.
color *preference* color2 *preference*	Sets the preference value to the specified value. The color and color2 preference values can be a number from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route.
damping *name*	Applies the specified route-damping parameters to the route. These parameters override BGP's default damping parameters. This action is useful only in import policies.
local-preference *value*	Sets the BGP local preference attribute. The preference can be a number from 0 through 4,294,967,295 ($2^{32} - 1$).

Action	Description
metric *metric* metric2 *metric* metric3 *metric* metric4 *metric*	Sets the metric. You can specify up to four metric values, starting with metric (for the first metric value) and continuing with metric2, metric3, and metric4. For BGP routes, metric corresponds to the MED, and metric2 corresponds to the IGP metric if the BGP next hop loops through another router.
next-hop *address*	Sets the next hop. If you specify *address* as self, the next-hop address is replaced by one of the local router's addresses. The advertising protocol determines which address to use.

Default and Final Actions

If none of the terms' match conditions evaluate to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either accept or reject) that overrides any action intrinsic to the protocol.

Applying Routing Policies

Once a policy is created, it must be applied before it is active. You apply routing policies using the import and export statements at the **Protocols >** *protocol-name* level in the configuration hierarchy.

In the import statement, list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the export statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

To specify more than one policy and create a policy chain, list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

Firewall Filter Overview

In a *stateful* firewall filter, all packets flowing from a trusted network to an untrusted network are allowed. Packets flowing from an untrusted network to a trusted network are allowed only if they are responses to a session originated by the trusted network, or if they are explicitly accepted by a term in the stateful firewall filter rule.

When Network Address Translation (NAT) is enabled, the source address of a packet flowing from a trusted network to an untrusted network is replaced with an address chosen from a specified range, or *pool*, of addresses. In addition, you can configure the Services Router to dynamically translate the source port of the packet—a process called *Network Address Port Translation (NAPT)*.

This section contains the following topics:

Stateful and Stateless Firewall Filters

A *stateless* firewall filter can filter packets transiting the Services Router from a source to a destination, or packets originating from, or destined for, the Routing Engine. Stateless firewall filters applied to the Routing Engine interface protect the processes and resources owned by the Routing Engine.

You can apply a stateless firewall filter to an input or output interface, or to both. Every packet, including fragmented packets, is evaluated against stateless firewall filters.

All firewall filters contain one or more terms, and each term consists of two components—match conditions and actions. The match conditions define the values or fields that the packet must contain to be considered a match. If a packet is a match, the corresponding action is taken. By default, a packet that does not match a firewall filter is discarded.

NOTE: A firewall filter with a large number of terms can adversely affect both the configuration commit time and the performance of the Routing Engine.

For more information about firewall filters, see "Configuring IPSec for Secure Packet Exchange" on page 483 and the *JUNOS Policy Framework Configuration Guide*. For more information about NAT, see the *JUNOS Services Interfaces Configuration Guide*.

Process for Configuring a Stateful Firewall Filter and NAT

To configure a stateful firewall filter and NAT, perform the following tasks:

- Define the stateful firewall filter output and input rules. You must define an output rule that allows all traffic (application and nonapplication) to flow from the trusted network to the untrusted network.

NOTE: If a packet does not match any terms in a stateful firewall filter rule, the packet is discarded.

To define the match condition in the term that allows application traffic to flow from the trusted network to the untrusted network, we recommend you specify the JUNOS default group junos-algs-outbound as the application set. To view the configuration of this group, enter the show groups junos-defaults applications application-set junos-algs-outbound configuration mode command. For more information about JUNOS default groups, see the *JUNOS System Basics Configuration Guide*.

You also must define an input rule to discard all traffic from the untrusted network that is not a response to a session originated by the trusted network.

- Define the NAT address and port pool.
- Define the NAT output and input rules.
- Define a service set that includes all stateful firewall filter and NAT rules and the service interface. You must specify the service interface as sp-0/0/0. This service interface is a virtual interface that must be included at the [edit interfaces] hierarchy level to support stateful firewall filter and NAT services.
- Apply the service set to the interfaces that make up the untrusted network.

NOTE: Do not apply the service set to the sp-0/0/0 interface.

For more information about match conditions and actions, see “Summary of Stateful Firewall Filter and NAT Match Conditions and Actions” on page 360.

Summary of Stateful Firewall Filter and NAT Match Conditions and Actions

Table 132 lists the match conditions you can specify in stateful firewall filter and NAT terms. Table 133 and Table 134 list actions you can specify in stateful firewall filter and NAT terms.

Table 132: Stateful Firewall Filter and NAT Match Conditions

Match Condition	Description
application-sets [*set-names*]	List of application set names. Application sets are defined at the [edit applications] hierarchy level.
applications [*application-names*]	List of applications. Applications are defined at the [edit applications] hierarchy level.
destination-address *address*	IP destination address field.
source-address *address*	IP source address field.

For more information about configuring applications and application sets for stateful firewall filters, see the *JUNOS Services Interfaces Configuration Guide*.

Table 133: Stateful Firewall Filter Actions

Actions	Description
accept	Accept the packet and send it to its destination.
allow-ip-options [*values*]	If the IP Option header of the packet contains a value that matches one of the specified values, accept the packet. If this action is not included, only packets without IP options are accepted. This action can be specified only with the accept action. You can specify the IP option as text or a numeric value: any (0), ip-security (130), ip-stream (8), loose-source-route (3), route-record (7), router-alert (148), strict-source-route (9), and timestamp (4).
discard	Do not accept the packet, and do not process it further.
reject	Do not accept the packet, and send a rejection message. UDP sends an ICMP unreachable code and RCP sends RST. Rejected packets can be logged or sampled.
syslog	Record information in the system logging facility. This action can be used with all options except discard.

Table 134: NAT Actions

Actions	Description
syslog	Record information in the system logging facility.
translated destination-pool *nat-pool-name*	Translate the destination address using the specified pool.
translated source-pool *nat-pool-name*	Translate the source address using the specified pool.
translation-type (destination *type* | source *type*)	Translate the destination and source port using the specified type: ■ destination static—Translate the destination address without port mapping. This type requires the size of the source address space to be the same as the size of the destination address space. You must specify a destination-pool name. The referenced pool must contain exactly one address and no port configuration at the [edit nat pool] hierarchy level. ■ source dynamic—Translate the source address with port mapping by means of NAPT. You must specify a source-pool name. The referenced pool must include a port configuration at the [edit nat pool] hierarchy level. ■ source static—Translate the source address without port mapping. This type requires the size of the source address space to be the same as the size of the destination address space. You must specify a source-pool name. The referenced pool must contain exactly one address and no port configuration at the [edit nat pool] hierarchy level.
syslog	Information is recorded in the system logging facility.

Planning a Stateless Firewall Filter

Before creating a stateless firewall filter and applying it to an interface, determine what you want the firewall filter to accomplish and how to use its match conditions and actions to achieve your goal. Also, make sure you understand how packets are matched and the default action of the resulting firewall filter.

CAUTION: If a packet does not match any terms in a stateless firewall filter rule, the packet is discarded. Take care that you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

To configure a stateless firewall filter, determine the following:

- Purpose of the firewall filter—for example, to limit traffic to certain protocols, IP source or destination addresses, or data rates, or to prevent denial-of-service (DoS) attacks.
- Appropriate match conditions. The packet header fields to match—for example, IP header fields (such as source and destination IP addresses, protocols, and IP options), TCP header fields (such as source and destination ports and flags), and ICMP header fields (such as ICMP packet type and code).
- Action to take if a match occurs—for example, accept, discard, or evaluate the next term.
- (Optional) Action modifiers. Additional actions to take if a packet matches—for example, count, log, rate limit, or police a packet.
- Interface on which the firewall filter is applied. The input or output side, or both sides, of the Routing Engine interface or a non-Routing Engine interface.

For more information about what a stateless firewall filter can include, see "Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers" on page 363. For more information about stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers

Table 135 lists the match conditions you can specify in stateless firewall filter terms. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a complete list of the synonyms, do any of the following:

- If you are using the J-Web interface, select the synonym from the appropriate drop-down list.
- If you are using the CLI, type a question mark (?) after the from statement.
- See the *JUNOS Policy Framework Configuration Guide*.

To specify a bit-field match condition with values, such as tcp-flags, you must enclose the values in quotation marks (" "). You can use bit-field logical operators to create expressions that are evaluated for matches. For example, if the following expression is used in a filter term, a match occurs if the packet is the initial packet of a TCP session:

```
tcp-flags "syn & !ack"
```

Table 136 lists the bit-field logical operators in order of highest to lowest precedence.

You can use text synonyms to specify some common bit-field matches. In the previous example, you can specify tcp-initial to specify the same match condition.

NOTE: When the Services Router compares the stateless firewall filter match conditions to a packet, it compares only the header fields specified in the match condition. There is no implied protocol match. For example, if you specify a match of destination-port ssh, the Services Router checks for a value of 0x22 in the 2-byte field that is two bytes after the IP packet header. The protocol field of the packet is not checked.

Table 135: Stateless Firewall Filter Match Conditions

Match Condition	Description
Numeric Range Match Conditions	
keyword-except	Negates a match. For example, destination-port-except *number*. The following keywords accept the -except extension: destination-port, dscp, esp-spi, forwarding-class, fragment-offset, icmp-code, icmp-type, interface-group, ip-options, packet-length, port, precedence, protocol and source-port.
destination-port *number*	TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the port and destination-port match conditions in the same term. Normally, you specify this match in conjunction with the protocol tcp or protocol udp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify telnet or 23.

Match Condition	Description
esp-spi *spi-value*	IPSec encapsulating security payload (ESP) security parameter index (SPI) value. Match on this specific SPI value. You can specify the ESP SPI value in either hexadecimal, binary, or decimal form.
forwarding-class *class*	Forwarding class. Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.
fragment-offset *number*	Fragment offset field.
icmp-code *number*	ICMP code field. Normally, you specify this match in conjunction with the protocol icmp match statement to determine which protocol is being used on the port. This value or keyword provides more specific information than icmp-type. Because the value's meaning depends on the associated icmp-type, you must specify icmp-type along with icmp-code. In place of the numeric value, you can specify a text synonym. For example, you can specify ip-header-bad or 0.
icmp-type *number*	ICMP packet type field. Normally, you specify this match in conjunction with the protocol icmp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify time-exceeded or 11.
interface-group *group-number*	Interface group on which the packet was received. An interface group is a set of one or more logical interfaces. For information about configuration interface groups, see the *JUNOS Policy Framework Configuration Guide*.
packet-length *bytes*	Length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
port *number*	TCP or UDP source or destination port field. You cannot specify both the port match and either the destination-port or source-port match conditions in the same term. Normally, you specify this match in conjunction with the protocol tcp or protocol udp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify bgp or 179.
precedence *ip-precedence-field*	IP precedence field. You can specify precedence in either hexadecimal, binary, or decimal form. In place of the numeric value, you can specify a text synonym. For example, you can specify immediate or 0x40.
protocol *number*	IP protocol field. In place of the numeric value, you can specify a text synonym. For example, you can specify ospf or 89.
source-port *number*	TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term. Normally, you specify this match in conjunction with the protocol tcp or protocol udp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify http or 80.
Address Match Conditions	
address *prefix*	IP source or destination address field. You cannot specify both the address and the destination-address or source-address match conditions in the same term.
destination-address *prefix*	IP destination address field. You cannot specify the destination-address and address match conditions in the same term.

Match Condition	Description
destination-prefix-list *prefix-list*	IP destination prefix list field. You cannot specify the destination-prefix-list and prefix-list match conditions in the same term.
prefix-list *prefix-list*	IP source or destination prefix list field. You cannot specify both the prefix-list and the destination-prefix-list or source-prefix-list match conditions in the same term.
source-address *prefix*	IP source address field. You cannot specify the source-address and address match conditions in the same rule.
source-prefix-list *prefix-list*	IP source prefix list field. You cannot specify the source-prefix-list and prefix-list match conditions in the same term.
Bit-Field Match Conditions with Values	
fragment-flags *number*	IP fragmentation flags. In place of the numeric value, you can specify a text synonym. For example, you can specify more-fragments or 0x2000.
ip-options *number*	IP options. In place of the numeric value, you can specify a text synonym. For example, you can specify record-route or 7.
tcp-flags *number*	TCP flags. Normally, you specify this match in conjunction with the protocol tcp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify syn or 0x02.
Bit-Field Text Synonym Match Conditions	
first-fragment	First fragment of a fragmented packet. This condition does not match unfragmented packets.
is-fragment	This condition matches if the packet is a trailing fragment. It does not match the first fragment of a fragmented packet. To match both first and trailing fragments, you can use two terms, or you can use fragment-offset 0-8191.
tcp-established	TCP packets other than the first packet of a connection. This match condition is a synonym for "(ack | rst)". This condition does not implicitly check that the protocol is TCP. To do so, specify the protocol tcp match condition.
tcp-initial	First TCP packet of a connection. This match condition is a synonym for "(syn & !ack)". This condition does not implicitly check that the protocol is TCP. To do so, specify the protocol tcp match condition.

Table 136: Stateless Firewall Filter Bit-Field Logical Operators

Logical Operator	Description
(...)	Grouping
!	Negation
& or +	Logical AND
| or ,	Logical OR

Table 137 lists the actions and action modifiers you can specify in stateless firewall filter terms.

Table 137: Stateless Firewall Filter Actions and Action Modifiers

Action or Action Modifier	Description
accept	Accepts a packet. This is the default if the packet matches. However, we strongly recommend that you always explicitly configure an action in the then statement.
discard	Discards a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Packets are available for logging and sampling before being discarded.
next term	Continues to the next term for evaluation.
reject <*message-type*>	Discards a packet, sending an ICMP destination unreachable message. Rejected packets are available for logging and sampling. You can specify one of the following message types: administratively-prohibited (default), bad-host-tos, bad-network-tos, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset. If you specify tcp-reset, a TCP reset is returned if the packet is a TCP packet. Otherwise, nothing is returned.
routing-instance *routing-instance*	Routes the packet using the specified routing instance.
Action Modifiers	
count *counter-name*	Counts the number of packets passing this term. The name can contain letters, numbers, and hyphens (-), and can be up to 24 characters long. A counter name is specific to the filter that uses it, so all interfaces that use the same filter increment the same counter.
forwarding-class *class-name*	Classifies the packet to the specified forwarding class.
log	Logs the packet's header information in the Routing Engine. You can access this information by entering the show firewall log command at the CLI.
loss-priority *priority*	Sets the scheduling priority of the packet. The priority can be low or high.
policer *policer-name*	Applies rate limits to the traffic using the named policer.
sample	Samples the traffic on the interface. Use this modifier only when traffic sampling is enabled. For more information, see the *JUNOS Policy Framework Configuration Guide*.
syslog	Records information in the system logging facility. This action can be used in conjunction with all options except discard.

Class-of-Service Overview

With the class-of-service (CoS) features on a Services Router, you can assign service levels with different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows. CoS is especially useful for networks supporting time-sensitive video and audio applications. To configure CoS features on a Services Router, see "Configuring Class of Service with DiffServ" on page 427.

This section contains the following topics. For more information about CoS and DiffServ, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Benefits of DiffServ CoS

IP routers normally forward packets independently, without controlling throughput or delay. This type of packet forwarding, known as best-effort service, is as good as your network equipment and links allow. Best-effort service is sufficient for many traditional IP data delivery applications, such as e-mail or Web browsing. However, newer IP applications such as real-time video and audio (or voice) require lower delay, jitter, and packet loss than simple best-effort networks can provide.

CoS features allow a Services Router to improve its processing of critical packets while maintaining best-effort traffic flows, even during periods of congestion. Network throughput is determined by a combination of available bandwidth and delay. CoS dedicates a guaranteed minimum bandwidth to a particular service class by reducing forwarding queue delays. (The other two elements of overall network delay, serial transmission delays determined by link speeds and propagation delays determined by media type, are not affected by CoS settings.)

Normally, packets are queued for output in their order of arrival, regardless of service class. Queueing delays increase with network congestion and often result in lost packets when queue buffers overflow. CoS packet classification assigns packets to forwarding queues by service class.

Because CoS must be implemented consistently end-to-end through the network, the CoS features on the Services Router are based on IETF Differentiated Services (DiffServ) standards, to interoperate with other vendors' CoS implementations.

DSCPs and Forwarding Service Classes

DiffServ specifications establish a 6-bit field in the IP packet header to indicate the forwarding service class to apply to the packet. The bit values in the DiffServ field form DiffServ code points (DSCPs) that can be set by the application or by a Services Router on the edge of a DiffServ-enabled network.

Each DiffServ forwarding service class has a well-known name and alias. Although not part of the specifications, the aliases are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

The 21 well-known DSCPs establish five DiffServ service classes. Table 138 identifies the forwarding service classes and aliases that correspond to the 21 DSCPs.

Table 138: Default Forwarding Service Class-to-DSCP Mapping

DiffServ Service Class Alias	IP DSCP	Forwarding Service Class and Use
ef	101110	**Expedited forwarding**—The Services Router delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class. Routers accept excess traffic in this class, but in contrast to assured forwarding, out-of-profile expedited-forwarding packets can be forwarded out of sequence or dropped.
af11	001010	**Assured forwarding**—The Services Router offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define. The router accepts excess traffic, but applies a random early discard (RED) drop profile to decide if the excess packets is dropped and not forwarded. Three drop probabilities (low, medium, and high) are defined for this service class.
af12	001100	
af13	001110	
af21	010010	
af22	010100	
af23	010110	
af31	011010	
af32	011100	
af33	011110	
af41	100010	
af42	100100	
af43	100110	
be	000000	**Best-effort**—The Services Router does not apply any special CoS handling to packets with 000000 in the DiffServ field, a backward compatibility feature. These packets are usually dropped under congested network conditions.
cs1	001000	**Conversational services**—The Services Router delivers assured (usually low) bandwidth with low delay and jitter for packets in this service class. Packets can be dropped, but are never delivered out of sequence. Packetized voice is a good example of a conversational service.
cs2	010000	
cs3	011000	
cs4	100000	
cs5	101000	
nc1/cs6	110000	**Network control**—The Services Router delivers packets in this service class with a low priority. (These packets are not delay sensitive.) Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, delay is preferable to discard. (See also the conversational services description in this table.)
nc2/cs7	111000	

JUNOS CoS Functions

Although the DiffServ CoS specifications define the position and length of the DSCP in the packet header, the DiffServ implementation is vendor specific. DiffServ CoS functions in JUNOS software are implemented by a series of components that you configure individually or in combination to define particular service offerings.

Figure 92 shows the components of the JUNOS CoS features, illustrating the sequence in which they interact. Table 139 defines the components and explains their use.

Figure 92: Packet Flow Through JUNOS CoS-Configurable Components

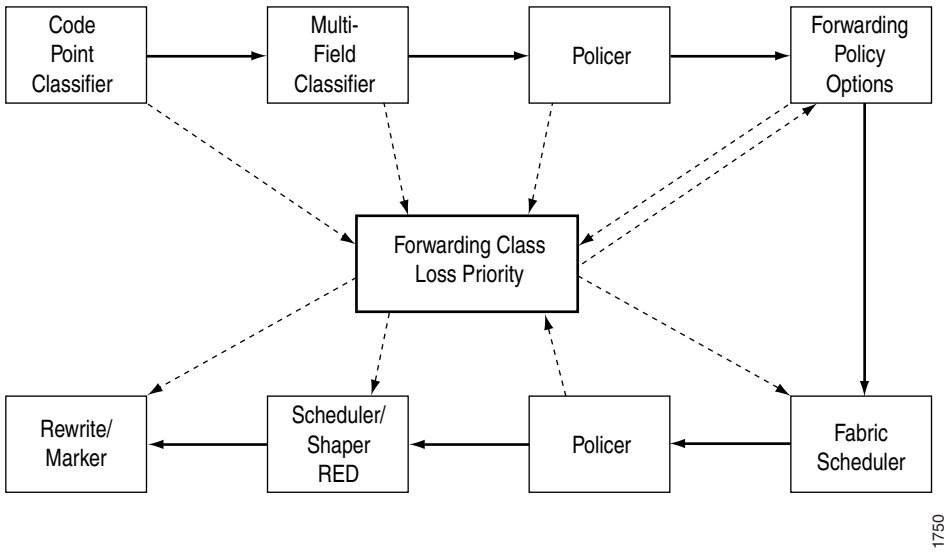

Table 139: JUNOS CoS Components

CoS Component	Use
Classifiers	Associate incoming packets with a forwarding class and packet loss priority (PLP). The following types of classifiers are available: ■ Behavior aggregate (BA) or code point traffic classifiers—Allow you to set the forwarding class and PLP based on DSCP. ■ Multifield (MF) traffic classifiers—Allow you to set the forwarding class and PLP based on firewall filter rules. This is usually done at the edge of the network for packets that do not have valid DSCPs in the packet headers.
Forwarding classes	Allow you to set the scheduling and marking of packets as they transit the Services Router. Known as ordered aggregates in the DiffServ architecture, the forwarding class plus the loss priority determine the router's per-hop behavior (PHB in DiffServ) for CoS.
Loss priorities	Allow you to set the priority of dropping a packet before it is sent. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering.

CoS Component	Use
Forwarding policy options	■ Allow you to associate forwarding classes with next hops. ■ Allow you to create classification overrides, which assign forwarding classes to sets of prefixes.
Transmission scheduling and rate control	Provide you with a variety of tools to manage traffic flows. The following types are available: ■ Schedulers—Allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular forwarding class for packet transmission. Drop profiles are useful for the assured forwarding service class. ■ Fabric schedulers—For M320 and T-series platforms only, fabric schedulers allow you to identify a packet as high or low priority based on its forwarding class, and to associate schedulers with the fabric priorities. ■ Policers for traffic classes—Allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class or to a different loss priority, or to both. You define policers with filters that can be associated with input or output interfaces. Policers are useful for the expedited forwarding service class.
Rewrite markers	Allow you to redefine the DSCP value of outgoing packets. Rewriting or marking outbound packets is useful when the routing platform is at the border of a network and must alter the code points to meet the policies of the targeted peer.

How Forwarding Classes and Schedulers Work

This section contains the following topics:

Default Forwarding Class Queue Assignments

J-series routers have only four queues built into the hardware. Other routing platforms can be configured for up to eight queues. If a classifier does not assign a packet to any other queue (for example, for other than well-known DSCPs that have not been added to the classifier), the packet is assigned by default to the class associated with queue 0.

Table 140 shows the four forwarding classes and queues that Juniper Networks classifiers assign to packets based on the DSCP values in arriving packet headers.

Table 140: Default Forwarding Class Queue Assignments

Forwarding Class	Forwarding Queue
best-effort	queue 0
expedited-forwarding	queue 1
assured-forwarding	queue 2
network-control	queue 3

Default Scheduler Settings

Each forwarding class has an associated scheduler priority. Only two forwarding classes, best-effort and network-control (queue 0 and queue 3), are used in the JUNOS default scheduler configuration.

By default, the best-effort forwarding class (queue 0) receives 95 percent of the output link bandwidth and buffer space, and the network-control forwarding class (queue 3) receives 5 percent of the output link bandwidth and buffer space. The default drop profile causes the buffer to fill and then discard all packets until it again has space.

The expedited-forwarding and assured-forwarding classes have no schedulers, because by default no resources are assigned to queue 1 and queue 2. However, you can manually configure resources for expedited-forwarding and assured-forwarding.

The default scheduler settings are implicit in the configuration, although they do not appear in the output of the show class-of-service command.

```
[edit class-of-service]
schedulers {
   network-control {
                transmit-rate percent 5;
                buffer-size percent 5;
                priority low;
                drop-profile-map loss-priority any protocol any;
                drop-profile terminal;
   }
   best-effort {
                transmit-rate percent 95;
                buffer-size percent 95;
                priority low;
                drop-profile-map loss-priority any protocol any;
                drop-profile terminal;
   }
}
drop-profiles {
   terminal {
      fill-level 100 drop-probability 100;
   }
}
```

Default Behavior Aggregate (BA) Classifiers

Table 141 shows the forwarding class and packet loss priority (PLP) that are assigned by default to each well-known DSCP. Although several DSCPs map to the expedited-forwarding (ef) and assured-forwarding (af) classes, by default no resources are assigned to these forwarding classes. All af classes other than af1*x* are mapped to best-effort, because RFC 2597, *Assured Forwarding PHB Group*, prohibits a node from aggregating classes. Assignment to best-effort implies that the node does not support that class.

You can modify the default settings through configuration. For instructions, see "Configuring Class of Service with DiffServ" on page 427.

Table 141: Default Behavior Aggregate (BA) Classification

DSCP Alias	Forwarding Class	Packet Loss Priority (PLP)
ef	expedited-forwarding	low
af11	assured-forwarding	low
af12	assured-forwarding	high
af13	assured-forwarding	high
af21	best-effort	low
af22	best-effort	low
af23	best-effort	low
af31	best-effort	low
af32	best-effort	low
af33	best-effort	low
af41	best-effort	low
af42	best-effort	low
af43	best-effort	low
be	best-effort	low
cs1	best-effort	low
cs2	best-effort	low
cs3	best-effort	low
cs4	best-effort	low
cs5	best-effort	low
nc1/cs6	network-control	low
nc2/cs7	network-control	low
other	best-effort	low

DSCP Rewrites

Typically, a router rewrites the DSCPs in outgoing packets once, when packets enter the DiffServ portion of the network, either because the packets do not arrive from the customer with the proper DSCP bit set or because the service provider wants to verify the that customer has set the DSCP properly. CoS implementations that accept the DSCP and classify and schedule traffic solely on DSCP value perform behavior aggregate (BA) DiffServ functions and do not usually rewrite the DSCP. DSCP rewrites typically occur in multifield (MF) DiffServ scenarios.

For instructions for configuring rewrite rules, see "Configuring and Applying Rewrite Rules" on page 435.

Sample BA Classification

Table 142 shows the router forwarding classes associated with each well-known DSCP code point and the resources assigned to their output queues for a sample DiffServ CoS implementation. This example assigns expedited forwarding to queue 1 and a subset of the assured forwarding classes (af1*x*) to queue 2, and distributes resources among all four forwarding classes.

Other DiffServ-based implementations are possible. For configuration information, see "Configuring Class of Service with DiffServ" on page 427.

Table 142: Sample BA Classification Forwarding Classes and Queues

DSCP Alias	DSCP Bits	Forwarding Class	PLP	Queue
ef	101110	expedited-forwarding	low	1
af11	001010	assured-forwarding	low	2
af12	001100	assured-forwarding	high	2
af13	001110	assured-forwarding	high	2
af21	010010	best-effort	low	0
af22	010100	best-effort	low	0
af23	010110	best-effort	low	0
af31	011010	best-effort	low	0
af32	011100	best-effort	low	0
af33	011110	best-effort	low	0
af41	100010	best-effort	low	0
af42	100100	best-effort	low	0
af43	100110	best-effort	low	0
be	000000	best-effort	low	0
cs1	0010000	best-effort	low	0
cs2	010000	best-effort	low	0
cs3	011000	best-effort	low	0

DSCP Alias	DSCP Bits	Forwarding Class	PLP	Queue
cs4	100000	best-effort	low	0
cs5	101000	best-effort	low	0
nc1/cs6	110000	network-control	low	3
nc2/cs7	111000	network-control	low	3
other	—	best-effort	low	0

Chapter 18
Configuring Routing Policies

Use routing policies as filters to control the information from routing protocols that a Services Router imports into its routing table and the information that the router exports (advertises) to its neighbors. To create a routing policy, you configure criteria against which routes are compared, and the action that is performed if the criteria are met.

You use either the J-Web configuration editor or CLI configuration editor to configure a routing policy.

This chapter contains the following topics. For more information about routing policies, see the *JUNOS Policy Framework Configuration Guide*.

Before You Begin

Before you begin configuring a routing policy, complete the following tasks:

- If you do not already have a basic understanding of routing policies, read "Routing Policy Overview" on page 353.
- Determine what you want to accomplish with the policy, and thoroughly understand how to achieve your goal using the various match conditions and actions.
- Make certain that you understand the default policies and actions for the policy you are configuring.
- Configure an interface on the router. See "Configuring Network Interfaces" on page 79.
- Configure an Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), if necessary. See "Configuring BGP Sessions" on page 331.
- Configure the router interface to reject or accept routes, if necessary. See "Configuring Firewall Filters and NAT" on page 389.
- Configure static routes, if necessary. See "Configuring Static Routes" on page 285.

Configuring a Routing Policy with a Configuration Editor

A routing policy has a major impact on the flow of routing information or packets within and through the Services Router. The match conditions and actions allow you to configure a customized policy to fit your needs.

To configure a routing policy, you must perform the following tasks marked *(Required)*. Perform additional tasks as needed for your router.

- (Required) "Configuring the Policy Name" on page 377
- (Required) "Configuring a Policy Term" on page 377
- (Optional) "Rejecting Known Invalid Routes" on page 378
- (Optional) "Injecting OSPF Routes into the BGP Routing Table" on page 380
- (Optional) "Grouping Source and Destination Prefixes in a Forwarding Class" on page 382
- (Optional) "Configuring Policy to Prepend the AS Path" on page 383
- (Optional) "Configuring Damping Parameters" on page 385

For information about using the J-Web and CLI configuration editors, see "Using J-series Configuration Tools" on page 127.

Configuring the Policy Name

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each routing policy name must be unique within a configuration.

To configure the policy name:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 143.
3. Go on to "Configuring a Policy Term" on page 377.

Table 143: Configuring the Policy Name

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Policy statement** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement**.	From the top of the CLI configuration hierarchy, enter edit policy-options
Enter the policy name.	In the Policy name box, type the name of the policy.	Type the policy-name value. For example: set policy-statement policy1
Apply your configuration changes.	Click **OK** to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the set command.

Configuring a Policy Term

Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

To configure a policy term:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 144.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to one of the following procedures:

- To remove useless routes, see “Rejecting Known Invalid Routes” on page 378.
- To advertise additional routes, see “Injecting OSPF Routes into the BGP Routing Table” on page 380.
- To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class” on page 382.
- To make a route less preferable to BGP, see “Configuring Policy to Prepend the AS Path” on page 383.
- To suppress route information, see “Configuring Damping Parameters” on page 385.

Table 144: Configuring a Policy Term

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Policy statement** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement**.	From the top of the CLI configuration hierarchy, enter edit policy-options policy-statement policy1
Create and name a policy term.	1. In the Term box, click **Add new entry**. 2. In the Term name box, type the name of a term and click **OK**.	Create and name a policy term. For example: set term term1

Rejecting Known Invalid Routes

You can specify known invalid (“bad”) routes to ignore by specifying matches on destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route, or a less precise match by using match types. You can configure either a common reject action that applies to the entire list, or an action associated with each prefix. Table 145 lists route list match types.

Table 145: Route List Match Types

Match Type	Match If ...
exact	The route shares the same most-significant bits (described by *prefix-length*), and *prefix-length* is equal to the route’s prefix length.
longer	The route shares the same most-significant bits (described by *prefix-length*), and *prefix-length* is greater than the route’s prefix length.

Match Type	Match If ...
orlonger	The route shares the same most-significant bits (described by *prefix-length*), and *prefix-length* is equal to or greater than the route's prefix length.
prefix-length-range *prefix-length2*-*prefix-length3*	The route shares the same most-significant bits (described by *prefix-length*), and the route's prefix length falls between *prefix-length2* and *prefix-length3*, inclusive.
through *destination-prefix*	All the following are true: ■ The route shares the same most-significant bits (described by *prefix-length*) of the first destination prefix. ■ The route shares the same most-significant bits (described by *prefix-length*) of the second destination prefix for the number of bits in the prefix length. ■ The number of bits in the route's prefix length is less than or equal to the number of bits in the second prefix. You do not use the through match type in most routing policy configurations. For more information, see the *JUNOS Policy Framework Configuration Guide*.
upto *prefix-length2*	The route shares the same most-significant bits (described by *prefix-length*) and the route's prefix length falls between *prefix-length* and *prefix-length2*.

For example, to reject routes with a mask of /8 and greater (/8, /9, /10, and so on) that have the first 8 bits set to 0 and accept routes less than 8 bits in length:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 146.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to one of the following procedures:
   - To advertise additional routes, see "Injecting OSPF Routes into the BGP Routing Table" on page 380.
   - To create a forwarding class, see "Grouping Source and Destination Prefixes in a Forwarding Class" on page 382.
   - To make a route less preferable to BGP, see "Configuring Policy to Prepend the AS Path" on page 383.
   - To suppress route information, see "Configuring Damping Parameters" on page 385.

Table 146: Creating a Policy to Reject Known Invalid Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Term** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement > Term**.	From the top of the CLI configuration hierarchy, enter edit policy-options policy-statement rejectpolicy1 term rejectterm1
Specify the routes to accept.	1. In the From option, click **Configure**. 2. In the Route filter box, click **Add new entry**. 3. In the Address box, enter the prefix of the routes. 4. Click **OK**.	Accept routes less than 8 bits in length: set from route-filter 0/0 up to /7 accept
Accept these routes.	1. In the Then option, click **Configure**. 2. In the Accept option, select the **Yes** check box. 3. Click **OK**.	
Specify the routes to reject.	1. In the configuration editor hierarchy, select **Policy options > Policy statement > Term**. 2. In the From option, click **Configure**. 3. In the Route filter box, click **Add new entry**. 4. In the Value box, enter the prefix of the routes to reject. 5. Click **OK**.	1. Specify routes less than 8 bits in length: set from route-filter /8 orlonger 2. Reject these routes: set then reject
Reject these routes.	1. In the Then option, click **Configure**. 2. In the Reject option, select the **Yes** check box. 3. Click **OK**.	

Injecting OSPF Routes into the BGP Routing Table

You can specify a match condition for policies based on procotols by naming a protocol from which the route is learned or to which the route is being advertised. You can specify one of the following protocols: aggregate, BGP, direct, DVMRP, IS-IS, local, OSPF, PIM-dense, PIM-sparse, RIP, or static

For example, you can inject or redistribute OSPF routes into the BGP routing table by creating a routing policy.

To redistribute OSPF routes from area 1 only into BGP and not advertise routes learned by BGP:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 147.
3. If you are finished configuring the policy injectpolicy1, commit the configuration.
4. Go on to one of the following procedures:
   - To create a forwarding class, see "Grouping Source and Destination Prefixes in a Forwarding Class" on page 382.
   - To make a route less preferable to BGP, see "Configuring Policy to Prepend the AS Path" on page 383.
   - To suppress route information, see "Configuring Damping Parameters" on page 385.

Table 147: Creating a Policy to Inject OSPF Routes into BGP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Term** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement > Term**.	From the top of the CLI configuration hierarchy, enter edit policy-options policy-statement injectpolicy1 term injectterm1
Specify the OSPF routes.	1. In the From option, click **Configure**. 2. In the Protocol box, click **Add new entry**. 3. In the Value drop box, select **OSPF**. 4. Click **OK**.	Specify the OSPF match condition: set from ospf
Specify the routes from a particular OSPF area.	1. In the Area option, type 1. 2. Click **OK**.	Specify Area 1 as a match condition: set from area 1
Specify that the route is to be accepted if the previous conditions are matched.	1. Next to Then, click **Configure**. 2. From the Accept reject box, Select **Accept**.	Specify the action to accept: set then accept

Task	J-Web Configuration Editor	CLI Configuration Editor
Set the default option to reject other OSPF routes.	1. In the configuration editor hierarchy, select **Policy options > Policy statement > Term**. 2. In the Then option, click **Configure**. 3. From the Accept reject box, Select **Reject**. 4. Click **OK**.	Changes in the CLI are applied automatically when you execute the set command.
Navigate to the **Protocol > Bgp**level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Protocols > Bgp**.	From the top of the CLI configuration hierarchy, enter: edit protocols bgp
Apply the routing policy **policy1** to BGP.	1. In the Export box, click **Add new entry**. 2. In the Value option, enter policy1. 3. Click **OK**.	Specify the OSPF match condition: set export policy1

Grouping Source and Destination Prefixes in a Forwarding Class

Create a forwarding class that includes packets based on both the destination address and the source address in the packet.

To configure and apply a routing policy to group source and destination prefixes in a forwarding class:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 148.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to one of the following procedures:
   - To make a route less preferable to BGP, see “Configuring Policy to Prepend the AS Path” on page 383.
   - To suppress route information, see “Configuring Damping Parameters” on page 385.

Table 148: Creating a Policy to Group Source and Destination Prefixes in a Forwarding Class

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Term** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement > Term**.	From the top of the CLI configuration hierarchy, enter edit policy-options policy-statement policy1 term term1
Specify the routes to include in the route filter.	1. In the From option, click **Configure**. 2. In the Route filter box, click **Add new entry**. 3. In the Value box, enter the source and destination prefixes. 4. Click **OK**.	1. Specify source routes 10.210.0.0/16 or longer: set from route-filter 10.210.0.0/16 orlonger 2. Specify destination routes 10.215.0.0/16 or longer: set from route-filter 10.215.0.0/16 orlonger
Group the source and destination prefixes.	1. In the configuration editor hierarchy, select **Policy options > Policy statement > Term**. 2. In the Then option, click **Configure**. 3. In the Forwarding class box, enter the forwarding class name. 4. Click **OK**.	Specify the forwarding class name: set then forwarding class forwarding-class-name1
Navigate to the **Forwarding table** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Routing options > Forwarding table**.	From the top of the CLI configuration hierarchy, enter edit routing-options forwarding-table
Apply the policy to the forwarding table.	1. In the Export box, click **Add new entry**. 2. In the Value box, enter the name of the policy. 3. Click **OK**. The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.	Specify source routes 10.210.0.0/16 or longer: set export policy1 You can refer to the same routing policy one or more times in the same or a different export statement.

Configuring Policy to Prepend the AS Path

You can *prepend* or add one or more autonomous system (AS) numbers at the beginning of an AS path. The AS numbers are added after the local AS number has been added to the path. Prepending an AS path makes a shorter AS path look longer and therefore less preferable to the Border Gateway Protocol (BGP).

For example, from AS 1, there are two equal paths (through AS 2 and AS 3) to reach AS 4. You might want packets from certain sources to use the path through AS 2. Therefore, you must make the path through AS 3 look less preferable so that BGP chooses the path through AS 2. In AS 1, you can prepend multiple AS numbers.

To prepend multiple AS numbers:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 149.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to "Configuring Damping Parameters" on page 385.

Table 149: Creating a Policy to Prepend AS Numbers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Term** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement > Term**.	From the top of the CLI configuration hierarchy, enter edit policy-options policy-statement prependpolicy1 term prependterm1
Specify the routes to prepend AS numbers to.	1. In the From option, click **Configure**. 2. In the Value box, enter the prefixes you wish to prepend. 3. In the Route filter box, click **Add new entry**. 4. Click **OK**.	1. Prepend routes 172.168.0.0/12 or longer: set from route-filter 172.16.0.0/12 orlonger 2. Prepend routes 192.168.0.0/16 or longer: set from route-filter 192.168.0.0/16 orlonger 3. Prepend routes 10.0.0.0/8 or longer: set from route-filter 10.0.0.0/8 orlonger
Specify the AS numbers to prepend.	1. In the configuration editor hierarchy, select **Policy options > Policy statement > Term**. 2. In the Then option, click **Configure**. 3. In the AS path prepend box, enter the string of AS numbers to prepend. Separate each AS number with a space. 4. Click **OK**.	Specify the AS numbers to prepend, and enclose them inside double quotation marks: set then as-path-prepend "1 1 1 1"

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Protocols > BGP >** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Protocols > BGP >** .	From the top of the CLI configuration hierarchy, enter edit protocols bgp
Apply the policy as an import policy for all BGP routes.	1. In the Import box, click **Add new entry**. 2. In the Value box, enter the name of the policy. 3. Click **OK**. The routing policy is evaluated when routes are being imported to the routing table.	Apply the policy: set import prependpolicy1 You can refer to the same routing policy one or more times in the same or a different import statement.

Configuring Damping Parameters

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Marking routes in this way leads to some delay, or *suppression* , in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (routes in different ASs). You can also apply flap damping within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.)

To change the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the damping action. For the damping routing policy to work, you also must enable BGP route flap damping.

To configure damping, perform these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 150.
3. If you are finished configuring the policy, commit the configuration.

Table 150: Creating a Policy to Accept and Apply Damping on Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Term** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options > Policy statement > Term**.	From the top of the CLI configuration hierarchy, enter edit policy-options policy-statement dampenpolicy1 term dampenterm1
Specify the routes to dampen.	1. In the From option, click **Configure**. 2. In the Value box, enter the prefixes you wish to dampen. 3. In the Route filter box, click **Add new entry**. 4. In the Value box, enter the prefixes you wish to dampen. 5. Click **OK**.	1. Dampen routes 172.168.0.0/16 or longer: set from route-filter 172.16.0.0/12 orlonger 2. Dampen routes 192.168.0.0/16 or longer: set from route-filter 192.168.0.0/16 orlonger 3. Dampen routes 10.0.0.0/8 or longer: set from route-filter 10.0.0.0/8 orlonger
Specify the damping parameters group to apply to the route filter.	1. In the configuration editor hierarchy, select **Policy options > Policy statement > Term**. 2. In the Then option, click **Configure**. 3. In the AS path prepend box, enter the string of AS numbers to prepend. Separate each AS number with a space. 4. Click **OK**.	Specify the AS numbers to prepend, and enclose inside them inside double quotation marks: set then as-path-prepend “1 1 1 1”
Navigate to the **Policy options** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Policy options**.	From the top of the CLI configuration hierarchy, enter edit policy-options

Task	J-Web Configuration Editor	CLI Configuration Editor
Create a damping parameter group.	1. In the Damping box, click **Add new entry**. 2. In the Damping object name box, enter the name of the damping parameter group. 3. Click **OK**.	Create and configure the damping parameter groups: edit damping group1 half-life 30 suppress 3000 reuse 750 max-suppress 60 edit damping group2 half-life 40 suppress 400 reuse 1000 max-suppress 45 edit damping group3 disable
Configure a damping parameter group.	1. In the Half life box, enter the half life duration, in minutes. 2. In the Max suppress box, enter the maximum holddown time, in minutes. 3. In the Reuse box, enter the reuse threshold, for this damping group. 4. In the Suppress box, enter the cutoff threshold, for this damping group. 5. To disable damping for this damping group, select the **Disable** check box. 6. Click **OK**.	
Navigate to the **BGP** level in the configuration hierarchy.	In the J-Web configuration editor hierarchy, select **Protocols > Bgp**.	From the top of the CLI configuration hierarchy, enter edit protocols bgp
Enable damping.	1. Select the **Damping** check box. 2. Click **OK**.	Enable damping: set damping
Navigate to the **Neighbor** level in the configuration hierarchy, for the BGP neighbor to which you want to apply the damping policy—for example, the neighbor at IP address 172.16.15.14.	In the J-Web configuration editor hierarchy, select **Protocols > Bgp > Group Group1 > Neighbor 172.16.15.14** .	From the top of the CLI configuration hierarchy, enter edit protocols bgp group group1 neighbor 172.16.15.14
Apply the policy as an import policy for the BGP neighbor.	1. In the Import box, click **Add new entry**. 2. In the Value box, enter the name of the policy. 3. Click **OK**. The routing policy is evaluated when routes are imported to the routing table.	Apply the policy: set import dampenpolicy1 You can refer to the same routing policy one or more times in the same or a different import statement.

Chapter 19
Configuring Firewall Filters and NAT

A *stateful* firewall filter inspects traffic flowing between a trusted network and an untrusted network. Contrasted with a *stateless* firewall filter that inspects packets in isolation, a stateful firewall filter provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions.

The Services Router uses the stateful firewall filter as a basis for performing Network Address Translation (NAT).

NOTE: You must have a license to configure a stateful firewall filter and NAT. For more information about licensing, see "Managing J-series Licenses" on page 69.

You can use either J-Web Quick Configuration or a configuration editor to configure stateful firewall filters and NAT. To configure a stateless firewall filter, use a configuration editor.

This chapter contains the following topics. For more information about firewall filters, see the *JUNOS Policy Framework Configuration Guide*. For more information about NAT, see the *JUNOS Services Interfaces Configuration Guide*.

Before You Begin

If you do not already have an understanding of firewall filters, read "Firewall Filter Overview" on page 358.

Before you begin configuring stateful firewall filters and NAT, you must configure the interfaces on which to apply these services. To configure an interface, see "Configuring Network Interfaces" on page 79.

Unlike a stateful firewall filter, you can configure a stateless firewall filter before configuring the interfaces on which they are applied.

Configuring a Stateful Firewall Filter with Quick Configuration

You can use the Firewall/NAT Quick Configuration pages to configure a stateful firewall filter and NAT. These Quick Configuration pages allow you to designate the interfaces that make up the untrusted network. In addition, you can designate the applications that are allowed to operate from the untrusted network to the trusted network.

Figure 93 and Figure 94 show the Firewall/NAT Quick Configuration main and application pages.

Figure 93: Firewall/NAT Quick Configuration Main Page

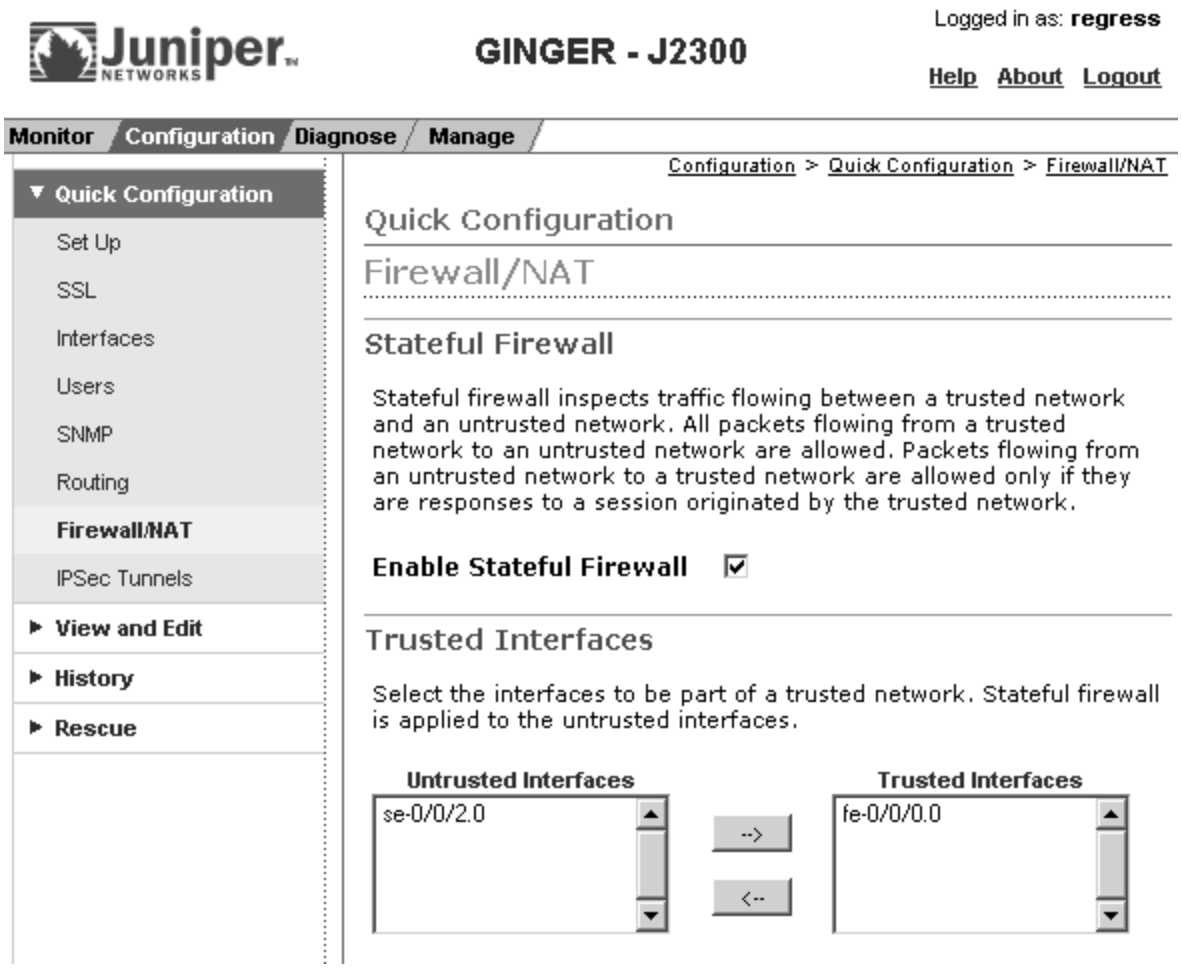

Figure 94: Firewall/NAT Quick Configuration Application Page

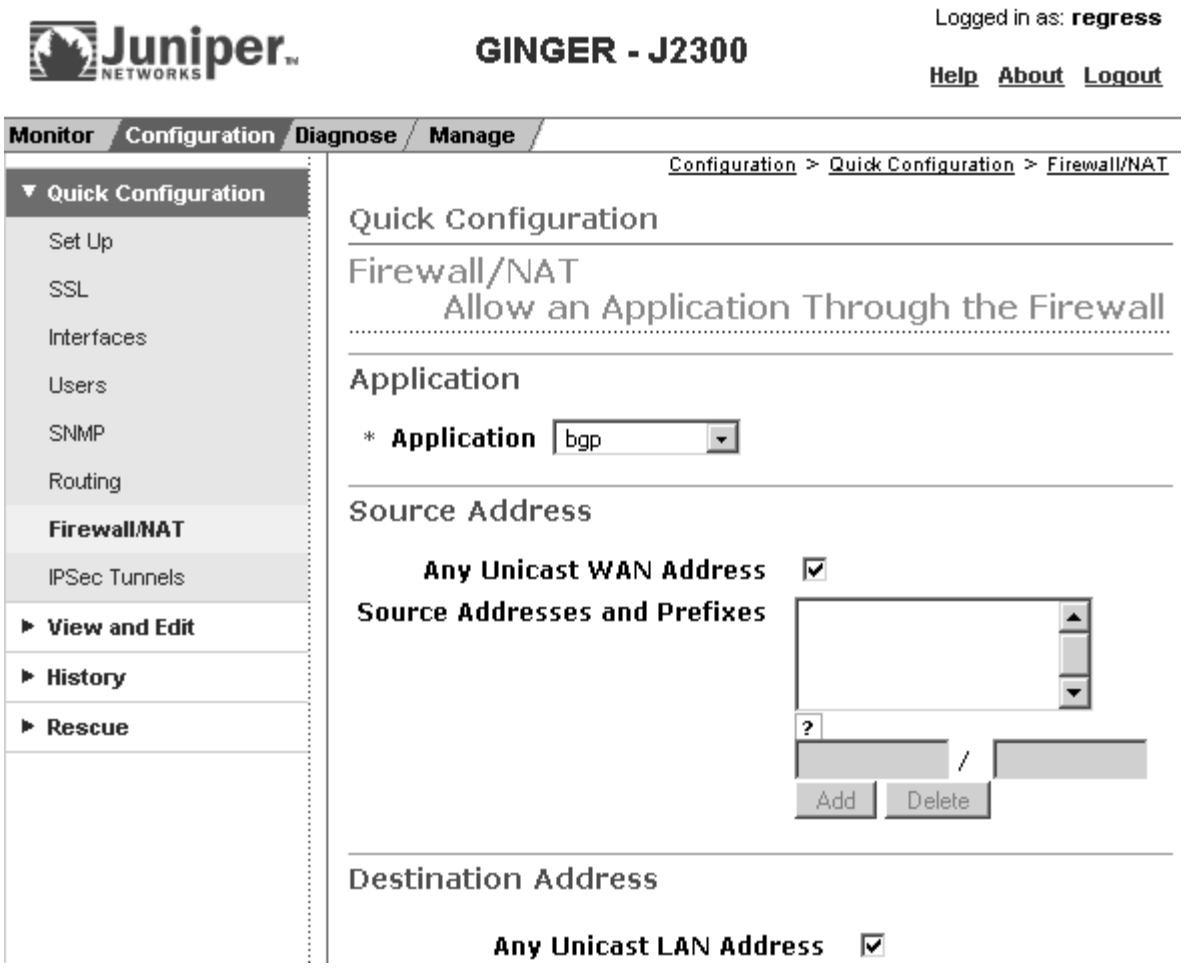

To configure a stateful firewall filter and NAT with Quick Configuration:

1. In the J-Web interface, select **Configuration > Firewall/NAT**.
2. Enter information into the Firewall/NAT Quick Configuration pages, as described in Table 151.
3. Click one of the following buttons on the Firewall/NAT Quick Configuration main page:
   - To apply the configuration and stay in the Firewall/NAT Quick Configuration main page, click **Apply**.
   - To apply the configuration and return to the Quick Configuration page, click **OK**.

- To cancel your entries and return to the Quick Configuration page, click **Cancel**.

4. Go on to one of the following procedures:

   - To display the configuration, see "Displaying Firewall Filter Configurations" on page 415.

   - To verify a stateful firewall filter, see "Verifying Firewall Filter Configuration" on page 415.

Table 151: Firewall/NAT Quick Configuration Pages Summary

Field	Function	Your Action
Stateful Firewall		
Enable Stateful Firewall	Enables stateful firewall filter configuration.	To enable stateful firewall filter configuration, select the check box.
Trusted Interfaces		
Trusted Interfaces	Designates the trusted and untrusted router interfaces. The stateful firewall filter is applied to the untrusted interfaces.	The Trusted Interfaces box displays a list of all the interfaces configured on the router. Do either of the following: ■ To *apply* a stateful firewall filter to an interface, click the interface in the Trusted Interfaces box to highlight it, and click the left arrow to add the interface to the Untrusted Interfaces list. You can select multiple interfaces by pressing Ctrl while you click the interface. ■ To *remove* a stateful firewall filter from an interface, click the interface in the Untrusted Interfaces box to highlight it, and click the right arrow to add the interface to the Trusted Interfaces list. You can select multiple interfaces by pressing Ctrl while you click the interface.
Network Address Translation (NAT)		
Enable NAT	Enables NAT configuration.	To enable NAT configuration, select the check box.
Low Address in Address Range (required)	Specifies the lowest address in the NAT pool address range. If a range of addresses is not specified, you can specify a single address or an IP prefix.	Type an IP address or prefix.
High Address in Address Range	Specifies the highest address in the NAT pool address range.	Type an IP address. The total range of addresses in the pool must be limited to a maximum of 32.
Outside Applications Allowed		

Field	Function	Your Action
	Add or delete applications that are allowed to operate from the untrusted network to the trusted network.	Click **Add** to move to the Firewall/NAT Quick Configuration application page. When you have finished entering information into this page, click **OK** to save it. To cancel your entries, click **Cancel**.
Application		
Application (required)	Designate which applications are allowed to operate from the untrusted network to the trusted network.	From the drop-down list, select the application you want to operate from the untrusted network to the trusted network.
Source Address		
Any Unicast WAN Address	Specifies that any unicast source address is allowed from the untrusted network.	To allow any unicast source address, select the check box.
Source Addresses and Prefixes	Designates the source addresses and prefixes that are allowed from the untrusted network.	To add an IP address and prefix, type them in the boxes above the **Add** button, then click **Add**. To delete an IP address and prefix, select them in the Source Addresses and Prefixes box, then click **Delete**.
Destination Address		
Any Unicast LAN Address	Specifies that any unicast destination address is allowed from the untrusted network.	To allow any unicast destination address, select the check box.
Destination Addresses and Prefixes	Designates the destination addresses and prefixes that are allowed from the untrusted network.	To add an IP address and prefix, type them in the boxes above the **Add** button, then click **Add**. To delete an IP address and prefix, select them in the Destination Addresses and Prefixes box, then click **Delete**.

Configuring a Stateful Firewall Filter with a Configuration Editor

To configure a stateful firewall filter and NAT with a configuration editor, you do the following:

- Define the filter's input and output rules.

NOTE: If a packet does not match any terms in a stateful firewall filter rule, the packet is discarded.

- Define an address pool and port pool for NAT.
- Define NAT input and output rules.
- Define a *service set* that includes the rules in the filter and NAT and the virtual sp-0/0/0 services interface.
- Finally, apply the service set to any interfaces on the Services Router that lead to or from the untrusted network.

The example in this section shows how to create a stateful firewall filter and NAT with the rules described in Table 152.

Table 152: Sample Stateful Firewall Filter and NAT Rules

Rule	Type	Term or Terms
to-wan-rule	Output	■ app-term—Accepts packets from any of the applications defined by the JUNOS default group junos-algs-outbound application set. ■ accept-all-term—Accepts packets that do not match app-term.
from-wan-rule	Input	■ wan-src-addr-term—Accepts input packets with a source prefix of 192.168.33.0/24. ■ discard-all-term—Discards all packets.
nat-to-wan-rule	Output	private-public-term—Translates the source address to an address within the pool 10.148.2.1 through 10.148.2.32 and dynamically translates the source port to a router-assigned port by means of NAPT

The example also assigns the name public-pool to the NAT address pool and NAPT router-assigned port.

In addition, the example creates the service set wan-service-set that includes the stateful firewall filter and NAT services and defines sp-0/0/0 as its service interface. Finally, wan-service-set is applied to the WAN interface to the untrusted network, t1-0/0/0.

For stateful firewall match conditions and actions, see “Summary of Stateful Firewall Filter and NAT Match Conditions and Actions” on page 360.

To configure a stateful firewall filter and NAT and apply them to the WAN interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web interface or the CLI configuration editor.
2. Perform the configuration tasks described in Table 153.
3. To apply the stateful firewall filter and NAT to the interface, perform the configuration tasks described in Table 154.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
   - To display the configuration, see "Displaying Firewall Filter Configurations" on page 415.
   - To verify the stateful firewall filter, see "Verifying a Stateful Firewall Filter" on page 420.

Table 153: Configuring a Stateful Firewall Filter and NAT

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Stateful firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Services > Stateful firewall**.	From the top of the configuration hierarchy, enter edit services stateful-firewall.
Define to-wan-rule and set its match direction.	1. Next to Rule, click **Add new entry**. 2. In the Rule name box, type to-wan-rule. 3. From the Match direction drop-down list, select **output**.	Set the rule name, match direction, term name, and match condition: set rule to-wan-rule match-direction output term app-term from application-sets junos-algs-outbound
Define app-term for the to-wan-rule rule.	1. Next to Term, click **Add new entry**. 2. In the Term name box, type app-term.	
Define the match condition for app-term—the default junos-algs-outbound application set.	1. Next to From, click **Configure**. 2. Next to Application sets, click **Add new entry**. 3. In the Application set name box, type junos-algs-outbound. 4. Click **OK** twice.	
Define an action for app-term.	1. On the Term app-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** twice.	Set the action: set rule to-wan-rule term app-term then accept

Task	J-Web Configuration Editor	CLI Configuration Editor
Define accept-all-term for to-wan-rule.	1. On the Rule to-wan-rule page, next to Term, click **Add new entry**. 2. In the Term name box, type accept-all-term.	Set the term name and the action: set rule to-wan-rule term accept-all-term then accept
Define an action for accept-all-term. The action is taken only if a packet does not match app-term.	1. Next to Then, click **Configure**. 2. From the Designation drop-down list, select **Accept**. 3. Next to Accept, select the check box. 4. Click **OK** three times.	
Define from-wan-rule and set its match direction.	1. On the Rule page, next to Rule, click **Add new entry**. 2. In the Rule name box, type from-wan-rule. 3. From the Match direction drop-down list, select **input**.	Set the rule name, match direction, term name, and the match condition: set rule from-wan-rule match-direction input term wan-src-addr-term from source-address 192.168.33.0/24
Define wan-src-addr-term for the from-wan-rule rule.	1. Next to Term, click **Add new entry**. 2. In the Term name box, type wan-src-addr-term.	
Define the match condition for wan-src-addr-term.	1. Next to From, click **Configure**. 2. Next to Source address, click **Add new entry**. 3. From the Address drop-down list, select **Enter Specific Value—>**. 4. In the Prefix box, type 192.168.33.0/24. 5. Click **OK** twice.	
Define an action for wan-src-addr-term.	1. On the Term wan-src-addr-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** twice.	Set the action: set rule from-wan-rule term wan-src-addr-term then accept
Define discard-all-term for from-wan-rule.	1. On the Rule from-wan-rule page, next to Term, click **Add new entry**. 2. In the Term name box, type discard-all-term.	Set the term name and the action: set rule from-wan-rule term discard-all-term then discard
Define an action for discard-all-term. The action is taken only if a packet does not match wan-src-addr-term.	1. Next to Then, click **Configure**. 2. From the Designation drop-down list, select **Discard**. 3. Click **OK** three times.	

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Nat** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Services**. 2. Next to NAT, click **Configure**.	From the top of the configuration hierarchy, enter edit services nat.
Define the public-pool address pool name and range.	1. Next to Pool, click **Add new entry**. 2. In the Pool name box, type public-pool. 3. From the Address choice drop-down list, select **Address range**. 4. In the High box, type 10.148.2.32. In the Low box, 10.148.2.1.	Set the address pool name and the range: set pool public-pool address-range low 10.148.2.1 high 10.148.2.32
Specify the NAT port pool to be automatically assigned by the router.	1. Next to Port, click **Configure**. 2. From the Port choice drop-down list, select **Automatic**. 3. Click **OK** twice.	Configure the source port translation to be automatic: set pool public-pool port automatic
Define nat-to-wan-rule and private-public-term.	1. On the Nat page, next to Rule, click **Add new entry**. 2. In the Rule name box, type nat-to-wan-rule. 3. From the Match direction drop-down list, select **output**. 4. Next to Term, select **Add new entry**. 5. In the Term name box, type private-public-term. 6. Next to Then, select **Configure**. 7. Next to Translated, select **Configure**. 8. In the Source pool box, type public-pool.	Set the rule name, match direction, term name, and the term's pool name: set rule nat-to-wan-rule match-direction output term private-public-term then translated source-pool public-pool
Set the NAT port translation type for private-public-term.	1. Next to Translation type, select the check box. 2. Select **Configure**. 3. From the Source drop-down list, select **dynamic**. 4. Click **OK** five times.	Set the NAT translation type: set rule nat-to-wan-rule match-direction output term private-public-term then translated translation-type source dynamic

Table 154: Applying a Stateful Firewall Filter and NAT to an Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Services** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Services**.	From the top of the configuration hierarchy, enter edit services.
Define wan-service-set and assign the stateful firewall filter rule to-wan-rule to the service set.	1. Next to Service set, click **Add new entry**. 2. In the Service set name box, type wan-service-set. 3. From the Stateful firewall rules choice drop-down list, select **Stateful firewall rules**. 4. Next to Stateful firewall rules, click **Add new entry**. 5. In the Rule name box, type to-wan-rule. 6. Click **OK**.	Define the service set and assign the rule: set service-set wan-service-set stateful-firewall-rules to-wan-rule
Assign the stateful firewall filter rule from-wan-rule to the service set.	1. Next to Stateful firewall rules, click **Add new entry**. 2. In the Rule name box, type from-wan-rule. 3. Click **OK**.	Define the service set and assign the rule: set service-set wan-service-set stateful-firewall-rules from-wan-rule
Assign the NAT rule nat-to-wan-rule to the service set.	1. From the Nat rules choice drop-down list, select **Nat rules**. 2. Next to Nat rules, click **Add new entry**. 3. In the Rule name box, type nat-to-wan-rule. 4. Click **OK**.	Assign the rule to the service set: set service-set wan-service-set nat-rules nat-to-wan-rule
Define the service set type and virtual interface sp–0/0/0 as the service interface for wan-service-set.	1. From the Service type choice drop-down list, select **Interface service**. 2. Next to Interface service, click **Configure**. 3. In the Service interface box, type sp-0/0/0. 4. Click **OK**.	Define the service set type and the service interface: set service-set wan-service-set interface-service service-interface sp-0/0/0

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the sp–0/0/0 service interface.	1. In the configuration editor hierarchy, select **interfaces**. 2. Next to Interface, click **Add new entry**. 3. In the Interface name box, type sp-0/0/0. 4. Next to Unit, click **Add new entry**. 5. In the Interface unit number box, type 0. 6. Next to Inet, select the check box. 7. Click **Configure**. 8. Click **OK**.	From the top of the configuration hierarchy, configure the interface: set interfaces sp-0/0/0 unit 0 family inet
From the Interfaces level of the configuration hierarchy, navigate to the **Inet** level of the T1 interface—the untrusted interface in this example—and apply wan-service-set to the input and output sides of the t1–0/0/0 interface.	1. In the configuration editor hierarchy, select **Interfaces > t1-0/0/0 > Unit > 0 > Family > Inet**. 2. Next to Service, click **Configure**. 3. Next to Input, click **Configure**. 4. Next to Service set, click **Add new entry**. 5. In the Service set name box, type wan-service-set. 6. Click **OK**. 7. Next to Output, click **Configure**. 8. Next to Service set, click **Add new entry**. 9. In the Service set name box, type wan-service-set. 10. Click **OK**.	From the top of the configuration hierarchy, apply the service set to the interface: set interfaces t1-0/0/0 unit 0 family inet service input service-set wan-service-set set interfaces t1-0/0/0 unit 0 family inet service output service-set wan-service-set

Configuring a Stateless Firewall Filter with a Configuration Editor

The section contains the following topics. For stateless firewall match conditions, actions, and modifiers, see “Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers” on page 363.

- Configuring a Routing Engine Firewall Filter to Handle Fragments on page 409
- Applying a Stateless Firewall Filter to an Interface on page 414

Stateless Firewall Filter Strategies

For best results, use the following sections to plan the purpose and contents of a stateless firewall filter before starting configuration.

CAUTION: If a packet does not match any terms in a stateless firewall filter rule, the packet is discarded. Take care that you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

Strategy for a Typical Stateless Firewall Filter

A primary goal of a typical stateless firewall filter is to protect the Routing Engine processes and resources from malicious or untrusted packets. You can configure a stateless firewall filter like the sample filter protect-RE to restrict traffic destined for the Routing Engine based on its source, protocol, and application. In addition, you can limit the traffic rate of packets destined for the Routing Engine to protect against flood, or *denial-of-service* (DoS), attacks.

For details, see "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 400 and "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 404.

Strategy for Handling Packet Fragments

You can configure a stateless firewall filter like the sample filter fragment-filter to address special circumstances associated with fragmented packets destined for the Routing Engine. Because the Services Router evaluates every packet against a firewall filter (including fragments), you must configure the filter to accommodate fragments that do not contain packet header information. Otherwise, the filter discards all but the first fragment of a fragmented packet.

For details, see "Configuring a Routing Engine Firewall Filter to Handle Fragments" on page 409.

Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources

The following example shows how to create a stateless firewall filter, protect-RE, that discards all traffic destined for the Routing Engine, except

SSH and BGP protocol packets from specified trusted sources. Table 155 lists the terms that are configured in this sample filter.

Table 155: Sample Stateless Firewall Filter protect-RE Terms to Allow Packets from Trusted Sources

Term	Purpose
ssh-term	Accepts TCP packets with a source address of 192.168.122.0/24 and a destination port that specifies SSH.
bgp-term	Accepts TCP packets with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol.
discard-rest-term	For all packets that are not accepted by ssh-term or bgp-term, creates a firewall filter log and system logging records, then discards all packets. To view the log, enter the show firewall log operational mode command. (For more information, see "Displaying Firewall Filter Logs" on page 421.)

By applying firewall filter protect-RE to the Routing Engine, you specify which protocols and services, or applications, are allowed to reach the Routing Engine, and you ensure the packets are from a trusted source. This protects processes running on the Routing Engine from an external attack.

To use the configuration editor to configure the stateless firewall filter:

1. Navigate to the top of the configuration hierarchy in either the J-Web interface or the CLI configuration editor.
2. Perform the configuration tasks described in Table 156.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To display the configuration, see "Displaying Firewall Filter Configurations" on page 415.
   - To apply the firewall filter to the Routing Engine, see "Applying a Stateless Firewall Filter to an Interface" on page 414.
   - To verify the firewall filter, see "Verifying a Services, Protocols, and Trusted Sources Firewall Filter" on page 423.

Table 156: Configuring a Protocols and Services Firewall Filter for the Routing Engine

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Firewall**.	From the top of the configuration hierarchy, enter edit firewall.

Task	J-Web Configuration Editor	CLI Configuration Editor
Define protect-RE and ssh-term, and define the protocol, destination port, and source address match conditions.	1. Next to Filter, click **Add new entry**. 2. In the Filter name box, type protect-RE. 3. Next to Term, click **Add New Entry**. 4. In the Rule name box, type ssh-term. 5. Next to From, click **Configure**. 6. In the Protocol choice drop-down list, select **Protocol**. 7. Next to Protocol, click **Add new entry**. 8. In the Value keyword drop-down list, select **tcp**. 9. Click **OK**. 10. In the Destination port choice drop-down list, select **Destination port**. 11. Next to Destination port, click **Add new entry**. 12. In the Value keyword drop-down list, select **ssh**. 13. Click **OK**. 14. Next to Source address, click **Add new entry**. 15. In the Address box, type 192.168.122.0/24. 16. Click **OK** twice.	Set the term name and define the match conditions: set family inet filter protect-RE term ssh-term from protocol tcp destination-port ssh source-address 192.168.122.0/24
Define the actions for ssh-term.	1. On the Term ssh-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** twice.	Set the actions: set family inet filter protect-RE term ssh-term then accept

Task	J-Web Configuration Editor	CLI Configuration Editor
Define bgp-term, and define the protocol, destination port, and source address match conditions.	1. On the Filter protect-RE page, next to Term, click **Add New Entry**. 2. In the Rule name box, type bgp-term. 3. Next to From, click **Configure**. 4. In the Protocol choice drop-down list, select **Protocol**. 5. Next to Protocol, click **Add new entry**. 6. In the Value keyword drop-down list, select **tcp**. 7. Click **OK**. 8. In the Destination port choice drop-down list, select **Destination port**. 9. Next to Destination port, click **Add new entry**. 10. In the Value keyword drop-down list, select **bgp**. 11. Click **OK**. 12. Next to Source address, click **Add new entry**. 13. In the Address box, type 10.2.1.0/24. 14. Click **OK** twice.	Set the term name and define the match conditions: set family inet filter protect-RE term bgp-term from protocol tcp destination-port bgp source-address 10.2.1.0/24
Define the action for bgp-term.	1. On the Term bgp-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** twice.	Set the action: set family inet filter protect-RE term bgp-term then accept
Define discard-rest-term and its action.	1. On the Filter protect-RE page, next to Term, click **Add New Entry**. 2. In the Rule name box, type discard-rest-term. 3. Next to Then, click **Configure**. 4. Next to Log, select the check box. 5. Next to Syslog, select the check box. 6. In the Designation drop-down list, select **Discard**. 7. Click **OK** four times.	Set the term name and define its actions: set family inet filter protect-RE term discard-rest-term then log syslog discard

Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods

The procedure in this section creates a sample stateless firewall filter, protect-RE, that limits certain TCP and ICMP traffic destined for the Routing Engine. A router without this kind of protection is vulnerable to TCP and ICMP flood attacks—also known as denial-of-service (DoS) attacks. For example:

- A TCP flood attack of SYN packets initiating connection requests can so overwhelm the Services Router that it can no longer process legitimate connection requests, resulting in denial of service.
- An ICMP flood can overload the Services Router with so many echo requests (ping requests) that it expends all its resources responding and can no longer process valid network traffic, also resulting in denial of service.

Applying a firewall filter like protect-RE to the Routing Engine protects against these types of attacks.

For each term in the sample filter, you first create a policer and then incorporate it into the action of the term. For more information about firewall filter policers, see the *JUNOS Policy Framework Configuration Guide*.

If you want to include the terms created in this procedure in the protect-RE firewall filter configured in the previous section (see "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 400), perform the configuration tasks in this section first, then configure the terms as described in the previous section. This approach ensures that the rate-limiting terms are included as the first two terms in the firewall filter.

NOTE: You can move terms within a firewall filter by using the insert CLI command. For more information, see "Inserting an Identifier" on page 152.

Table 157 lists the terms that are configured in this sample filter.

Table 157: Sample Stateless Firewall Filter protect-RE Terms to Protect Against Floods

Term	Purpose	Policer
tcp-connection-term	Polices the following types of TCP packets with a source address of 192.168.122.0/24 or 10.2.1.0/24: ■ Connection request packets (SYN and ACK flag bits equal 1 and 0) ■ Connection release packets (FIN flag bit equals 1) ■ Connection reset packets (RST flag bit equals 1)	tcp-connection-policer—Limits the traffic rate and burst size of these TCP packets to 500,000 bps and 15,000 bytes. Packets that exceed the traffic rate are discarded.
icmp-term	Polices the following types of ICMP packets. All are counted in counter icmp-counter. ■ Echo request packets ■ Echo response packets ■ Unreachable packets ■ Time-exceeded packets	icmp-policer—Limits the traffic rate and burst size of these ICMP packets to 1,000,000 bps and 15,000 bytes. Packets that exceed the traffic rate are discarded.

To use the configuration editor to configure the policers and the stateless firewall filter:

1. Navigate to the top of the configuration hierarchy in either the J-Web interface or the CLI configuration editor.
2. To configure the firewall filter policers, perform the configuration tasks described in Table 158.
3. To configure the prefix lists and the firewall filter, perform the configuration tasks described in Table 159.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
   - To display the configuration, see “Displaying Firewall Filter Configurations” on page 415.
   - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 414.
   - To verify the firewall filter, see “Verifying a TCP and ICMP Flood Firewall Filter” on page 424.

Table 158: Configuring Policers for TCP and ICMP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Firewall**.	From the top of the configuration hierarchy, enter edit firewall.
Define tcp-connection-policer and set its rate limits. You can use the following abbreviations when specifying the bandwidth limit: ■ k (1000) ■ m (1,000,000) ■ g (1,000,000,000)	1. Next to Policer, click **Add new entry**. 2. In the Policer name box, type tcp-connection-policer. 3. Next to Filter specific, select the check box. 4. Next to If Exceeding, select the check box and click **Configure**. 5. In the Burst size limit box, type 15k. The burst size limit can be from 1,500 through 100,000,000 bytes. 6. In the Bandwidth drop-down list, select **Bandwidth limit**. 7. In the Bandwidth limit box, type 500k. The bandwidth limit can be from 32,000 through 32,000,000,000 bps. 8. Click **OK**.	Set the policer name and its rate limits: set policer tcp-connection-policer filter-specific if-exceeding burst-size-limit 15k bandwidth-limit 500k
Define the policer action for tcp-connection-policer.	1. On the Policer tcp-connection-policer page, next to Then, click **Configure**. 2. Next to Discard, select the check box. 3. Click **OK** twice.	Set the policer action: set policer tcp-connection-policer then discard

Task	J-Web Configuration Editor	CLI Configuration Editor
Define icmp-policer and set its rate limits. You can use the following abbreviations when specifying the bandwidth limit: ■ k (1000) ■ m (1,000,000) ■ g (1,000,000,000)	1. On the Firewall page, next to Policer, click **Add new entry**. 2. In the Policer name box, type icmp-policer. 3. Next to Filter specific, select the check box. 4. Next to If Exceeding, select the check box and click **Configure**. 5. In the Burst size limit box, type 15k. The burst size limit can be from 1,500 through 100,000,000 bytes. 6. In the Bandwidth drop-down list, select **Bandwidth limit**. 7. In the Bandwidth limit box, type 1m. The bandwidth limit can be from 32,000 through 32,000,000,000 bps. 8. Click **OK**.	Set the policer name and its rate limits: set policer icmp-policer filter-specific if-exceeding burst-size-limit 15k bandwidth-limit 1m
Define the policer action for icmp-policer.	1. On the Policer icmp-policer page, next to Then, click **Configure**. 2. Next to Discard, select the check box. 3. Click **OK** three times.	Set the policer action: set policer icmp-policer then discard

Table 159: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Policy options** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Policy options**.	From the top of the configuration hierarchy, enter edit policy-options.
Define the prefix list trusted-addresses.	1. Next to Prefix list, click **Add new entry**. 2. In the Name box, type trusted-addresses. 3. Next to Prefix list item, click **Add new entry**. 4. In the Prefix box, type 192.168.122.0/24. 5. Click **OK**. 6. Next to Prefix list item, click **Add new entry**. 7. In the Prefix box, type 10.2.1.0/24. 8. Click **OK** three times.	Set the prefix list: set prefix-list trusted-addresses 192.168.122.0/24 set prefix-list trusted-addresses 10.2.1.0/24

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Firewall**.	From the top of the configuration hierarchy, enter edit firewall.
Define protect-RE and tcp-connection-term, and define the source prefix list match condition.	1. Next to Filter, click **Add new entry**. 2. In the Filter name box, type protect-RE. 3. Next to Term, click **Add New Entry**. 4. In the Rule name box, type tcp-connection-term. 5. Next to From, click **Configure**. 6. Next to Source prefix list, click **Add new entry**. 7. In the Name box, type trusted-addresses. 8. Click **OK**.	Set the term name and define the source address match condition: set family inet filter protect-RE term tcp-connection-term from source-prefix-list trusted-addresses
Define the TCP flags and protocol match conditions for tcp-connection-term.	1. In the TCP flags box, type (syn & !ack) | fin | rst. 2. In the Protocol choice drop-down list, select **Protocol**. 3. Next to Protocol, click **Add new entry**. 4. In the Value keyword drop-down list, select **tcp**. 5. Click **OK**.	Set the TCP flags and protocol and protocol match conditions for the term: set family inet filter protect-RE term tcp-connection-term from protocol tcp tcp-flags "(syn & !ack) | fin | rst"
Define the actions for tcp-connection-term.	1. On the Term tcp-connection-term page, next to Then, click **Configure**. 2. In the Policer box, type tcp-connection-policer. 3. In the Designation drop-down list, select **Accept**. 4. Click **OK** twice.	Set the actions: set family inet filter protect-RE term tcp-connection-term then policer tcp-connection-policer accept
Define icmp-term, and define the protocol.	1. On the Filter protect-RE page, next to Term, click **Add New Entry**. 2. In the Rule name box, type icmp-term. 3. Next to From, click **Configure**. 4. In the Protocol choice drop-down list, select **Protocol**. 5. Next to Protocol, click **Add new entry**. 6. In the Value keyword drop-down list, select **icmp**. 7. Click **OK**.	Set the term name and define the protocol: set family inet filter protect-RE term icmp-term from protocol icmp

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the ICMP type match conditions.	1. In the Icmp type choice drop-down list, select **Icmp type**. 2. Next to Icmp type, click **Add new entry**. 3. In the Value keyword drop-down list, select **echo-request**. 4. Click **OK**. 5. Next to Icmp type, click **Add new entry**. 6. In the Value keyword drop-down list, select **echo-reply**. 7. Click **OK**. 8. Next to Icmp type, click **Add new entry**. 9. In the Value keyword drop-down list, select **unreachable**. 10. Click **OK**. 11. Next to Icmp type, click **Add new entry**. 12. In the Value keyword drop-down list, select **time-exceeded**. 13. Click **OK**.	Set the ICMP type match conditions: set family inet filter protect-RE term icmp-term from icmp-type [echo-request echo-reply unreachable time-exceeded]
Define the actions for icmp-term.	1. On the icmp-term page, next to Then, click **Configure**. 2. In the Count box, type icmp-counter. 3. In the Policer box, type icmp-policer. 4. In the Designation drop-down list, select **Accept**. 5. Click **OK** four times.	Set the actions: set family inet filter protect-RE term icmp-term then policer icmp-policer count icmp-counter accept

Configuring a Routing Engine Firewall Filter to Handle Fragments

The procedure in this section creates a sample stateless firewall filter, fragment-RE, that handles fragmented packets destined for the Routing Engine. By applying fragment-RE to the Routing Engine, you protect against the use of IP fragmentation as a means to disguise TCP packets from a firewall filter.

Table 160 lists the terms that are configured in this sample filter.

Table 160: Sample Stateless Firewall Filter fragment-RE Terms

Term	Purpose
small-offset-term	Discards IP packets with a fragment offset of 1 through 5, and adds a record to the system logging facility.
not-fragmented-term	Accepts unfragmented TCP packets with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol. A packet is considered unfragmented if its MF flag and its fragment offset in the TCP header equal 0.
first-fragment-term	Accepts the first fragment of a fragmented TCP packet with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol.
fragment-term	Accepts all packet fragments with an offset of 6 through 8191.

For example, consider an IP packet that is fragmented into the smallest allowable fragment size of 8 bytes (a 20-byte IP header plus an 8-byte payload). If this IP packet carries a TCP packet, the first fragment (fragment offset of 0) that arrives at the Services Router contains only the TCP source and destination ports (first 4 bytes), and the sequence number (next 4 bytes). The TCP flags, which are contained in the next 8 bytes of the TCP header, arrive in the second fragment (fragment offset of 1). The fragment-RE filter works as follows:

- Term small-offset-term discards small offset packets to ensure that subsequent terms in the firewall filter can be matched against all the headers in the packet.
- Term fragment-term accepts all fragments that were not discarded by small-offset-term. However, only those fragments that are part of a packet containing a first fragment accepted by first-fragment-term are reassembled by the Services Router.

For more information about IP fragment filtering, see RFC 1858, *Security Considerations for IP Fragment Filtering*.

To use the configuration editor to configure the stateless firewall filter:

1. Navigate to the top of the configuration hierarchy in either the J-Web interface or the CLI configuration editor.
2. To configure the firewall filter, perform the configuration tasks described in Table 161.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To display the configuration, see “Displaying Firewall Filter Configurations” on page 415.
   - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 414.
   - To verify the firewall filter, see “Verifying a Firewall Filter That Handles Fragments” on page 425.

Table 161: Configuring a Fragments Firewall Filter for the Routing Engine

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Firewall**.	From the top of the configuration hierarchy, enter edit firewall.
Define fragment-RE and small-offset-term, and define the fragment offset match condition. The fragment offset can be from 1 through 8191.	1. Next to Filter, click **Add new entry**. 2. In the Filter name box, type fragment-RE. 3. Next to Term, click **Add New Entry**. 4. In the Rule name box, type small-offset-term. 5. Next to From, click **Configure**. 6. In the Fragment offset choice drop-down list, select **Fragment offset**. 7. Next to Fragment offset, select **Add New Entry**. 8. In the Range box, type 1-5. 9. Click **OK** twice.	Set the term name and define the fragment offset match condition: set family inet filter fragment-RE term small-offset-term from fragment-offset 1-5
Define the action for small-offset-term.	1. On the Term small-offset-term page, next to Then, click **Configure**. 2. Next to Syslog, select the check box. 3. In the Designation drop-down list, select **Discard**. 4. Click **OK** twice.	Set the action: set family inet filter fragment-RE term small-offset-term then syslog discard

Task	J-Web Configuration Editor	CLI Configuration Editor
Define not-fragmented-term, and define the fragment, protocol, destination port, and source address match conditions.	1. On the Filter fragment-RE page, next to Term, click **Add New Entry**. 2. In the Term name box, type not-fragmented-term. 3. Next to From, click **Configure**. 4. In the Fragment flags box, type 0x0. 5. In the Fragment offset choice drop-down list, select **Fragment offset**. 6. Next to Fragment offset, select **Add New Entry**. 7. In the Range box, type 0. 8. Click **OK**. 9. In the Protocol choice drop-down list, select **Protocol**. 10. Next to Protocol, click **Add new entry**. 11. In the Value keyword drop-down list, select **tcp**. 12. Click **OK**. 13. In the Destination port choice drop-down list, select **Destination port**. 14. Next to Destination port, click **Add new entry**. 15. In the Value keyword drop-down list, select **bgp**. 16. Click **OK**. 17. Next to Source address, click **Add new entry**. 18. In the Address box, type 10.2.1.0/24. 19. Click **OK** twice.	Set the term name and define match conditions: set family inet filter fragment-RE term not-fragmented-term from fragment-flags 0x0 fragment-offset 0 protocol tcp destination-port bgp source-address 10.2.1.0/24
Define the action for not-fragmented-term.	1. On the Term not-fragmented-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** twice.	Set the action: set family inet filter fragment-RE term not-fragmented-term then accept

Task	J-Web Configuration Editor	CLI Configuration Editor
Define first-fragment-term, and define the fragment, protocol, destination port, and source address match conditions.	1. On the Filter fragment-RE page, next to Term, click **Add New Entry**. 2. In the Rule name box, type first-fragment-term. 3. Next to From, click **Configure**. 4. Next to First fragment, select the check box. 5. In the Protocol choice drop-down list, select **Protocol**. 6. Next to Protocol, click **Add new entry**. 7. In the Value keyword drop-down list, select **tcp**. 8. Click **OK**. 9. In the Destination port choice drop-down list, select **Destination port**. 10. Next to Destination port, click **Add new entry**. 11. In the Value keyword drop-down list, select **bgp**. 12. Click **OK**. 13. Next to Source address, click **Add new entry**. 14. In the Address box, type 10.2.1.0/24. 15. Click **OK** twice.	Set the term name and define match conditions: set family inet filter fragment-RE term first-fragment-term from first-fragment protocol tcp destination-port bgp source-address 10.2.1.0/24
Define the action for first-fragment-term.	1. On the Term first-fragment-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** twice.	Set the action: set family inet filter fragment-RE term first-fragment-term then accept

Task	J-Web Configuration Editor	CLI Configuration Editor
Define fragment-term and define the fragment match condition.	1. On the Filter fragment-RE page, next to Term, click **Add New Entry**. 2. In the Rule name box, type fragment-term. 3. Next to From, click **Configure**. 4. In the Fragment offset choice drop-down list, select **Fragment offset**. 5. Next to Fragment offset, select **Add New Entry**. 6. In the Range box, type 6-8191. 7. Click **OK** twice.	Set the term name and define match conditions: set family inet filter fragment-RE term fragment-term from fragment-offset 6–8191
Define the action for fragment-term.	1. On the Term fragment-term page, next to Then, click **Configure**. 2. In the Designation drop-down list, select **Accept**. 3. Click **OK** four times.	Set the action: set family inet filter fragment-RE term fragment-term then accept

Applying a Stateless Firewall Filter to an Interface

You can apply a stateless firewall to the input or output sides, or both, of an interface. To filter packets transiting the router, apply the firewall filter to any non-Routing Engine interface. To filter packets originating from, or destined for, the Routing Engine, apply the firewall filter to the loopback (lo0) interface.

For example, to apply a stateless firewall filter protect-RE to the input side of the Routing Engine interface, follow this procedure:

1. Perform the configuration tasks described in Table 162.
2. If you are finished configuring the network, commit the configuration.

Table 162: Applying a Firewall Filter to the Routing Engine Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Inet** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Interfaces > lo0 > Unit > 0 > Family > Inet**.	From the top of the configuration hierarchy, apply the filter to the interface: set interfaces lo0 unit 0 family inet filter input protect-RE
Apply protect-RE as an input filter to the lo0 interface.	1. Next to Filter, click **Configure**. 2. In the Input box, type protect-RE. 3. Click **OK** five times.	

To view the configuration of the Routing Engine interface, enter the show interfaces lo0 command. For example:

```
user@host# show interfaces lo0
unit 0 {
   family inet {
      filter {
                     input protect-RE;
      }
                  address 127.0.0.1/32;
   }
}
```

Verifying Firewall Filter Configuration

To verify a firewall filter configuration, perform these tasks:

- Displaying Firewall Filter Configurations on page 415
- Verifying a Stateful Firewall Filter on page 420
- Displaying Firewall Filter Logs on page 421
- Displaying Firewall Filter Statistics on page 422
- Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 423
- Verifying a TCP and ICMP Flood Firewall Filter on page 424
- Verifying a Firewall Filter That Handles Fragments on page 425

Displaying Firewall Filter Configurations

Purpose Verify the configuration of the firewall filter. You can analyze the flow of the firewall filter terms by displaying the entire configuration.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the show services or show firewall command for stateful and stateless firewall filters.

The sample output in this section displays the following firewall filters (in order):

- Stateful firewall filter and NAT configured in “Configuring a Stateful Firewall Filter with a Configuration Editor” on page 393
- Stateless protect-RE filter configured in “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 400
- Stateless protect-RE filter configured in “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 404
- Stateless fragment-RE filter configured in “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 409

Sample Output

```
[edit]
user@host# show services
stateful-firewall
   rule to-wan-rule {
                match-direction output;
      term app-term {
         from {
                        application-sets junos-algs-outbound;
         }
         then {
                        accept;
         }
      }
      term accept-all-term {
         then {
                        accept;
         }
      }
   }
   rule from-wan-rule {
                match-direction input;
      term wan-src-addr-term {
         from {
            source-address {
                           192.168.33.0/24;
            }
         }
         then {
                        accept;
         }
      }
      term discard-all-term {
         then {
                        discard;
         }
      }
   }
}
nat {
   pool public-pool {
```

```
                address-range low 10.148.2.1 high 10.148.2.32;
                port automatic;
    }
    rule nat-to-wan-rule {
                match-direction output;
        term private-public-term {
            then {
                translated {
                            source-pool public-pool;
                            translation-type source dynamic;
                }
            }
        }
    }
}
service-set wan-service-set {
              stateful-firewall-rules to-wan-rule;
              stateful-firewall-rules from-wan-rule;
              nat-rules nat-to-wan-rule;
    interface-service {
                service-interface sp-0/0/0;
    }
}


[edit]
user@host# show firewall
firewall {
    family inet {
        filter protect-RE {
            term ssh-term {
                from {
                    source-address {
                                192.168.122.0/24;
                    }
                            protocol tcp;
                            destination-port ssh;
                }
                        then accept;
            }
            term bgp-term {
                from {
                    source-address {
                                10.2.1.0/24;
                    }
                            protocol tcp;
                            destination-port bgp;
                }
                        then accept;
            }
            term discard-rest-term {
                then {
                            log;
                            syslog;
                            discard;
```

```
            }
          }
        }
      }
    }

[edit]
user@host# show firewall
firewall {
    policer tcp-connection-policer {
                filter-specific;
        if-exceeding {
                    bandwidth-limit 500k;
                    burst-size-limit 15k;
        }
                then discard;
    }
    policer icmp-policer {
                filter-specific;
        if-exceeding {
                    bandwidth-limit 1m;
                    burst-size-limit 15k;
        }
                then discard;
    }
    family inet {
        filter protect-RE {
            term tcp-connection-term {
                from {
                    source-prefix-list {
                                trusted-addresses;
                    }
                            protocol tcp;
                            tcp-flags “(syn & !ack) | fin | rst”;
                }
                then {
                            policer tcp-connection-policer;
                            accept;
                }
            }
            term icmp-term {
                from {
                            protocol icmp;
                            icmp-type [ echo-request echo-reply unreachable time-exceeded ];
                }
                then {
                            policer icmp-policer;
                            count icmp-counter;
                            accept;
                }
            }
                        additional terms ...
        }
    }
```

```
}


[edit]
user@host# show firewall
firewall {
   family inet {
      filter fragment-RE {
         term small-offset-term {
            from {
                        fragment-offset 1-5;
            }
            then {
                        syslog;
                        discard;
            }
         }
         term not-fragmented-term {
            from {
               source-address {
                           10.2.1.0/24;
               }
                        fragment-offset 0;
                        fragment-flags 0x0;
                        protocol tcp;
                        destination-port bgp;
            }
                     then accept;
         }
         term first-fragment-term {
            from {
               source-address {
                           10.2.1.0/24;
               }
                        first-fragment;
                        protocol tcp;
                        destination-port bgp;
            }
                     then accept;
         }
         term fragment-term {
            from {
                        fragment-offset 6-8191;
            }
                     then accept;
         }
                  additional terms ...
      }
   }
}
```

What It Means Verify that the output shows the intended configuration of the firewall filter. For more information about the format of a configuration file, see "Viewing the Configuration Text" on page 136.

Verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the insert CLI command. For more information, see “Inserting an Identifier” on page 152.

Verifying a Stateful Firewall Filter

Purpose Verify the firewall filter configured in “Configuring a Stateful Firewall Filter with a Configuration Editor” on page 393.

Action To verify that the actions of the firewall filter terms are taken, send packets to and from the untrusted network that match the terms. In addition, verify that actions are *not* taken for packets that do not match.

- Send packets—associated with the junos-algs-outbound application set—from a host in the trusted network to a host in the untrusted network. Verify that packets received from the host in the untrusted network are responses only to the session originated by the host in the trusted network. To ensure that packets from the host are not accepted because of rule from-wan-rule, do not send packets to the host in the untrusted network with an IP address that matches 192.168.33.0/24.

  For example, send a ping request from host trusted-nw-trusted-host to host untrusted-nw-untrusted-host, and verify that a ping response is returned. Ping requests and responses use ICMP, which belongs to the junos-algs-outbound application set.

NOTE: To view the configuration of junos-algs-outbound, enter the show groups junos-defaults applications application-set junos-algs-outbound configuration mode command.

- Send packets from a host in the untrusted network to a host in the trusted network. Verify that the host in the trusted network receives packets only from the host in the untrusted network with an IP address that matches 192.168.33.0/24.

  For example, send a ping request from host untrusted-nw-trusted-host with an IP address that matches 192.168.33.0/24 to host trusted-nw-trusted-host, and verify that a ping response is returned.

  Verify that the ping response displays an IP address from the configured NAT pool.

Sample Output

```
user@trusted-nw-trusted-host> ping untrusted-nw-untrusted-host

PING untrusted-nw-untrusted-host.acme.net (172.69.13.5): 56 data bytes
64 bytes from 192.169.13.5: icmp_seq=0 ttl=22 time=8.238 ms
64 bytes from 192.169.13.5: icmp_seq=1 ttl=22 time=9.116 ms
64 bytes from 192.169.13.5: icmp_seq=2 ttl=22 time=10.875 ms
...
```

```
user@untrusted-nw-trusted-host> ping trusted-nw-trusted-host

PING trusted-nw-trusted-host-fe-000.acme.net (112.148.2.3): 56 data bytes
64 bytes from 10.148.2.3: icmp_seq=0 ttl=253 time=18.248 ms
64 bytes from 10.148.2.3: icmp_seq=1 ttl=253 time=10.906 ms
64 bytes from 10.148.2.3: icmp_seq=2 ttl=253 time=12.845 ms
...
```

What It Means Verify the following information:

- A ping request from host trusted-nw-trusted-host returns a ping response from host untrusted-nw-untrusted-host.
- A ping request from host untrusted-nw-trusted-host returns a ping response from host trusted-nw-trusted-host. Verify that the ping response displays an IP address from the configured NAT pool of 10.148.2.1 through 10.148.2.32.

For information about using the J-Web interface to ping a host, see "Using the J-Web Ping Host Tool" on page 218.

For more information about the ping command, see "Using the ping Command" on page 226 or the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Displaying Firewall Filter Logs

Purpose Verify that packets are being logged. If you included the log or syslog action in a term, verify that packets matching the term are recorded in the firewall log or your system logging facility.

Action From operational mode in the CLI, enter the show firewall log command.

The log of discarded packets generated from the firewall filter configured in "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 400 is displayed in the following sample output.

Sample Output

```
user@host> show firewall log

Log :
Time      Filter    Action Interface     Protocol Src Addr      Dest Addr
15:11:02  pfe       D      fe-0/0/0.0    TCP      172.17.28.19  192.168.70.71
15:11:01  pfe       D      fe-0/0/0.0    TCP      172.17.28.19  192.168.70.71
15:11:01  pfe       D      fe-0/0/0.0    TCP      172.17.28.19  192.168.70.71
15:11:01  pfe       D      fe-0/0/0.0    TCP      172.17.28.19  192.168.70.71
...
```

What It Means Each record of the output contains information about the logged packet. Verify the following information:

- Under Time, the time of day the packet was filtered is shown.
- The Filter output is always pfe.
- Under Action, the configured action of the term matches the action taken on the packet—A (accept), D (discard), R (reject).
- Under Interface, the ingress interface on which the packet arrived is appropriate for the filter.
- Under Protocol, the protocol in the IP header of the packet is appropriate for the filter.
- Under Src Addr, the source address in the IP header of the packet is appropriate for the filter.
- Under Dest Addr, the destination address in the IP header of the packet is appropriate for the filter.

For more information about the show firewall log command, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Displaying Firewall Filter Statistics

Purpose Verify that packets are being policed and counted.

Action From operational mode in the CLI, enter the show firewall filter *filter-name* command.

The value of the counter, icmp-counter, and the number of packets discarded by the policers in the firewall filter configured in "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 404 are displayed in the following sample output.

Sample Output

```
user@host> show firewall filter protect-RE

Filter: protect-RE
Counters:
Name                                                Bytes              Packets
icmp-counter                                      1040000                 5600
Policers:
Name                                                Packets
tcp-connection-policer                            643254873
icmp-policer                                         7391
```

What It Means Verify the following information:

- Next to Filter, the name of the firewall filter is correct.
- Under Counters:
  - Under Name, the names of any counters configured in the firewall filter are correct.
  - Under Bytes, the number of bytes that match the filter term containing the count *counter-name* action are shown.
  - Under Packets, the number of packets that match the filter term containing the count *counter-name* action are shown.
- Under Policers:
  - Under Name, the names of any policers configured in the firewall filter are correct.
  - Under Packets, the number of packets that match the conditions specified for the policer are shown.

For more information about the show firewall filter command, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying a Services, Protocols, and Trusted Sources Firewall Filter

Purpose Verify the firewall filter configured in "Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources" on page 400.

Action To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Use the ssh *host-name* command from a host at an IP address that matches 192.168.122.0/24 to verify that you can log in to the Services Router using only SSH from a host with this address prefix.
- Use the show route summary command to verify that the routing table on the Services Router does not contain any entries with a protocol other than Direct, Local, BGP, or Static.

Sample Output

```
% ssh 192.168.249.71

%ssh host
user@host's password:
--- JUNOS 6.4-20040518.0 (JSERIES) #0: 2004-05-18 09:27:50 UTC

user@host>
```

```
user@host> show route summary

Router ID: 192.168.249.71

inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
              Direct:     10 routes,      9 active
               Local:      9 routes,      9 active
                 BGP:     10 routes,     10 active
              Static:      5 routes,      5 active
...
```

What It Means Verify the following information:

- You can successfully log in to the Services Router using SSH.
- The show route summary command does not display a protocol other than Direct, Local, BGP, or Static.

For more information about the show route summary command, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying a TCP and ICMP Flood Firewall Filter

Purpose Verify the firewall filter configured in "Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods" on page 404.

Action To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Verify that the Services Router can establish only TCP sessions with a host at an IP address that matches 192.168.122.0/24 or 10.2.1.0/24. For example, log in to the router with the telnet *host-name* command from another host with one of these address prefixes.
- Use the ping *host-name* command to verify that the Services Router responds only to ICMP packets (such as ping requests) that do not exceed the policer traffic rates.
- Use the ping *host-name* size *bytes* command to exceed the policer traffic rates by sending ping requests with large data payloads.

Sample Output

```
user@host> telnet 192.168.249.71

Trying 192.168.249.71...
Connected to host.acme.net.
Escape character is '^]'.

host (ttyp0)

login: user
Password:

--- JUNOS 6.4-20040521.1 built 2004-05-21 09:38:12 UTC
```

```
user@host>

user@host> ping 192.168.249.71

PING host-fe-000.acme.net (192.168.249.71): 56 data bytes
64 bytes from 192.168.249.71: icmp_seq=0 ttl=253 time=11.946 ms
64 bytes from 192.168.249.71: icmp_seq=1 ttl=253 time=19.474 ms
64 bytes from 192.168.249.71: icmp_seq=2 ttl=253 time=14.639 ms
...


user@host> ping 192.168.249.71 size 20000

PING host-fe-000.acme.net (192.168.249.71): 20000 data bytes
^C
--- host-fe-000.acme.net ping statistics ---
12 packets transmitted, 0 packets received, 100% packet loss
```

What It Means Verify the following information:

- You can successfully log in to the Services Router using Telnet.
- The Services Router sends responses to the ping host command.
- The Services Router does not send responses to the ping host size 20000 command.

For more information about the ping command, see "Using the ping Command" on page 226 or the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

For information about using the J-Web interface to ping a host, see "Using the J-Web Ping Host Tool" on page 218.

For more information about the telnet command, see "Using the telnet Command" on page 195 or the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying a Firewall Filter That Handles Fragments

Purpose Verify the firewall filter configured in "Configuring a Routing Engine Firewall Filter to Handle Fragments" on page 409.

Action To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Verify that packets with small fragment offsets are recorded in the router's system logging facility.
- Use the show route summary command to verify that the routing table does not contain any entries with a protocol other than Direct, Local, BGP, or Static.

Sample Output

```
user@host> show route summary

Router ID: 192.168.249.71

inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
              Direct:     10 routes,      9 active
               Local:      9 routes,      9 active
                 BGP:     10 routes,     10 active
              Static:      5 routes,      5 active
...
```

What It Means Verify that the show route summary command does not display a protocol other than Direct, Local, BGP, or Static. For more information about the show route summary command, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Chapter 20

Configuring Class of Service with DiffServ

You configure class of service (CoS) with Differentiated Services (DiffServ) when you need to override the default packet forwarding behavior of a Services Router—especially in the three areas identified in Table 163.

Table 163: Reasons to Configure Class of Service (Cos) with DiffServ

Default Behavior to Override with CoS	CoS Configuration Area
Packet classification—By default, the Services Router does not use DiffServ to classify packets. Packet classification applies to incoming traffic.	Classifiers
Scheduling queues—By default, the Services Router has only two queues enabled. Scheduling queues apply to outgoing traffic.	Schedulers
Packet headers—By default, the Services Router does not rewrite CoS bits in packet headers. Rewriting packet headers applies to outgoing traffic.	Rewrite rules

You can use either the J-Web configuration editor or CLI configuration editor to configure CoS with DiffServ. The J-Web interface does not include Quick Configuration pages for CoS or DiffServ.

This chapter contains the following topics. For more information about CoS and DiffServ, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Before You Begin

Before you begin configuring a Services Router for CoS with DiffServ, complete the following tasks:

- If you do not already have a basic understanding of CoS and DiffServ, read "Policy, Firewall Filter, and Class-of-Service Overview" on page 351.
- Determine whether the Services Router needs to support different traffic streams, such as voice or video. If so, CoS with DiffServ helps to make sure this traffic receives more than basic best-effort packet delivery service.
- Determine whether the Services Router is directly attached to any applications that send DiffServ packets. If no sources are enabled for DiffServ, you must configure and apply rewrite rules on the interfaces to the sources.
- Determine whether the Services Router must support DiffServ assured forwarding (AF) classes. Assured forwarding usually requires random early detection (RED) drop profiles to be configured and applied.
- Determine whether the Services Router must support DiffServ expedited forwarding (EF) classes with a policer. Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a high loss priority, so that packets exceeding the policer limits are discarded first.

Configuring CoS with DiffServ with a Configuration Editor

To configure the Services Router as a node in a network supporting CoS with DiffServ, you must perform the following tasks marked *(Required)*.

- (Required) "Configuring a Policer for a Firewall Filter" on page 429
- (Required) "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 430
- (Required) "Assigning Forwarding Classes to Output Queues" on page 434
- (Required) "Configuring and Applying Rewrite Rules" on page 435
- (Required) "Configuring and Applying Behavior Aggregate Classifiers" on page 440
- (Required) "Configuring RED Drop Profiles for Assured Forwarding Congestion Control" on page 443
- (Optional) "Configuring Schedulers" on page 446
- (Optional) "Configuring and Applying Scheduler Maps" on page 450
- (Optional) "Configuring and Applying Virtual Channels" on page 453

For information about using the J-Web and CLI configuration editors, see "Using J-series Configuration Tools" on page 127.

Configuring a Policer for a Firewall Filter

You configure a policer to detect packets that exceed the limits established for DiffServ expedited forwarding. For DiffServ, packets that exceed these limits are given a higher loss priority than packets within the bandwidth and burst size limits.

The following example shows how to configure a policer called ef-policer that identifies for likely discard expedited forwarding packets with a burst size greater than 2000 bytes and a bandwidth greater than 10 percent.

For more information about firewall filters, see "Configuring Firewall Filters and NAT" on page 389 and the *JUNOS Policy Framework Configuration Guide*.

To configure an expedited forwarding policer for a firewall filter for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 164.
3. Go on to "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 430.

Table 164: Configuring a Policer for a Firewall Filter

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Firewall**.	From the top of the configuration hierarchy, enter edit firewall
Create and name the policer for expedited forwarding.	1. Click **Add new entry** next to Policer. 2. In the Policer name box, type a name for the EF policer—for example, ef-policer.	Enter edit policer ef-policer

Task	J-Web Configuration Editor	CLI Configuration Editor
Enter the burst limit and bandwidth for the policer.	1. Click **Configure** next to If exceeding. 2. In the Burst size limit box, type a limit for the burst size allowed—for example, 2k. 3. From the **Bandwidth** list, select a limit or percentage—for example, **bandwidth-percent**. 4. In the Bandwidth percent box, type a percentage for the bandwidth allowed for this type of traffic—for example, 10. 5. Click **OK**.	Enter set if-exceeding burst-limit-size 2k set if-exceeding bandwidth-percent 10
Enter the loss priority for packets exceeding the limits established by the policer.	1. Click **Configure** next to Then. 2. From the Loss priority list, select **high**. 3. Click **OK** three times.	Enter set then loss-priority high

Configuring and Applying a Firewall Filter for a Multifield Classifier

You configure a multifield (MF) classifier to detect packets of interest to CoS and assign the packet to the proper forwarding class independently of the DiffServ code point (DSCP). To configure a multifield classifier on a customer-facing or host-facing link, configure a firewall filter to classify traffic. Packets are classified as they arrive on an interface.

One common way to detect packets of CoS interest is by source or destination address. The destination address is used in this example, but many other matching criteria for packet detection are available to firewall filters.

This example shows how to configure the firewall filter mf-classifier and apply it to the Services Router's Fast Ethernet interface fe-0/0/0. The firewall filter consists of the rules (terms) listed in Table 165.

Table 165: Sample mf-classifier Firewall Filter Terms

Rule (Term)	Purpose	Contents
assured forwarding	Detects packets destined for 192.168.44.55, assigns them to an assured forwarding class, and gives them a low likelihood of being dropped.	Match condition: destination address 192.168.44.55 Forwarding class: af-class Loss priority: low

Rule (Term)	Purpose	Contents
expedited-forwarding	Detects packets destined for 192.168.66.77, assigns them to an expedited forwarding class, and subjects them to the EF policer configured in "Configuring a Policer for a Firewall Filter" on page 429.	Match condition: destination address 192.168.66.77 Forwarding class: ef-class Policer: ef-policer
network control	Detects packets with a network control precedence and forwards them to the network control class.	Match condition: precedence net-control Forwarding class: nc-class
best-effort-data	Detects all other packets and assigns them to the best effort class.	Forwarding class: be-class

For more information about firewalls filters see "Configuring Firewall Filters and NAT" on page 389 and the *JUNOS Policy Framework Configuration Guide* .

To configure a firewall filter for a multifield classifier for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 166.
3. Go on to "Assigning Forwarding Classes to Output Queues" on page 434.

Table 166: Configuring and Applying a Firewall Filter for a Multifield Classifier

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Firewall** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Firewall**.	From the top of the configuration hierarchy, enter edit firewall
Create and name the multifield classifier filter.	1. Click **Add new entry** next to Filter. 2. In the Filter name box, type a name for the multifield classifier filter—for example, mf-classifier. 3. Select the check box next to Interface specific.	Enter edit filter mf-classifier set interface-specific
Create and name the term for the assured forwarding traffic class.	1. Click **Add new entry** next to Term. 2. In the Rule name box, type a name for the assured forwarding term—for example, assured-forwarding.	Enter edit term assured-forwarding

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the match condition for the assured forwarding traffic class.	1. Click **Configure** next to From. 2. Click **Add new entry**next to Destination address. 3. In the Address box, type the destination address for assured forwarding traffic in dotted decimal notation—for example, 192.168.44.55. 4. Click **OK** three times.	Enter set from destination-address 192.168.44.55
Create the priority for the assured forwarding traffic class.	1. Click **Configure** next to Then. 2. In the Forwarding class box, type the forwarding class for assured forwarding DiffServ traffic—for example, af-class. 3. From the Loss priority list, select **low**. 4. Click **OK** twice.	From the top of the configuration hierarchy, enter edit firewall filter mf-classifier term assured-forwarding set then forwarding-class af-class set then loss-priority low
Create and name the term for the expedited forwarding traffic class.	1. Click **Add new entry** next to Term. 2. In the Rule name box, type a name for the expedited term—for example, expedited-forwarding.	Enter edit term expedited-forwarding
Create the match condition for the assured forwarding traffic class.	1. Click **Configure** next to From. 2. Click **Add new entry** next to Destination address. 3. In the Address box, type the destination address for assured forwarding traffic in dotted decimal notation—for example, 192.168.66.77. 4. Click **OK** twice.	Enter set from destination-address 192.168.66.77
Create the priority and apply the policer for the expedited forwarding traffic class.	1. Click **Configure** next to Then. 2. In the Forwarding class box, type the forwarding class for expedited forwarding DiffServ traffic—for example, ef-class. 3. In the Policer box, type the name of the EF policer previously configured for expedited forwarding DiffServ traffic—ef-policer. (See "Configuring a Policer for a Firewall Filter" on page 429.) 4. Click **OK** twice.	From the top of the configuration hierarchy, enter edit firewall filter mf-classifier term expedited-forwarding set then forwarding-class ef-class set then policer ed-policer

Task	J-Web Configuration Editor	CLI Configuration Editor
Create and name the term for the network control traffic class.	1. Click **Add new entry** next to Term. 2. In the Rule name box, type a name for the network control term—for example, network-control.	Enter edit term network-control
Create the match condition for the network control traffic class.	1. Click **Configure** next to From. 2. From the Precedence choice list, select **Precedence**. 3. Click **Add new entry** next to Precedence. 4. From the Value keyword list, select **net-control**. 5. Click **OK** twice.	Enter set from traffic-class net-control
Create the forwarding class for the network control traffic class.	1. Click **Configure** next to Then. 2. In the Forwarding class box, type the forwarding class for network control traffic—for example, nc-class. 3. Click **OK** twice.	From the top of the configuration hierarchy, enter edit firewall filter mf-classifier term network-control set then forwarding-class nc-class
Create and name the term for the best-effort traffic class.	1. Click **Add new entry** next to Term. 2. In the Rule name box, type a name for the best-effort term—for example, best-effort-data.	Enter edit term best-effort-data
Create the forwarding class for the best-effort traffic class. (Because this is the last term in the filter, it has no match condition.)	1. Click **Configure** next to Then. 2. In the Forwarding class box, type the forwarding class for best effort traffic—for example, be-class. 3. Click **OK** four times.	From the top of the configuration hierarchy, enter set then forwarding-class be-class
Navigate to the **Interfaces** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Interfaces**.	From the top of the configuration hierarchy, enter edit interfaces
Apply the multifield classifier firewall filter as an input filter on the customer-facing or host-facing interfaces.	1. Click the Interface and Unit of each interface needing the filter—for example, **fe-0/0/0**, unit **0**. 2. Click **Configure** next to Inet. 3. Click **Configure** next to Filter. 4. In the Input box, type the name of the previously configured filter—for example, mf-classifier. 5. Click **OK** five times.	Enter set interfaces fe-0/0/0 unit 0 family inet filter input mf-classifier

Assigning Forwarding Classes to Output Queues

You must assign the forwarding classes established by the mf-classifier multifield classifier to output queues. This example assigns output queues as shown in Table 167.

Table 167: Sample Output Queue Assignments for mf-classifier Forwarding Queues

mf-classifier Forwarding Class	For Traffic Type	Output Queue
be-class	Best-effort traffic	Queue 0
ef-class	Expedited forwarding traffic	Queue 1
af-class	Assured forwarding traffic	Queue 2
nc-class	Network control traffic	Queue 3

For multifield classifier details, see "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 430.

To assign forwarding classes to output queues for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 168.
3. Go on to "Configuring and Applying Rewrite Rules" on page 435.

Table 168: Assigning Forwarding Classes to Output Queues

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class-of-service**.	From the top of the configuration hierarchy, enter edit class-of-service
Assign best-effort traffic to queue 0.	1. Click **Configure** next to Forwarding classes. 2. Click **Add new entry** next to Queue. 3. In the Queue num box, type 0. 4. In the Class name box, type the previously configured name of the best-effort class—be-class. 5. Click **OK**.	Enter set forwarding-classes queue 0 be-class

Task	J-Web Configuration Editor	CLI Configuration Editor
Assign expedited forwarding traffic to queue 1.	1. Click **Add new entry** next to Queue. 2. In the Queue num box, type 1. 3. In the Class name box, type the previously configured name of the expedited forwarding class—ef-class. 4. Click **OK**.	Enter set forwarding-classes queue 1 ef-class
Assign assured forwarding traffic to queue 2.	1. Click **Add new entry** next to Queue. 2. In the Queue num box, type 2. 3. In the Class name box, type the previously configured name of the assured forwarding class—af-class. 4. Click **OK**.	Enter set forwarding-classes queue 2 af-class
Assign network control traffic to queue 3.	1. Click **Add new entry** next to Queue. 2. In the Queue num box, type 3. 3. In the Class name box, type the previously configured name of the expedited forwarding class—nc-class. 4. Click **OK** twice.	Enter set forwarding-classes queue 3 nc-class

Configuring and Applying Rewrite Rules

You optionally configure rewrite rules to replace DiffServ code points (DSCPs) on packets received from the customer or host with the values expected by other routers. You do not have to configure rewrite rules if the received packets already contain valid DSCPs. Rewrite rules apply the forwarding class information and packet loss priority used internally by the Services Router to establish the DSCP on outbound packets. Once configured, you must apply the rewrite rules to the correct interfaces.

The following example shows how to create the rewrite rules rewrite-dscps, and apply them to the Services Router's Fast Ethernet interface fe-0/0/0. The rewrite rules replace the DSCPs on packets in the four forwarding classes, as shown in Table 169.

Table 169: Sample rewrite-dscps Rewrite Rules to Replace DSCPs

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
be-class	Best-effort traffic	Low-priority code point: 000000 High-priority code point: 000001

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
ef-class	Expedited forwarding traffic	Low-priority code point: 101110 High-priority code point: 101111
af-class	Assured forwarding traffic	Low-priority code point: 001010 High-priority code point: 001100
nc-class	Network control traffic	Low-priority code point: 110000 High-priority code point: 110001

To configure and apply rewrite rules for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 170.
3. If you are finished configuring the network, commit the configuration.
4. Go on to "Configuring and Applying Behavior Aggregate Classifiers" on page 440.

Table 170: Configuring and Applying Rewrite Rules

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class of service**.	From the top of the configuration hierarchy, enter edit class-of-service
Configure rewrite rules for DiffServ CoS.	1. Click **Configure** next to Rewrite rules. 2. Click **Add new entry** next to Dscp. 3. In the Name box, type the name of the rewrite rules—for example, rewrite-dscps.	Enter edit rewrite-rules dscp rewrite-dscps

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure best-effort forwarding class rewrite rules.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured best-effort forwarding class—be-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **low**. 5. In the Code point box, type the value of the low-priority code point for best-effort traffic—for example, 000000. 6. Click **OK**. 7. Click **Add new entry** next to Loss priority. 8. From the Loss val list, select **high**. 9. In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, 000001. 10. Click **OK** twice.	Enter set forwarding-class be-class loss-priority low code points 000000 set forwarding-class be-class loss-priority high code points 000001

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure expedited forwarding class rewrite rules.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured expedited forwarding class—ef-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **low**. 5. In the Code point box, type the value of the low-priority code point for expedited forwarding traffic—for example, 101110. 6. Click **OK**. 7. Click **Add new entry** next to Loss priority. 8. From the Loss val list, select **high**. 9. In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, 101111. 10. Click **OK** twice.	Enter set forwarding-class ef-class loss-priority low code points 101110 set forwarding-class ef-class loss-priority high code points 101111

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure assured forwarding class rewrite rules.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured assured forwarding class—af-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **low**. 5. In the Code point box, type the value of the low-priority code point for assured forwarding traffic—for example, 001010. 6. Click **OK**. 7. Click **Add new entry** next to Loss priority. 8. From the Loss val list, select **high**. 9. In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, 001100. 10. Click **OK** twice.	Enter set forwarding-class af-class loss-priority low code points 001010 set forwarding-class af-class loss-priority high code points 001100

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure network control class rewrite rules.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured network control forwarding class—nc-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **low**. 5. In the Code point box, type the value of the low-priority code point for network control traffic—for example, 110000. 6. Click **OK**. 7. Click **Add new entry** next to Loss priority. 8. From the Loss val list, select **high**. 9. In the Code point box, type the value of the high-priority code point for network control traffic—for example, 110001. 10. Click **OK** twice.	Enter set forwarding-class nc-class loss-priority low code points 110000 set forwarding-class nc-class loss-priority high code points 110001
Apply rewrite rules to an interface.	1. Click **Add new entry** next to Interfaces 2. In the Interface name box, type the name of the interface—for example, fe-0/0/0. 3. In the Rewrite rules box, type the name of the previously configured rewrite rules—rewrite-dscps. 4. Click **OK**.	Enter set interfaces fe-0/0/0 unit 0 rewrite-rules rewrite-dscps

Configuring and Applying Behavior Aggregate Classifiers

You configure DiffServ behavior aggregate (BA) classifiers to classify packets that contain valid DSCPs to appropriate queues. Once configured, you must apply the BA classifier to the correct interfaces.

The following example shows how to configure the DSCP BA classifier ba-classifier as the default DSCP map, and apply it to the Services Router's Fast Ethernet interface fe-0/0/0. The BA classifier assigns loss priorities, as shown in Table 171, to incoming packets in the four forwarding classes.

Table 171: Sample ba-classifier Loss Priority Assignments

mf-classifier Forwarding Class	For CoS Traffic Type	ba-classifier Assignments
be-class	Best-effort traffic	High-priority code point: 000001
ef-class	Expedited forwarding traffic	High-priority code point: 101111
af-class	Assured forwarding traffic	High-priority code point: 001100
nc-class	Network control traffic	High-priority code point: 110001

To configure and apply BA classifiers for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 172.
3. If you are finished configuring the network, commit the configuration.
4. Go on to "Configuring RED Drop Profiles for Assured Forwarding Congestion Control" on page 443.

Table 172: Configuring and Applying Behavior Aggregate Classifiers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class of service**.	From the top of the configuration hierarchy, enter edit class-of-service
Configure BA classifiers for DiffServ CoS.	1. Click **Configure** next to Classifiers. 2. Click **Add new entry** next to Dscp. 3. In the Name box, type the name of the BA classifier—for example, ba-classifier. 4. In the Import box, type the name of the default DSCP map, default.	Enter edit classifiers dscp ba-classifier set import default

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a best-effort forwarding class classifier.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured best-effort forwarding class—be-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **high**. 5. In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, 00001. 6. Click **OK** three times.	Enter set forwarding-class be-class loss-priority high code points 000001
Configure an expedited forwarding class classifier.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured expedited forwarding class—ef-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **high**. 5. In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, 101111. 6. Click **OK** three times.	Enter set forwarding-class ef-class loss-priority high code points 101111
Configure an assured forwarding class classifier.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured assured forwarding class—af-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **high**. 5. In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, 001100. 6. Click **OK** three times.	Enter set forwarding-class af-class loss-priority high code points 001100

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a network control class classifier.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured network control forwarding class—nc-class. 3. Click **Add new entry** next to Loss priority. 4. From the Loss val list, select **high**. 5. In the Code point box, type the value of the high-priority code point for network control traffic—for example, 110001. 6. Click **OK** three times.	Enter set forwarding-class nc-class loss-priority high code points 110001
Apply the BA classifier to an interface.	1. Click **Add new entry** next to Interfaces. 2. In the Interface name box, type the name of the interface—for example, fe-0/0/0. 3. In the Classifiers box, type the name of the previously configured BA classifier—ba-classifier. 4. Click **OK**.	Enter set interfaces fe-0/0/0 unit 0 classifiers dscp ba-classifier

Configuring RED Drop Profiles for Assured Forwarding Congestion Control

If the Services Router must support DiffServ assured forwarding (AF), you can control congestion by configuring random early detection (RED) drop profiles. RED drop profiles use drop probabilities for different levels of buffer fullness to determine which scheduling queue on the router is likely to drop DiffServ assured forwarding (AF) packets under congested conditions. The router can drop packets when the queue buffer becomes filled to the configured percentage.

Assured forwarding traffic with the PLP (packet loss priority) bit set is more likely to be discarded than traffic without the PLP bit set. This example shows how to configure a drop probability and a queue fill level for both PLP and non-PLP assured forwarding traffic. It is only one example of how to use RED drop profiles.

The example shows how to configure the RED drop profiles listed in Table 173.

Table 173: Sample RED Drop Profiles

Drop Profile	Drop Probability	Queue Fill Level
af-normal—For non-PLP (normal) assured forwarding traffic	Between 0 (never dropped) and 100 percent (always dropped)	Between 95 and 100 percent
af-with-plp—For PLP (aggressive packet dropping) assured forwarding traffic	Between 95 and 100 percent (always dropped)	Between 80 and 95 percent

To configure RED drop profiles for assured forwarding congestion control on the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 174.
3. Go on to one of the following tasks:
   - "Configuring Schedulers" on page 446
   - "Verifying a DiffServ Configuration" on page 457

Table 174: Configuring RED Drop Profiles for Assured Forwarding Congestion Control

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class of service**.	From the top of the configuration hierarchy, enter edit class-of-service

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the lower drop probability for normal, non-PLP traffic.	1. Click **Add new entry** next to Drop profiles. 2. In the Profile name box, type the name of the drop profile—for example, af-normal. 3. Click **Configure** next to Interpolate. 4. Click **Add new entry** next to Drop probability. 5. In the Value box, type a number for the first drop point—for example, 0. 6. Click **OK**. 7. Click **Add new entry** next to Drop probability again. 8. In the Value box, type a number for the next drop point—for example, 100. 9. Click **OK**.	Enter edit drop-profiles af-normal interpolate set drop-probability 0 set drop-probability 100
Configure a queue fill level for the lower non-PLP drop probability.	1. Click **Add new entry** next to Fill level. 2. In the Value box, type a number for the first fill level—for example, 95. 3. Click **OK**. 4. In the Value box, type a number for the next fill level—for example, 100. 5. Click **OK** three times.	Enter set fill-level 95 set fill-level 100

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the higher drop probability for PLP traffic.	1. Click **Add new entry** next to Drop profiles. 2. In the Profile name box, type the name of the drop profile—for example, af-with-plp. 3. Click **Configure** next to Interpolate. 4. Click **Add new entry** next to Drop probability. 5. In the Value box, type a number for the first drop point—for example, 95. 6. Click **OK**. 7. In the Value box, type a number for the next drop point—for example, 100. 8. Click **OK**.	Enter edit drop-profiles af-with-PLP interpolate set drop-probability 95 set drop-probability 100
Configure a queue fill level for the higher PLP drop probability.	1. Click **Add new entry** next to Fill level. 2. In the Value box, type a number for the first fill level—for example, 80. 3. Click **OK**. 4. In the Value box, type a number for the next fill level—for example, 95. 5. Click **OK**.	Enter set fill-level 80 set fill-level 95

Configuring Schedulers

You configure schedulers to assign resources, priorities, and drop profiles to output queues. By default, only queues 0 and 4 have resources assigned.

This example creates the schedulers listed in Table 175.

Table 175: Sample Schedulers

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer	Assigned Bandwidth (Transmit Rate)
be-scheduler	Best-effort traffic	Low	40 percent	10 percent
ef-scheduler	Expedited forwarding traffic	High	10 percent	10 percent

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer	Assigned Bandwidth (Transmit Rate)
af-scheduler	Assured forwarding traffic	High	45 percent	45 percent
nc-scheduler	Network control traffic	Low	5 percent	5 percent

To configure schedulers for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 176.
3. Go on to "Configuring and Applying Scheduler Maps" on page 450.

Table 176: Configuring Schedulers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class-of-service**.	From the top of the configuration hierarchy, enter edit class-of-service
Configure a best-effort scheduler.	1. Click **Add new entry** next to Schedulers. 2. In the Scheduler name box, type the name of the best-effort scheduler—for example, be-scheduler.	Enter edit schedulers be-scheduler
Configure a best-effort scheduler priority and buffer size.	1. In the Priority box, type low. 2. Click **Configure** next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, **percent**. 4. In the Percent box, type the percentage of the buffer to be used by the best-effort scheduler—for example, 40. 5. Click **OK**.	Enter set priority low set buffer-size percent 40

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a best-effort scheduler transmit rate.	1. Click **Configure** next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, **percent**. 3. In the Percent box, type the percentage of the bandwidth to be used by the best-effort scheduler—for example, 10. 4. Click **OK** twice.	Enter set transmit-rate percent 10
Configure an expedited forwarding scheduler.	1. Click **Add new entry** next to Schedulers. 2. In the Scheduler name box, type the name of the expedited forwarding scheduler—for example, ef-scheduler.	Enter edit schedulers ef-scheduler
Configure an expedited forwarding scheduler priority and buffer size.	1. In the Priority box, type high. 2. Click **Configure** next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, **percent**. 4. In the Percent box, type the percentage of the buffer to be used by the expedited forwarding scheduler—for example, 10. 5. Click **OK**.	Enter set priority high set buffer-size percent 10
Configure an expedited forwarding scheduler transmit rate.	1. Click **Configure** next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, **percent**. 3. In the Percent box, type the percentage of the bandwidth to be used by the expedited forwarding scheduler—for example, 10. 4. Click **OK** twice.	Enter set transmit-rate percent 10
Configure an assured forwarding scheduler.	1. Click **Add new entry** next to Schedulers. 2. In the Scheduler name box, type the name of the assured forwarding scheduler—for example, af-scheduler.	Enter edit schedulers af-scheduler

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure an assured forwarding scheduler priority and buffer size.	1. In the Priority box, type high. 2. Click **Configure** next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, **percent**. 4. In the Percent box, type the percentage of the buffer to be used by the assured forwarding scheduler—for example, 45. 5. Click **OK**.	Enter set priority high set buffer-size percent 45
Configure an assured forwarding scheduler transmit rate.	1. Click **Configure** next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, **percent**. 3. In the Percent box, type the percentage of the bandwidth to be used by the assured forwarding scheduler—for example, 45. 4. Click **OK**.	Enter set transmit-rate percent 45
(Optional) Configure a drop profile map for assured forwarding low and high priority. (DiffServ can have a RED drop profile associated with assured forwarding.)	1. Click **Add new entry** next to Drop profile map. 2. From the Loss priority box, select **Low**. 3. From the Protocol box, select **Any**. 4. In the Drop profile box, type the name of the drop profile—for example, af-normal. 5. Click **OK**. 6. Click **Add new entry** next to Drop profile map. 7. From the Loss priority box, select **High**. 8. From the Protocol box, select **Any**. 9. In the Drop profile box, type the name of the drop profile—for example, af-with-PLP. 10. Click **OK**.	Enter set drop-profile-map loss-priority low protocol any drop-profile af-normal set drop-profile-map loss-priority high protocol any drop-profile af-with-PLP

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a network control scheduler.	1. Click **Add new entry** next to Schedulers. 2. In the Scheduler name box, type the name of the network control scheduler—for example, nc-scheduler.	Enter edit schedulers nc-scheduler
Configure a network control scheduler priority and buffer size.	1. In the Priority box, type low. 2. Click **Configure** next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, **percent**. 4. In the Percent box, type the percentage of the buffer to be used by the network control scheduler—for example, 5. 5. Click **OK**.	Enter set priority low set buffer-size percent 5
Configure a network control scheduler transmit rate.	1. Click **Configure** next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, **percent**. 3. In the Percent box, type the percentage of the bandwidth to be used by the network control scheduler—for example, 5. 4. Click **OK** twice.	Enter set transmit-rate percent 5

Configuring and Applying Scheduler Maps

You configure a scheduler map to assign a forwarding class to a scheduler, then apply the scheduler map to any interface that must enforce DiffServ CoS.

The following example shows how to create the scheduler map diffserv-cos-map and apply it to the Services Router's Fast Ethernet interface fe-0/0/0. The map associates the mf-classifier forwarding classes configured in "Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 430 to the schedulers configured in "Configuring Schedulers" on page 446, as shown in Table 177.

Table 177: Sample diffserv-cos-map Scheduler Mapping

mf-classifier Forwarding Class	For CoS Traffic Type	diffserv-cos-map Scheduler
be-class	Best-effort traffic	be-scheduler
ef-class	Expedited forwarding traffic	ef-scheduler
af-class	Assured forwarding traffic	af-scheduler
nc-class	Network control traffic	nc-scheduler

To configure and apply scheduler maps for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 178.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see "Verifying a DiffServ Configuration" on page 457.

Table 178: Configuring Scheduler Maps

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class of service**.	From the top of the configuration hierarchy, enter edit class-of-service
Configure a scheduler map for DiffServ CoS.	1. Click **Add new entry** next to Scheduler maps. 2. In the Map name box, type the name of the scheduler map—for example, diffserv-cos-map.	Enter edit scheduler-maps diffserv-cos-map
Configure a best-effort forwarding class and scheduler.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured best-effort forwarding class—be-class. 3. In the Scheduler box, type the name of the previously configured best-effort scheduler—be-scheduler. 4. Click **OK**.	Enter set forwarding-class be-class scheduler be-scheduler

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure an expedited forwarding class and scheduler.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured expedited forwarding class—ef-class. 3. In the Scheduler box, type the name of the previously configured expedited forwarding scheduler—ef-scheduler. 4. Click **OK**.	Enter set forwarding-class ef-class scheduler ef-scheduler
Configure an assured forwarding class and scheduler.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured assured forwarding class—af-class. 3. In the Scheduler box, type the name of the previously configured assured forwarding scheduler—af-scheduler. 4. Click **OK**.	Enter set forwarding-class af-class scheduler af-scheduler
Configure a network control class and scheduler.	1. Click **Add new entry** next to Forwarding class. 2. In the Class name box, type the name of the previously configured network control class—nc-class. 3. In the Scheduler box, type the name of the previously configured network control scheduler—nc-scheduler. 4. Click **OK** twice.	Enter set forwarding-class nc-class scheduler nc-scheduler
Apply the scheduler map to an interface.	1. Click **Add new entry** next to Interfaces. 2. In the Interface name box, type the name of the interface—for example, fe-0/0/0. 3. In the Scheduler map box, type the name of the previously configured scheduler map—diffserv-cos-map. 4. Click **OK**.	Enter set interfaces fe-0/0/0 scheduler-map diffserv-cos-map

Configuring and Applying Virtual Channels

You configure a virtual channel to set up queuing, packet scheduling, and accounting rules to be applied to one or more logical interfaces. You then must apply the virtual channel to a particular logical interface.

The following example shows how to create the virtual channels branch1–vc, branch2–vc, and branch3–vc and apply them in the firewall filter choose-vc to the Services Router's T3 interface t3-1/0/0.

To configure and apply virtual channels for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 179.
3. If you are finished configuring the network, commit the configuration.

Table 179: Configuring and Applying Virtual Channels

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Class-of-service** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Class of service**.	From the top of the configuration hierarchy, enter edit class-of-service
Define the virtual channels branch1–vc, branch2–vc, branch3–vc, and the default virtual channel. You must specify a default virtual channel.	1. Click **Add new entry** next to Virtual channels. 2. In the Channel name box, type the name of the virtual channel—for example, branch1–vc. 3. Click **OK**. 4. Create additional virtual channels for branch2–vc, branch3–vc, and default-vc.	Enter set virtual-channels branch1–vc Repeat this statement for branch2–vc, branch3–vc, and default-vc.

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the virtual channel group wan-vc-group to include the four virtual channels, and assign each virtual channel the scheduler map bestscheduler.	1. Click **Add new entry** next to Virtual channel groups. 2. In the Group name box, type the name of the virtual channel group—wan-vc-group. 3. Click **Add new entry** next to Channel. 4. In the Channel name box, enter the name of the previously configured virtual channels—branch1–vc. 5. In the Scheduler map box, enter the name of the previously configured scheduler map—bestscheduler. 6. Click **OK**. 7. Add the virtual channels branch2–vc, branch3–vc, and default-vc. Select the **Default** box when adding the virtual channel default-vc.	1. Enter set virtual-channel-groups wan-vc-group branch1–vc scheduler-map bestscheduler 2. Repeat this statement for branch2–vc, branch3–vc, and default-vc. 3. Enter set virtual-channel-groups wan-vc-group default–vc default
Specify a shaping rate of 1.5 Mbps for each virtual channel within the virtual channel group.	1. Click **branch1–vc** in the list of virtual channels. 2. Select the **Shaping rate** box. 3. Click **Configure**. 4. Select **Absolute rate** from the Rate choice box.. 5. In the Absolute rate box, enter the shaping rate—1.5m. 6. Add the shaping rate for the branch2–vc and branch3–vc virtual channels. 7. Click **OK**.	1. Enter set virtual-channel-groups wan-vc-group branch1–vc shaping-rate 1.5m 2. Repeat this statement for branch2–vc and branch3–vc.

Task	J-Web Configuration Editor	CLI Configuration Editor
Apply the virtual channel group to the logical interface t3–1/0/0.0.	1. Click **Add new entry** next to Interfaces. 2. In the Interface name box, type the name of the interface—t3–1/0/0. 3. Click **Add new entry** next to Unit. 4. In the Unit number box, type the logical interface unit number—0. 5. In the Virtual channel group box, type the name of the previously configured virtual channel group—wan-vc-group. 6. Click **OK**.	Enter set interfaces t3–1/0/0 unit 0 virtual-channel-group wan-vc-group

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the firewall filter choose-vc to select the traffic that is transmitted on a particular virtual channel.	1. Navigate to the top of the configuration hierarchy and select **Firewall**. 2. Click **Add new entry** next to Filter. 3. In the Filter name box, enter the name of the firewall filter—choose-vc. 4. Click **Add new entry** next to Term. 5. In the Rule name box, enter the name of the firewall term—branch1. 6. Click **Configure** next to From. 7. Click **Add new entry** next to Destination address. 8. In the Address box, enter the IP address of the destination host—192.168.10.0/24. 9. Click **OK** twice. 10. On the firewall term page, click **Configure** next to Then. 11. Select **Accept** from the Designation box. 12. In the Virtual channel box, enter the name of the previously configured virtual channel—branch1–vc. 13. Click **OK**. 14. Repeat these steps for the virtual channels branch2–vc and branch3–vc.	1. From the top of the configuration hierarchy, enter edit firewall 2. Enter set family inet filter choose-vc term branch1 from destination 192.168.10.0/24 3. Enter set family inet filter choose-vc term branch1 then accept 4. Enter set family inet filter choose-vc term branch1 then virtual-channel branch1–vc 5. Repeat these steps for virtual channels branch2–vc and branch3–vc.
Apply the firewall filter choose-vc to output traffic on the t3–1/0/0.0 interface.	1. Navigate to the top of the configuration hierarchy and select **Interfaces**. 2. Click **t3–1/0/0** in the list of configured interfaces. 3. Click **0** in the list of configured logical units for the interface. 4. Click **Edit** next to Inet. 5. Click **Configure** next to Filter. 6. In the Output box, enter the name of the previously configured firewall filter—choose–vc. 7. Click **OK**.	1. From the top of the configuration hierarchy, enter edit interfaces 2. Enter set t3–1/0/0 unit 0 family inet filter output choose-vc

Verifying a DiffServ Configuration

To verify a DiffServ configuration, perform the following task.

Verifying Multicast Session Announcements

Purpose Verify that the Services Router is listening to the appropriate groups for multicast Session Announcement Protocol (SAP) session announcements.

Action From the CLI, enter the show sap listen command.

Sample Output

```
user@host> show sap listen

Group Address   Port
224.2.127.254   9875
```

What It Means The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:

- Each group address configured, especially the default 224.2.127.254, is listed.
- Each port configured, especially the default 9875, is listed.

For more information about show sap listen, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Part 7

Managing Multicast Transmissions

Chapter 21
Multicast Overview

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning that the destinations needing to receive the information from a particular source receive the traffic stream.

IP network destinations (clients) do not often communicate directly with sources (servers), so the routers between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. The multicast router must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities.

This chapter contains the following topics. For more information about multicast, see the *JUNOS Multicast Protocols Configuration Guide*. For configuration instructions, see "Configuring a Multicast Network" on page 471.

Multicast Terms

To understand multicast routing, you must be familiar with the terms defined in Table 180. See Figure 95 for a general view of some of the elements commonly used in an IP multicast network architecture.

Table 180: Multicast Terms

Term	Definition
administrative scoping	Multicast routing strategy that limits the routers and interfaces used to forward a multicast packet by reserving a range of multicast addresses.
any-source multicast (ASM)	
Auto-RP	Cisco multicast routing protocol that allows sparse-mode routing protocols to find rendezvous points (RPs) within a routing domain.
Bootstrap Router (BSR) protocol	Multicast routing protocol that allows sparse-mode routing protocols to find rendezvous points (RPs) within a routing domain.
branch	Part of a multicast network that is formed when a leaf subnetwork is joined to the multicast distribution tree. Branches with no interested receivers are pruned from the tree so that multicast packets are no longer replicated on the branch.
broadcast routing protocol	Protocol that distributes traffic from a particular source to all destinations.
dense mode	Multicast routing mode appropriate for LANs with many interested receivers.
Distance Vector Multicast Routing Protocol (DVMRP)	Distributed multicast routing protocol that dynamically generates IP multicast distribution trees using reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces.
distribution tree	Path linking multicast receivers (listeners) to sources. The root of the tree is at the source, and the branches connect subnetworks of interested receivers (leaves). Multicast packets are replicated only where a distribution tree branches. To shorten paths to a source at the edge of a network, sparse mode multicast protocols can use a *shared* distribution tree located more centrally in the network backbone.
downstream interface	Interface on a multicast router that is leading toward the receivers. You can configure all the logical interfaces except one as downstream interfaces.
group address	Multicast destination address. A multicast network uses the Class D IP address of a logical group of multicast receivers to identify a destination. IP multicast packets have a multicast group address as the destination address and a unicast source address.
Internet Group Management Protocol (IGMP)	Multicast routing protocol that runs between receiver hosts and routers to determine whether group members are present. Services Routers support IGMPv1, IGMPv2, and IGMPv3.
leaf	IP subnetwork that is connected to a multicast router and that includes at least one host interested in receiving IP multicast packets. The router must send a copy of its multicast packets out on each interface with a leaf, and its action is unaffected by the number of leaves on the interface.
listener	Another name for a receiver in a multicast network.

Term	Definition
multicast routing protocol	Protocol that distributes traffic from a particular source to only the destinations needing to receive it. Typical multicast routing protocols are the Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM).
Multicast Source Discovery Protocol (MSDP)	Multicast routing protocol that connects multicast routing domains and allows them to find rendezvous points (RPs).
Pragmatic General Multicast (PGM)	Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic.
Protocol Independent Multicast (PIM) protocol	Protocol-independent multicast routing protocol that can be used in either sparse or dense mode. In sparse mode, PIM routes to multicast groups that might span WANs and interdomain Internets. In dense mode, PIM is a flood-and-prune protocol.
pruning	Removing from a multicast distribution tree branches that no longer include subnetworks with interested hosts. Pruning ensures that packets are replicated only as needed.
reverse-path forwarding (RPF)	Multicast routing strategy that allows a router to receive packets through an interface if it is the same interface a unicast packet uses as the shortest path back to the source.
rendezvous point (RP)	Core router operating as the root of a shared distribution tree in a multicast network.
Session Announcement Protocol (SAP)	Multicast routing protocol used with other multicast protocols—typically Session Description Protocol (SDP)—to handle session conference announcements.
Session Description Protocol (SDP)	Session directory protocol that advertise multimedia conference sessions and communicates setup information to participants who want to join the session.
shortest-path tree (SPT)	Multicast routing strategy for sparse mode multicast protocols. SPT uses a shared distribution tree rooted in the network backbone to shorten paths to sources at the edge of a network.
source-specific multicast (SSM)	Service that allows a client to receive multicast traffic directly from the source, without the help of a rendezvous point (RP).
sparse mode	Multicast routing mode appropriate for WANs with few interested receivers.
unicast routing protocol	Protocol that distributes traffic from one source to one destination.
upstream interface	Interface on a multicast router that is leading toward the source. To minimize bandwidth use, configure only one upstream interface on a router receiving multicast packets.

Multicast Architecture

Multicast-capable routers replicate packets on the multicast network, which has exactly the same topology as the unicast network it is based on. Multicast routers

use a multicast routing protocol to build a distribution tree that connects receivers (also called *listeners*) to sources.

Multicast architecture includes the following topics:

Upstream and Downstream Interfaces

A single upstream interface on the router leads toward the source to receive multicast packets. The downstream interfaces on the router lead toward the receivers to transmit packets. A router can have as many downstream interfaces as it has logical interfaces, minus 1. To prevent looping, the router's upstream interface must never receive copies of its own downstream multicast packets.

Subnetwork Leaves and Branches

On a multicast router, each subnetwork of hosts that includes at least one interested receiver is a leaf on the multicast distribution tree (see Figure 95). The router must send out a copy of the IP multicast packet on each interface with a leaf. When a new leaf subnetwork joins the tree, a new branch is built so that the router can send out replicated packets on the interface. The number of leaves on an interface does not affect the router. The action is the same for one leaf or a hundred.

A branch that no longer has leaves is pruned from the distribution tree. No multicast packets are sent out on a router interface leading to an IP subnetwork with no interested hosts. Because packets are replicated only where the distribution tree branches, no link ever carries a duplicate flow of packets.

In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast group address instead of a unicast destination address. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

Figure 95: Multicast Elements in an IP Network

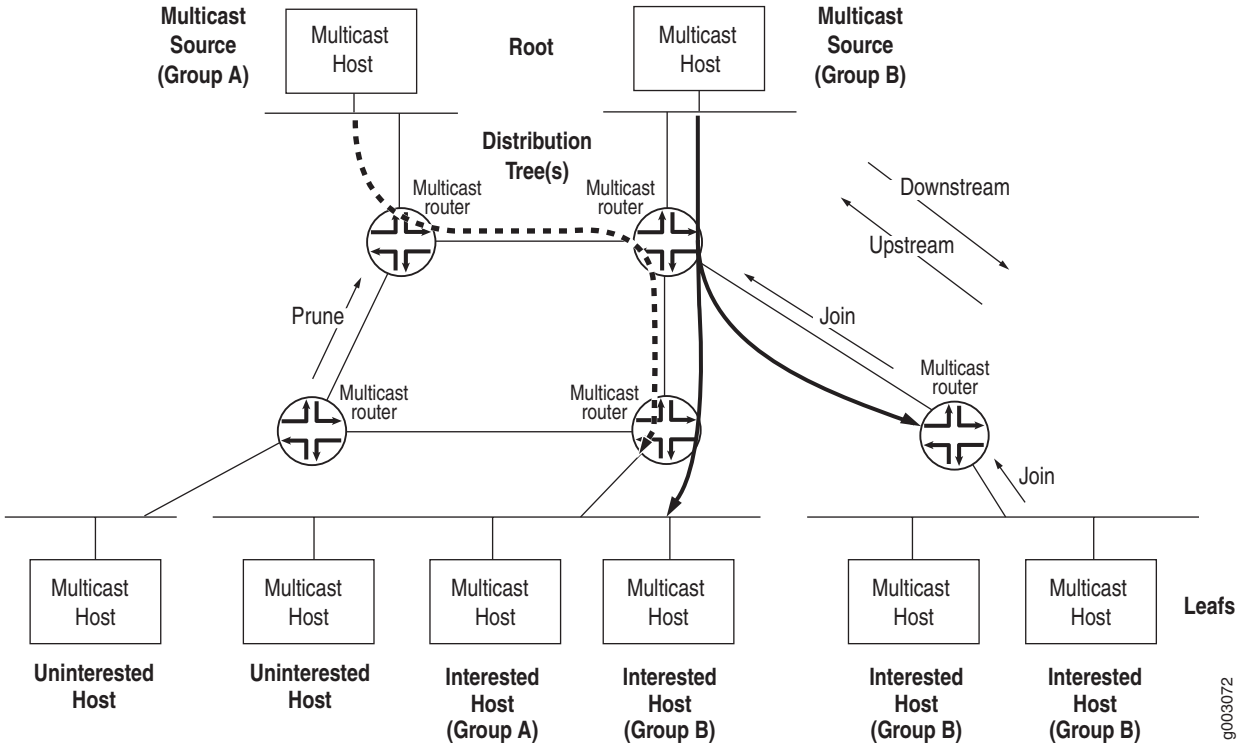

Multicast IP Address Ranges

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices, and can appear only as the destination in an IP packet, never as the source address.

Notation for Multicast Forwarding States

The multicast forwarding state in a router is usually represented by one of the following notations:

- (S,G) notation—S refers to the unicast IP address of the source for the multicast traffic and G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.

- (*, G) notation—The asterisk (*) is a wildcard for the address of any multicast application source sending to group G. For example, if two sources are originating exactly the same content for multicast group 224.1.1.2, a router can use (*, 224.1.1.2) to represent the state of a router forwarding traffic from both sources to the group.

Dense and Sparse Routing Modes

To keep packet replication to a minimum, multicast routing protocols use the two primary modes shown in Table 181.

CAUTION: A common multicast guideline is *not to run dense mode on a WAN under any circumstances*.

Table 181: Primary Multicast Routing Modes

Multicast Mode	Description	Appropriate Network for Use
Dense mode	Network is flooded with traffic on all possible branches, then pruned back as branches explicitly (by message) or implicitly (time-out silence) eliminate themselves.	LANs—Networks in which all possible subnets are likely to have at least one receiver.
Sparse mode	Network establishes and sends packets only on branches that have at least one leaf indicating (by message) a need for the traffic.	WANs—Network in which very few of the possible receivers require packets from this source.

Strategies for Preventing Routing Loops

Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets, which can overwhelm a network. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols. Three multicast strategies help prevent routing loops by defining routing paths in different ways:

- Reverse-Path Forwarding for Loop Prevention on page 466
- Shortest-Path Tree for Loop Prevention on page 467
- Administrative Scoping for Loop Prevention on page 467

Reverse-Path Forwarding for Loop Prevention

The router's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In reverse-path forwarding (RPF), every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the router verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. Routers can build and maintain separate tables for RPF purposes.

Shortest-Path Tree for Loop Prevention

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast router operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

Administrative Scoping for Loop Prevention

Scoping limits the routers and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. Routers at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

Multicast Protocol Building Blocks

Multicast is not a single protocol, but a collection of protocols working together to form trees, prune branches, locate sources and groups, and prevent routing loops:

- Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM) operate between routers. PIM can operate in dense mode and sparse mode.
- Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routers.
- Several other protocols enhance multicast networks by providing useful functions not included in other protocols. These include the Bootstrap Router (BSR) and Auto-RP protocols, Multicast Source Discovery Protocol (MSDP), Session Announcement Protocol (SAP) and Session Discovery Protocol (SDP), and Pragmatic General Multicast (PGM) protocol.

Table 182 lists and summarizes these protocols.

Table 182: Multicast Protocol Building Blocks

Multicast Protocol	Description	Uses
DVMRP	Dense-mode-only protocol that uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVRMP uses source-based distribution trees in the form (S,G) and builds its own multicast routing tables for RPF checks.	Not appropriate for large-scale Internet use.
PIM dense mode	Sends an *implicit* join message, so routers use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are. PIM dense mode uses source-based distribution trees in the form (S,G), and also supports sparse-dense mode, with mixed sparse and dense groups. Both PIM modes use unicast routing information for RPF checks.	Most promising multicast protocol in use for LANs.
PIM sparse mode	Sends an *explicit* join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to a rendezvous point (RP) router, which is the initial source of multicast group traffic. PIM sparse mode builds distribution trees in the form (*,G), but migrates to an (S,G) source-based tree if that path is shorter than the path through the RP router for a particular multicast group's traffic. Both PIM modes use unicast routing information for RPF checks.	Most promising multicast protocol in use for WANs.
PIM source-specific multicast (SSM)	Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of a rendezvous point (RP).	Used with IGMPv3 to create a shortest-path tree between receiver and source.
IGMPv1	The original protocol defined in RFC 1112, *Host Extensions for IP Multicasting*. IGMPv1 sends an explicit join message to the router, but uses a time-out to determine when hosts leave a group.	
IGMPv2	Defined in RFC 2236, *Internet Group Management Protocol, Version 2*. Among other features, IGMPv2 adds an explicit leave message to the join message.	Used by default.

Multicast Protocol	Description	Uses
IGMPv3	Defined in RFC 3376, *Internet Group Management Protocol, Version 3*. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM).*	Used with PIM SSM to create a shortest-path tree between receiver and source.
BSR Auto-RP	Allow sparse-mode routing protocols to find rendezvous points (RPs) within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.	
MSDP	Allows groups located in one multicast routing domain to find rendezvous points (RPs) in other routing domains. MSDP is not used on an RP if all receivers and sources are located in the same routing domain.	Typically runs on the same router as PIM sparse mode rendezvous point (RP). Not appropriate if all receivers and sources are located in the same routing domain.
SAP and SDP	Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.	
PGM	Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.	

Chapter 22
Configuring a Multicast Network

You configure a router network to support multicast applications with a related family of protocols. To use multicast, you must understand the basic components of a multicast network and their relationships, and then configure the J-series Services Router to act as a node in the network.

NOTE: The J-series Services Router supports both PIM version 1 and PIM version 2. In this chapter, the term *PIM* refers to both versions of the protocol.

You use either the J-Web configuration editor or CLI configuration editor to configure multicast protocols. The J-Web interface does not include Quick Configuration pages for multicast protocols.

This chapter contains the following topics. For more information about multicast, see the *JUNOS Multicast Protocols Configuration Guide*.

Before You Begin

Before you begin configuring a multicast network, complete the following tasks:

- If you do not already have a basic understanding of multicast, read “Multicast Overview” on page 461.
- Determine whether the Services Router is directly attached to any multicast sources. Receivers must be able to locate these sources.
- Determine whether the Services Router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
- Determine whether to use the sparse, dense, or sparse-dense mode of multicast operation. Each mode has different configuration considerations.
- Determine the address of the rendezvous point (RP) if sparse or sparse-dense mode is used.
- Determine whether to locate the RP with the static configuration, bootstrap router (BSR), or Auto-RP method.
- Determine whether to configure multicast to use its own reverse-path forwarding (RPF) routing table when configuring PIM in sparse, dense, or sparse-dense modes.

Configuring a Multicast Network with a Configuration Editor

To configure the Services Router as a node in a multicast network, you must perform the following tasks marked *(Required)*.

- (Optional) “Configuring SAP and SDP” on page 472
- (Required) “Configuring IGMP” on page 473
- (Optional) “Configuring the PIM Static RP” on page 474
- (Optional) “Configuring a PIM RPF Routing Table” on page 476

For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 127.

Configuring SAP and SDP

Multicast session announcements are handled by two protocols, the Session Announcement Protocol (SAP) and Session Description Protocol (SDP). These two protocols display multicast session names and correlate the names with multicast traffic. Enabling SDP and SAP allows the router to receive announcements about multimedia and other multicast sessions from sources. Enabling SAP automatically enables SDP.

For more information on SAP and SDP, see the *JUNOS Multicast Protocols Configuration Guide*.

The Services Router listens for session announcements on one or more addresses and ports. By default, the router listens to address and port 224.2.127.254:9875.

To configure SAP and SDP for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 183.
3. Go on to "Configuring IGMP" on page 473.

Table 183: Configuring SAP and SDP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Listen** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Protocols > Sap**. 2. Click **Add new entry** next to Listen.	From the top of the configuration hierarchy, enter edit protocols sap
(Optional) Enter one or more addresses and ports for the Services Router to listen to session announcements on. By default, the Services Router listens to address and port 224.2.127.254:9875.	1. In the Address box, type the multicast address the Services Router can listen to session announcements on, in dotted decimal notation. 2. In the Port box, type the port number in decimal notation. 3. Click **OK**.	1. Set the address value to the IP address that the Services Router can listen to session announcements on, in dotted decimal notation. For example: set listen 224.2.127.254 2. Set the port value to the number of the port that the Services Router can listen to session announcements on, in decimal notation. For example: set listen 224.2.127.254 port 9875.

Configuring IGMP

The Internet Group Membership Protocol (IGMP) manages the membership of hosts and routers in multicast groups. IGMP is an integral part of IP and must be enabled on all routers and hosts that need to receive IP mulitcasts. IGMP is automatically enabled on all broadcast interfaces when you configure PIM or DVMRP.

For more information on IGMP, see *JUNOS Multicast Protocols Configuration Guide*.

By default, the Services Router runs IGMPv2. However, you might still want to set the IGMP version explicitly on an interface, or all interfaces. Routers running different versions of IGMP negotiate the lowest common version of IGMP supported

by hosts on their subnet. One host running IGMPv1 forces the Services Router to use that version and lose features important to other hosts.

To explicitly configure the IGMP version, perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 184.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
   - To configure PIM sparse mode, see "Configuring the PIM Static RP" on page 474 and Table 186.
   - To check the configuration, see "Verifying a Multicast Configuration" on page 478.

Table 184: Explicitly Configuring the IGMP version

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Interface** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Protocols > Igmp**. 2. Click **Add new entry** next to Interface.	From the top of the configuration hierarchy, enter edit protocols igmp
Set the IGMP version. By default, the Services Router uses IGMPv2, but this version can be changed through negotiation with hosts unless explicitly configured.	1. In the Interface name box, type the name of the interface, or all. 2. In the Version box, type the version number: 1, 2, or 3. 3. Click **OK**.	1. Set the interface value to the interface name, or all. For example: set igmp interface all 2. Set the version value to 1, 2, or 3. For example: set igmp interface all version 2

Configuring the PIM Static RP

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the Services Router. However, because PIM must not be configured on the network management interface of the Services Router, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about

all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the Services Router must determine the IP address of the RP for the source.

One common way for the Services Router to locate RPs is by static configuration of the IP address of the RP. For information about alternate methods of locating RPs, see the *JUNOS Multicast Protocols Configuration Guide*.

To configure PIM sparse mode, disable PIM on fe-0/0/0, and configure the IP address of the RP perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 185.
3. Go on to "Configuring a PIM RPF Routing Table" on page 476.

Table 185: Configuring PIM Sparse Mode and the RP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Interface** level in the configuration hierarchy.	1. In the configuration editor hierarchy, select **Protocols > Pim**. 2. Click **Add new entry** next to Interface.	From the top of the configuration hierarchy, enter edit protocols pim
Enable PIM on all network interfaces.	In the Interface name box, type all.	Set the interface value to all. For example: set pim interface all
Apply your configuration changes.	Click **OK** to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the set command.
Remain at the **Interface** level in the configuration hierarchy.	Click **Add new entry** next to Interface.	Remain at the edit protocols pim interface configuration hierarchy level.
Disable PIM on the network management interface.	1. In the Interface name box, type fe-0/0/0. 2. Select the check box next to Disable.	Disable the fe-0/0/0 interface: set pim interface fe—0/0/0 unit 0 disable
Apply your configuration changes.	Click **OK** to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the set command.

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Rp** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Pim > Rp**.	From the top of the configuration hierarchy, enter edit protocols pim rp
Configure the IP address of the RP.	1. Click **Configure** next to Static. 2. Click **Add new entry** next to Address. 3. In the Addr box, type the IP address of the RP in dotted decimal notation. 4. Click **OK**.	Set the address value to the IP address of the RP in dotted decimal notation. For example: set static address 192.168.14.27

Configuring a PIM RPF Routing Table

By default, PIM uses inet.0 as its reverse-path forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and for the RP address. PIM can optionally use inet.2 as its RPF routing table group. The inet.2 routing table is organized more efficiently for RPF checks.

Once configured, the RPF routing table must be applied to PIM as a routing table group.

To configure and apply a PIM RPF routing table, perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 186.
3. To check the configuration, see “Verifying a Multicast Configuration” on page 478.

Table 186: Configuring a PIM RPF Routing Table

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Routing options** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Routing options**.	From the top of the configuration hierarchy, enter edit routing-options
Configure a new group for the RPF routing table.	Next to Rib groups, click **Add new entry**.	Enter edit rib-groups

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a name for the RPF routing table group, and use inet.2 for its export routing table.	1. In the Ribgroup name box, type a name for the RPF routing table group—for example, multicast-rfp-rib. 2. In the Export rib box, type inet.2.	Type the name for the RPF routing table and set the export routing table to inet.2. For example: set multicast-rpf-rib export-rib inet.2
Configure an import routing table routing information base (RIB) group for the RPF routing table.	1. Click **Add new entry** next to Import rib. 2. In the Value box, type inet.2. 3. Click **OK** three times.	Set the import routing table to inet.2. For example: set multicast-rpf-rib import-rib inet.2
Navigate to the **Rib group** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Protocols > Pim > Rib group**.	From the top of the configuration hierarchy, enter edit protocols pim
Apply the RPF routing table to PIM.	1. In the Inet box, type the name of the RPF routing table group—for example, multicast-rpf-rib. 2. Click **OK** three times.	Enter set rib-group multicast-rpf-rib
Create a RIB group for the interface routes.	1. Navigate to the **Routing options** level in the configuration hierarchy. 2. Next to Rib groups, click **Add new entry**.	From the top of the configuration hierarchy, enter edit routing-options rib-groups.
Configure a name for the RPF routing table group, and use inet.2 and inet.0 for its import routing table.	1. In the Ribgroup name box, type a name for the RPF routing table group—for example, if-rib. 2. Click **Add new entry** next to Import rib. 3. In the Value box, type inet.2 inet.0. 4. Click **OK** twice.	Type the name for the RPF routing table and set the export routing table to inet.2 and inet.0. For example: set if-rib import-rib inet.2 set if-rib import-rib inet.0
Add the RIB group to the interface routes.	1. On the **Routing options** page, select **Interface routes > Rib group**. 2. In the Inet box, type the name of the interface RIB group—for example, if-rib. 3. Click **OK**.	From the top of the configuration hierarchy, enter edit routing-options interface-routes set rib-group inet if-rib

Verifying a Multicast Configuration

To verify a multicast configuration, perform these tasks:

Verifying SAP and SDP Addresses and Ports

Purpose Verify that SAP and SDP are configured to listen on the correct group addresses and ports.

Action From the CLI, enter the show sap listen command.

Sample Output

```
user@host> show sap listen

Group Address   Port
224.2.127.254   9875
```

What It Means The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:

- Each group address configured, especially the default 224.2.127.254, is listed.
- Each port configured, especially the default 9875, is listed.

For more information about show sap listen, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying the IGMP Version

Purpose Verify that IGMP version 2 is configured on all applicable interfaces.

Action From the CLI, enter the show igmp interface command.

Sample Output

```
user@host> show igmp interface

Interface: fe-0/0/0.0
    Querier: 192.168.4.36
    State:         Up Timeout:     197 Version:  2 Groups:      0

Configured Parameters:
```

```
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

What It Means The output shows a list of the Services Router interfaces that are configured for IGMP. Verify the following information:

- Each interface on which IGMP is enabled is listed.
- Next to Version, the number 2 appears.

For more information about show igmp interface, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying the PIM Mode and Interface Configuration

Purpose Verify that PIM sparse mode is configured on all applicable interfaces.

Action From the CLI, enter the show pim interfaces command.

Sample Output

```
user@host> show pim interfaces

Instance: PIM.master
Name                 Stat Mode       IP V State Count DR address
lo0.0                Up   Sparse      4 2 DR        0 127.0.0.1
pime.32769           Up   Sparse      4 2 P2P       0
```

What It Means The output shows a list of the Services Router interfaces that are configured for PIM. Verify the following information:

- Each interface on which PIM is enabled is listed.
- The network management interface, fe–0/0/0, is *not* listed.
- Under Mode, the word Sparse appears.

For more information about show pim interfaces, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying the PIM RP Configuration

Purpose Verify that the PIM RP is statically configured with the correct IP address.

Action From the CLI, enter the show pim rpscommand.

Sample Output

```
user@host> show pim rps

Instance: PIM.master
Address family INET
RP address      Type      Holdtime Timeout Active groups Group prefixes
192.168.14.27   static           0    None             2 224.0.0.0/4
```

What It Means The output shows a list of the RP addresses that are configured for PIM. At least one RP must be configured. Verify the following information:

- The configured RP is listed with the proper IP address.
- Under Type, the word static appears.

Verifying the RPF Routing Table Configuration

Purpose Verify that the PIM RPF routing table is configured correctly.

Action From the CLI, enter the show multicast rpf command.

Sample Output

```
user@host> show multicast rpf

Multicast RPF table: inet.0 , 2 entries...
```

What It Means The output shows the multicast RPF table that is configured for PIM. If no multicast RPF routing table is configured, RPF checks use inet.0. Verify the following information:

- The configured multicast RPF routing table is inet.0.
- The inet.0 table contains entries.

For more information about show multicast rpf, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Part 8

Managing Packet Security

Chapter 23

Configuring IPSec for Secure Packet Exchange

IPSec Tunnel Overview

An IPSec tunnel allows access to a private network through a secure tunnel. This feature is particularly useful when a private network is divided among multiple sites, and transit between the sites must occur on a public network. To ensure secure transport of packets across the public network to the multiple sites, individual tunnels are configured. Each tunnel is defined by a local tunnel endpoint and a remote tunnel endpoint.

Packets with a destination address matching the private network prefix are encrypted and encapsulated in a tunnel packet that is routable through the outside network. The source address of the tunnel packet is the local gateway, and the destination address is the remote gateway. Once the encapsulation packet reaches the other side, the remote end determines how to route the packet.

Security Associations

An IPSec security association (SA) is a set of rules used by IPSec tunnel gateways by which traffic is transported. IPSec security associations are established either manually, through configuration statements, or by Internet Key Exchange (IKE). In the case of manually configured security associations, the connection is established when both ends of the tunnel are configured, and the connections last until one of the endpoints is taken offline. For IKE security associations, connections are established only when traffic is sent through the tunnel, and they dissolve after a preset amount of time or traffic.

Securing IncomingTraffic

Incoming (ingress) traffic across the tunnel must be secured to ensure that the IPSec tunnel is protected. Typically, you secure incoming traffic by configuring a stateful firewall filter that acts on the incoming flow through the tunnel. By filtering all traffic that does not match the remote gateway address, you ensure that only traffic sent by the tunnel endpoint reaches destinations through the IPSec tunnel.

Translating Outgoing Traffic

Outgoing (egress) traffic across the tunnel must be marked with the outbound tunnel endpoint address so that it can be filtered by the stateful firewall filter on the opposite side of the tunnel. Packet tagging is performed by Network Address Translation (NAT). The source address for outbound packets is translated to the local gateway address so that, to the remote gateway, all packets appear to originate from the local endpoint. Address translation enables the remote gateway to filter packets based on source address to determine which packets are to be transported through the tunnel.

Before You Begin

Before you begin configuring an IPSec tunnel, you must have completed these tasks:

- Establish basic connectivity. See “Establishing Basic Connectivity” on page 47.
- Configure network interfaces. See “Configuring Network Interfaces” on page 79.
- Configure one or more routing protocols. See “Configuring Static Routes” on page 285, “Configuring a RIP Network” on page 297, “Configuring an OSPF Network” on page 309, or “Configuring BGP Sessions” on page 331.

Configuring an IPSec Tunnel with Quick Configuration

J-Web Quick Configuration allows you to create IPSec tunnels. Figure 96 shows the Quick Configuration page for IPSec tunnels.

Figure 96: Quick Configuration Page for IPSec Tunnels

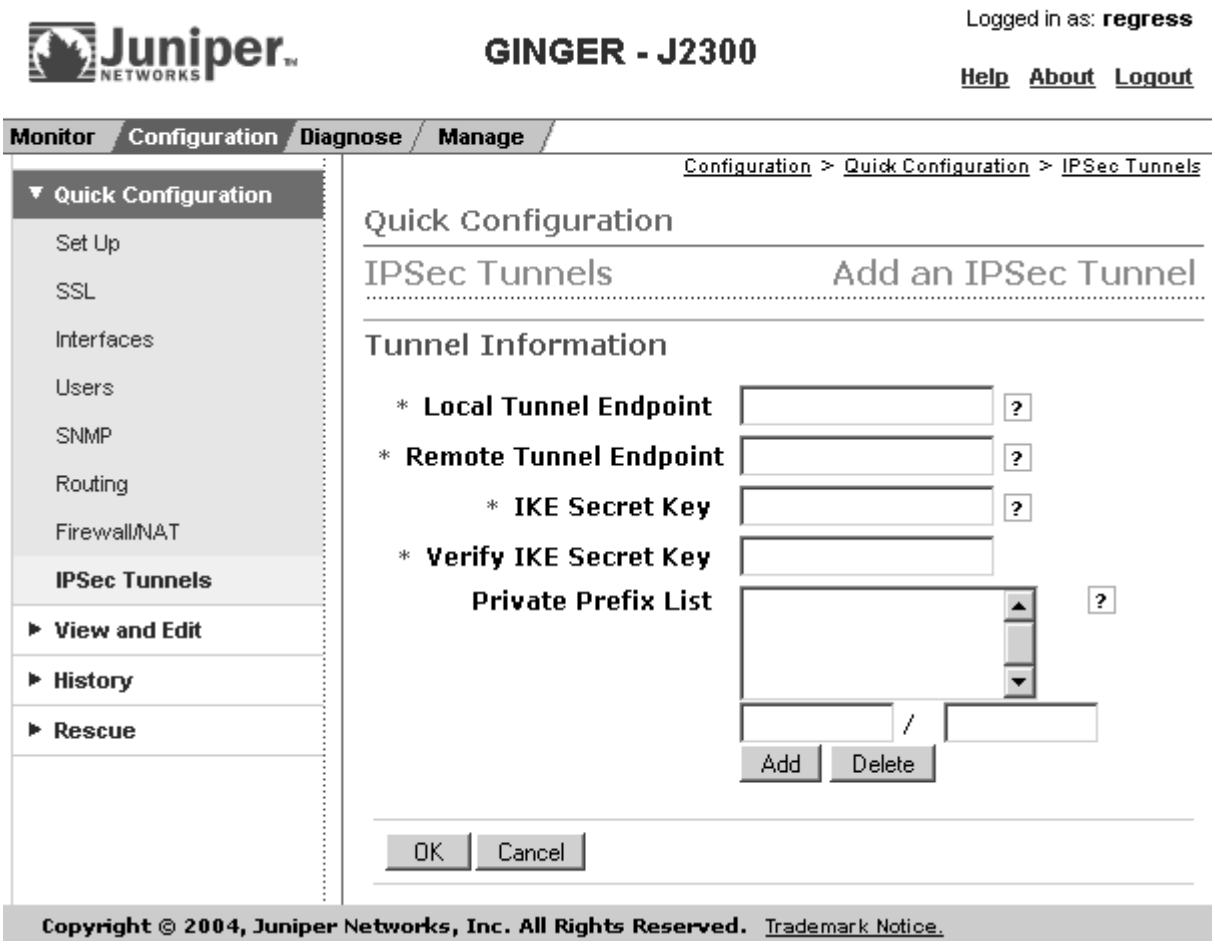

To configure an IPSec tunnel with Quick Configuration:

1. In the J-Web user interface, select **Configuration > IPSec Tunnels**.
2. Enter information into the Quick Configuration page for IPSec Tunnels, as described in Table 187.
3. From the IPSec Tunnels Quick Configuration page, click one of the following buttons:
   - To apply the configuration and return to the Quick Configuration IPSec Tunnels page, click **OK**.
   - To cancel your entries and return to the Quick Configuration for IPSec Tunnels page, click **Cancel**.

4. To check the configuration, see “Verifying the IPSec Tunnel Configuration” on page 496.

Table 187: IPSec Tunnels Quick Configuration Summary

Field	Function	Your Action
Tunnel Information		
Local Tunnel Endpoint (required)	Externally routable IP address that is the local endpoint of the IPSec tunnel	Type the IPSec tunnel’s local endpoint 32-bit IP address, in dotted decimal notation.
Remote Tunnel Endpoint (required)	Externally routable IP address that is the peer endpoint of the IPSec tunnel	Type the IPSec tunnel’s peer endpoint 32-bit IP address, in dotted decimal notation.
IKE Secret Key (required)	Internet Key Exchange key that is preshared to ensure authentication across the IPSec tunnel	Type the IKE key to be used for authentication across the IPSec tunnel. Characters are disguised as you type.
Verify IKE Secret Key (required)	Internet Key Exchange key that is preshared to ensure authentication across the IPSec tunnel	Verify the IKE key by retyping the key to be used for authentication across the IPSec tunnel. Characters are disguised as you type.
Private Prefix List	List of addresses or address prefixes for which the IPSec tunnel is used. Packets whose destination address matches any of the addresses or prefixes in this list are transported through the IPSec tunnel to the remote tunnel endpoint.	1. In the text box at the bottom of the list, enter an IP address or address prefix, in dotted decimal notation. 2. Click **Add**.

Configuring an IPSec Tunnel with a Configuration Editor

To configure a Services Router to transport traffic across a secure IPSec tunnel, you must define the tunnel and configure its components. To configure an IPSec tunnel, perform the following tasks:

- Configuring IPSec Services Interfaces on page 487
- Configuring IPSec Service Sets on page 488
- Configuring an IPSec Stateful Firewall Filter on page 492
- Configuring a NAT Pool on page 494

Configuring IPSec Services Interfaces

To configure an IPSec tunnel, you must configure the following services interfaces:

- *Inside services interface*—Logical interface used to apply the service sets that define the behavior of the IPSec tunnel for outbound traffic (traffic whose next hop is inside the IPSec tunnel).
- *Outside services interface*—Logical interface used to apply the service sets that define the behavior of the IPSec tunnel for inbound traffic (traffic whose next hop is outside the IPSec tunnel).

For the services to be applied, you must first define the logical interfaces to be used.

To configure IPSec inside services interfaces and outside services interfaces:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 188.
3. If you are finished configuring the network, commit the configuration.
4. Go on to "Configuring IPSec Service Sets" on page 488.

Table 188: Configuring IPSec Interfaces

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the **Interfaces** level in the configuration hierarchy.	In the configuration editor hierarchy, select **Interfaces**.	From the top of the configuration hierarchy, enter edit interfaces

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the inside services interface for the IPSec tunnel. On the J-series Services Router, the services interface is always sp-0/0/0.*unit*. The logical interface must have a unit number other than 0. By default, the J-Web Quick Configuration uses the unit number 1001 for inside-service logical interfaces.	1. In the Interface field, click **Add new entry**. 2. In the Interface name field, type sp-0/0/0, and click **OK**. 3. In the Interface field, click **sp-0/0/0**. 4. In the Unit field, click **Add new entry**. 5. In the Interface unit number field, type 1001. 6. In the Service domain box, select **inside** from the drop-down menu. 7. In the Family field, click **inet**. 8. Select the **Primary** box, and click **OK**.	1. Configure the services interface as an inside-service interface: set sp-0/0/0 unit 1001 service-domain inside 2. Configure the services interface as an inet interface: set sp-0/0/0 unit 1001 family inet
Configure the outside services interface for the IPSec tunnel. On the J-series Services Router, the services interface is always sp-0/0/0.*unit*. The logical interface must have a unit number other than 0. By default, the J-Web Quick Configuration uses the unit number 2001 for outside-service logical interfaces.	1. In the Interface field, click **Add new entry**. 2. In the Interface name field, type sp-0/0/0, and click **OK**. 3. In the Interface field, click **sp-0/0/0**. 4. In the Unit field, click **Add new entry**. 5. In the Interface unit number field, type 2001. 6. In the Service domain box, select **outside** from the drop-down menu. 7. In the Family field, click **inet**. 8. Select the **Primary** box, and click **OK**.	1. Configure the services interface as an outside-service interface: set sp-/0/0/0 unit 2001 service-domain outside 2. Configure the services interface as an inet interface: set sp-0/0/0 unit 2001 family inet

Configuring IPSec Service Sets

The next-hop service set defines which services interface to use for all inside-service next hops and all outside-service next hops (traffic inside the network and outside the network). The unit numbers used to define the next-hop interfaces must match exactly the unit numbers used in the interfaces configuration.

When you configure an IPSec service set, you must also configure the local gateway. You then configure an IPSec rule to set the remote gateway on all traffic, configure a security association (SA) with a static IKE key, and configure another rule to act on input traffic. This configuration allows you to set the remote gateway address and perform IKE validation on all incoming traffic through the IPSec tunnel.

Finally, you apply the entire service set.

To configure IPSec service sets:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 189.
3. If you are finished configuring the network, commit the configuration.
4. Go on to "Configuring an IPSec Stateful Firewall Filter" on page 492.

Table 189: Configuring IPSec Service Sets

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the next-hop service set for the IPSec tunnel.	1. From the top of the configuration hierarchy, click **Services**. 2. In the Service sets field, click **Add new entry**. 3. In the Service set name field, type the name of the service set. The name can be any unique string. 4. In the Service type choice field, select **Next hop service** from the drop-down menu. 5. In the Nested configuration field, click **Next hop service**. 6. In the Inside service interface field, type the services interface, including unit number, for the inside-service interface—for example, sp-0/0/0.1001. 7. Click **OK**. 8. In the Nested configuration field, click **Next hop service**. 9. In the Outside service interface field, type the services interface, including the unit number—for example, sp–0/0/0.2002. 10. Click **OK**.	1. From the top of the configuration hierarchy, enter edit services 2. Set the inside-service interface: set service-set *service-set-name* next-hop-service inside-service-interface sp-0/0/0.1001 3. Set the outside-service interface: set service-set *service-set-name* next-hop-service outside-service-interface sp-0/0/0.2001

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the local gateway for the IPSec service set.	1. In the Ipsec vpn options field, click **Configure**. 2. In the Local gateway box, type the IP address of the local tunnel endpoint, in dotted decimal notation—for example, 1.1.1.1.	Set the local gateway address for the service set: set service-set *service-set-name* ipsec-vpn-options local-gateway 1.1.1.1
Configure IPSec rules to set the remote gateway on all traffic to 2.2.2.2. Because the rule applies to all traffic, you must only configure the action (or then statement) for the term.	1. From the top of the configuration hierarchy, click **Services > Ipsec-vpn**. 2. In the Rule field, click **Add new entry**. 3. In the Rule name field, type the name of the rule. The rule name can be any unique string. 4. In the term field, click **Add new entry**. 5. In the Term name field, type the name of the term. It can be any unique string. 6. To configure an action, click **Then**. 7. In the Remote gateway field, type the remote gateway address, in dotted decimal notation—for example, 2.2.2.2. 8. Click **OK**.	1. From the top of the configuration hierarchy, enter edit services ipsec-vpn 2. Configure a rule with a term that sets the remote gateway to 2.2.2.2: set rule *rule-name* term *term-name* then remote-gateway 2.2.2.2

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure an security association with a static IKE key. The IKE key is a preshared key and must be configured exactly the same way at both the local and remote endpoints of the IPSec tunnel. The IKE key is configured as ike policy and then applied using the dynamic statement.	1. From the top of the configuration hierarchy, select **Services > Ipsec-vpn > Ike**. 2. In the Policy field, click **Add new entry**. 3. In the Name box, type the name of the IKE policy. It can be any unique string. 4. Click **Pre-shared key**. 5. In the Key choice field, select **Ascii text** from the drop-down menu. 6. In the Ascii text box, enter the IKE key in plain text. 7. Click **OK**. 8. Navigate to the IPSec rule configured previously. From the top of the configuration hierarchy, click **Services > Ipsec-vp >** *rule-name* **> term** *term-name* **> then**. 9. Click **Dynamic**. 10. In the Ike-policy box, type the name of the IKE policy you configured. 11. Click **OK**.	1. From the top of the configuration hierarchy, enter edit services ipsec-vpn ike 2. Configure the IKE pre-shared key in ASCII text format: set policy *policy-name* pre-shared-key ascii-text *ike-key* 3. Navigate to the IPSec rule configured previously. From the top of the configuration hierarchy, enter edit services ipsec-vpn *rule-name* term *term-name* then. 4. Configure a dynamic security association that applies the IKE policy: set dynamic ike-policy *policy-name*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the IPSec rule so that it acts on input traffic.	1. From the top of the configuration hierarchy, click **Services > Ipsec-vpn > Rule >** *rule-name* . 2. In the Match direction field, select **Input** from the drop-down menu. 3. Click **OK**.	1. From the top of the configuration hierarchy, enter edit services ipsec-vpn rule *rule-name* 2. Set the match direction for the rule: set match-direction input
Apply the IPSec rule to all traffic through the previously configured service set.	1. From the top of the configuration hierarchy, click **Services > Service-set >** *service-set-name* . 2. In the Ipsec vpn rules choice field, select **Ipsec vpn rules** from the drop-down menu. 3. In the Ipsec vpn rules field, click **Add new entry**. 4. In the Rule name box, type the name of the previously configured IPSec rule. 5. Click **OK**.	1. From the top of the configuration hierarchy, enter edit services service-set *service-set-name* 2. Apply the IPSec rule previously configured: set ipsec-vpn-rules *rule-name*

Configuring an IPSec Stateful Firewall Filter

Configure stateful firewall filter rules to ensure that only desired traffic is permitted. This firewall is applied to all inbound traffic from the WAN. For this IPSec tunnel, desired traffic must be from the remote tunnel endpoint, destined for the local tunnel endpoint, and using either IPSec or IKE as an application protocol.

For more information about firewall filters, see “Configuring Firewall Filters and NAT” on page 389.

To configure an IPSec stateful firewall filter:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 190.
3. If you are finished configuring the network, commit the configuration.
4. Go on to “Configuring a NAT Pool” on page 494.

Table 190: Configuring an IPSec Stateful Firewall Filter

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the stateful firewall rule and apply it to inbound traffic.	1. From the top of the configuration hierarchy, click **Services > Stateful firewall**. 2. In the rule field, click **Add new entry**. 3. In the Rule name box, type the name of the rule. It can be any unique string. 4. In the Match direction field, select **Input** from the drop-down menu.	1. From the top of the configuration hierarchy, enter edit services stateful-firewall 2. Create the firewall rule and apply it to input traffic: set rule *rule-name* match-direction input
Create the firewall term to match only desired traffic.	1. In the Term field, click **Add new entry**. 2. In the Term name box, type the name of the term. It can be any unique string. 3. Click **From**. 4. In the Destination address field, click **Add new entry**. 5. In the address field, select **Enter specific value** from the drop-down menu. 6. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation, and click **OK**. 7. In the Source address field, click **Add new entry**. 8. In the address field, select **Enter specific value** from the drop-down menu. 9. In the Address box, type the IP address of the remote tunnel endpoint, in dotted decimal notation, and click **OK**. 10. In the Applications field, click **Add new entry**. 11. In the Application name field, type junos-ipsec-esp, and click **OK**. 12. In the Applications field, click **Add new entry**. 13. In the Application name field, type junos-ike, and click **OK**.	1. Create the firewall term and match all packets with a destination address that matches the local tunnel endpoint: set term *term-name* from destination-address *local-tunnel-end-point-address* 2. Match all packets with a source address that matches the remote tunnel endpoint: set term *term-name* from source-address *remote-tunnel-end-point-address* 3. Match all packets using IPSec as an application protocol: set term *term-name* from applications junos-ipsec-esp 4. Match all packets using IKE as an application protocol: set term *term-name* from applications junos-ike

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the firewall term to accept only desired traffic.	1. Click **OK** to return to the Term name page, and click **Then**. 2. In the Designation field, select **Accept** from the drop-down menu, select the **Yes** box. 3. Click **OK**.	Set the match action to accept: set term *term-name* then accept
Create the firewall term to reject all other traffic.	1. From the top of the configuration hierarchy, click **Services > Stateful firewall > Rule >** *rule-name* 2. In the Term field, click **Add new entry**. 3. In the Term name field, type the name of the term. The name can be any unique string. 4. Click **Then**. 5. In the Designation field, select **Discard** from the drop-down menu.	1. From the top of the configuration hierarchy, enter edit services stateful-firewall rule *rule-name* 2. Configure a term to discard all traffic: set term *term-name* then discard

Configuring a NAT Pool

To hide internal IP addresses from the rest of the Internet, you configure the local tunnel endpoint as the only address in a Network Address Translation (NAT) pool, to ensure that it is the address used for address translation.

To configure a NAT pool for IPSec:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 191.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying the IPSec Tunnel Configuration” on page 496.

Table 191: Configuring a NAT Pool for IPSec

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the NAT pool from which the addresses for network address translation are taken.	1. From the top of the configuration hierarchy, click **Services > Nat**. 2. In the Pool field, click **Add new entry**. 3. In the Pool name field, type the name of the NAT pool. It can be any unique string less than 64 characters long. 4. In the Address choice field, select **Address** from the drop-down menu. 5. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation.	1. From the top of the configuration hierarchy, enter edit services nat 2. Add the local tunnel endpoint to the NAT address pool: set pool *pool-name* address 1.1.1.1

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the router so that all outgoing traffic is matched against the IP address of the local tunnel endpoint.	1. From the top of the configuration hierarchy, click **Services > Nat**. 2. In the Rule field, click **Add new entry**. 3. In the Rule name field, type the name of the rule. The name can be any unique string. 4. In the Match direction field, select **Output** from the drop-down menu. 5. In the Term field, click **Add new entry**. 6. In the Term name field, type the name of the term. The name can be any unique string. 7. Click **From**. 8. In the Source address field, click **Add new entry**. 9. In the address field, select **Enter specific value** from the drop-down menu. 10. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation, and click **OK**.	1. From the top of the configuration hierarchy, enter edit services nat 2. Configure a NAT rule and apply it to all output traffic: set rule *rule-name* match-direction output 3. Configure the rule to match traffic with a source address that is the same as the local tunnel endpoint: set rule *rule-name* term *term-name* from source-address 1.1.1.1
Configure the router so that the source address for traffic through the local endpoint is translated to the local endpoint address.	1. From the top of the configuration hierarchy, click **Services > Nat > Rule >** *rule-name* **Term >** *term-name* 2. Click **Then**. 3. Click **Translated**. 4. In the Source pool field, type the name of the NAT pool in which the local tunnel endpoint is configured. 5. In the Source field, select **Static** from the drop-down menu.	1. From the top of the configuration hierarchy, enter edit services nat rule *rule-name* term *term-name* 2. Configure the source pool: set then translated source-pool *pool-name* 3. Configure the type of translation: set then translated translation-type source static

Verifying the IPSec Tunnel Configuration

To verify the IPSec tunnel configuration, perform the following task.

Verifying IPSec Tunnel Statistics

Purpose Verify that traffic is being sent through the configured IPSec tunnel.

Action From the CLI, enter the show services ipsec-vpn ipsec statistics command.

Sample Output

```
user@host> show services ipsec-vpn ipsec statistics

PIC: sp-0/0/0, Service set: service-set-1

Local gateway: 1.1.1.1, Remote gateway: 2.2.2.2, Tunnel index: 1
ESP Statistics:
 Encrypted bytes:                0
 Decrypted bytes:                0
 Encrypted packets:              0
 Decrypted packets:              0
AH Statistics:
 Input bytes:                    0
 Output bytes:                   0
 Input packets:                  0
 Output packets:                 0
Errors:
 AH authentication failures: 0, Replay errors: 0
 ESP authentication failures: 0, Decryption errors: 0
 Bad headers: 0 Bad trailers: 0
```

What It Means The output shows the statistics for the particular service set that defines the IPSec tunnel, including the local and remote gateway addresses, the number of packets that have been encrypted and transported, and the number of errors and failures. Verify the following information:

- The local and remote tunnel endpoints are configured correctly.
- The number of Authentication Header (AH) and Encapsulation Security Payload (ESP) errors is zero. If these numbers are nonzero, the Services Router might be having a problem either transmitting or receiving encrypted packets through the IPSec tunnel.

For more information about show services ipsec-vpn ipsec statistics, see the *JUNOS Network and Services Interfaces Command Reference*.

Part 9

Upgrading the Services Router

Chapter 24

Performing Software Upgrades and Reboots

You can upgrade the JUNOS Internet software on a Services Router by installing a new version that you download from the Web to a remote server or your computer. Use either the J-Web interface or the CLI to perform the upgrade.

If you need to replace the primary boot device or add a backup boot device on the router, you can configure a boot device with the CLI or with a UNIX or Microsoft Windows computer. You can also configure a boot device to receive core dumps.

Use either the J-Web interface or the CLI to schedule a reboot or system halt on the router, or to perform one immediately. For more information about installing and upgrading JUNOS software, see the *JUNOS System Basics Configuration Guide*.

Upgrade Overview

The Services Router is delivered with the JUNOS Internet software preinstalled. To upgrade the software, you use the J-Web interface or CLI commands to copy a set of software images over the network to memory storage on the Routing Engine.

All junos-jseries software is delivered in signed packages that contain Secure Hash Algorithm 1 (SHA-1) checksums. A package is installed only if the SHA-1 checksum within it matches the SHA-1 hash recorded in its corresponding .sha1 file. (For example, -export.tgz contains -export.tgz and -export.tgz.sha1. The junos-jseries-*release*-export.tgz package is installed only if the SHA-1 hashes match in the two -export.tgz.sha1 files.)

The junos-jseries package completely reinstalls the software. This package rebuilds the file system but retains configuration files, log files, and similar information from the previous version.

Before You Begin

To download software upgrades, you must have a Web account with Juniper Networks. To obtain an account, complete the registration form at the Juniper Networks Web site: https://www.juniper.net/registration/Register.jsp.

Before upgrading, be sure to back up the currently running and active file system and configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. To back up the file system, you must have a removable compact flash drive installed on the J4300 or J6300 Services Router, or a USB drive installed on any J-series Services Router.

To back up the file system to the removable compact flash drive, issue the following command:

```
user@host> request system snapshot media removable-compact-flash
```

To back up the file system to the removable USB drive, issue the following command:

```
user@host> request system snapshot media usb
```

For details about the request system snapshot command, see “Configuring Boot Devices with the CLI” on page 508.

Downloading Software Upgrades from Juniper Networks

Follow these steps to download software upgrades from Juniper Networks:

1. Using a Web browser, follow the links to the download URL on the Juniper Networks Web page. Choose either **Canada and U.S. Version** or **Worldwide Version**:

   - https://www.juniper.net/support/csc/swdist-domestic/

- https://www.juniper.net/support/csc/swdist-ww/

2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Using either the J-Web interface or the CLI, choose the software package for your application.
4. Download the software to a local host.

Installing Software Upgrades with J-Web Quick Configuration

You can use the J-Web interface to install software upgrades from a remote server using FTP or HTTP, or by uploading the file to the router.

- Installing Software Upgrades from a Remote Server on page 503
- Installing Software Upgrades by Uploading Files on page 505

Installing Software Upgrades from a Remote Server

You can use the J-Web interface to install software packages on the Services Router that are retrieved with FTP or HTTP from the location specified.

Figure 97 shows the Install Remote page for the router.

Figure 97: Install Remote Page

To install software upgrades from a remote server:

1. Download the software package as described in "Downloading Software Upgrades from Juniper Networks" on page 502.
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Download the software to your local host or internal software distribution site.
4. In the J-Web interface, select **Manage > Software > Install Remote**.
5. On the Install Remote Quick Configuration page, enter information into the fields described in Table 192.
6. Click **OK**. The software is activated after the router has rebooted.

Table 192: Install Remote Quick Configuration Summary

Field	Function	Your Action
Package Location (required)	Specify the FTP or HTTP server on which the software package resides.	Type the full address of the software package location on the FTP or HTTP server.
User	Specify the username, if the server requires one.	Type the username.
Password	Specify the password, if the server requires one.	Type the password.
Reboot If Required	If this box is checked, the router is automatically rebooted when the upgrade is complete.	Check the box if you want the router to reboot automatically when the upgrade is complete.

Installing Software Upgrades by Uploading Files

You can use the J-Web interface to install software packages uploaded from your computer to the Services Router.

Figure 98 shows the Upload Package page for the router.

Figure 98: Upload Package Page

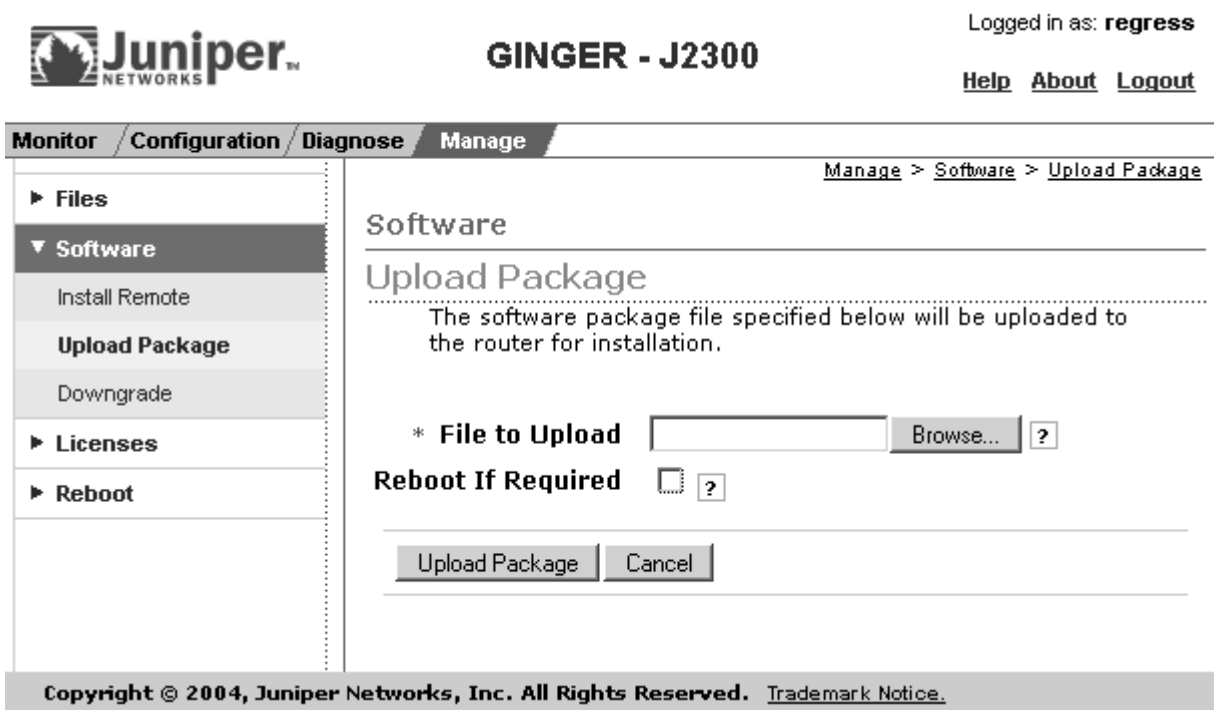

To install software upgrades by uploading files:

1. Download the software package as described in "Downloading Software Upgrades from Juniper Networks" on page 502.
2. In the J-Web interface, select **Manage > Software > Upload Package**.
3. Enter information into the fields described in Table 193 into the Upload Package Quick Configuration page.
4. Click **Upload Package**. The software is activated after the router has rebooted.

Table 193: Upload Package Quick Configuration Summary

Field	Function	Your Action
File to Upload (required)	Specify the location of the software package.	Type the location of the software package, or click **Browse** to navigate to the location.
Reboot If Required	If this box is checked the router is automatically rebooted when the upgrade is complete.	Select the check box if you want the router to reboot automatically when the upgrade is complete.

Installing Software Upgrades with the CLI

To install software upgrades using the CLI:

1. Download the software package as described in "Downloading Software Upgrades from Juniper Networks" on page 502.
2. Copy the software package to the router. We recommend that you copy it to the /var/tmp directory.
3. Install the new package on the Services Router:

   - Customers in the United States and Canada use the following command:

     user@host> **request system software add validate** *path* **/junos-jseries** *release* **-domestic.tgz**

   - All other customers use the following command:

     user@host> **request system software add validate** *path* **/junos-jseries** *release* **-export.tgz**

   Replace *path* with the full pathname to the bundle. Replace *release* with the software release version of the bundle.

4. Reboot the router to activate the junos-jseries software:

   user@host> **request system reboot**

```
Reboot the system ? [yes,no] (no) yes
Shutdown NOW!
```

All the software is activated when you issue the reboot command.

The router then reboots from the primary boot device on which you just installed the software. When the reboot is complete, the router displays the login prompt.

5. If your compact flash is running out of space and you do not wish to downgrade the software to a previous version, you can recover up to 30 MB of space by using the request system software delete-backup CLI command. This command deletes the backup software package.

Downgrading the Software with the J-Web Interface

You can downgrade the software from the J-Web interface. When you downgrade the software to a previous version, the software version that is saved in junos.old is the version of JUNOS that your router is downgraded to. For your changes to take effect, you must reboot the router.

To downgrade software:

1. Go to **Manage > Software > Downgrade**. The previous version (if any) is displayed on this page. For example, you can downgrade to the previously installed version of the router software, /cf/packages/junos-7.0120040930_1745-domestic.

NOTE: Once you perform this operation, you cannot undo it.

2. Select **Downgrade** to downgrade to the previous version of the software or **Cancel** to cancel the downgrade process.

3. When the downgrade process is complete, for the new software to take effect, click **Manage > Reboot** to reboot the router at your convenience.

Downgrading the Software with the CLI

You can revert to the previous set of software using the request system software rollback command in the CLI. Rollback fails if the junos-jseries software bundle cannot be found in /var/sw/pkg.

You can roll back only to the software release that was installed on the Services Router before the current release. Once you issue the request system software rollback command, the old release is loaded and you can not reload it again. Issuing the request system software rollback command again results in an error.

To downgrade to an earlier version of software, follow the procedure for upgrading, using the junos-jseries software bundle labeled for the appropriate release.

Configuring Boot Devices

You can configure boot devices to replace the primary boot device on your Services Router, or to act as a backup boot device.

For more information about installing boot devices, see "Removing and Installing the Primary Compact Flash Disk" on page 523, "Removing and Installing the Removable Compact Flash Disk" on page 525, and "Removing and Installing the USB Drive" on page 527.

This section contains the following topics:

Configuring Boot Devices with the CLI

You can use the request system snapshot CLI command to create a boot device:

```
user@host> request system snapshot <as-primary> <config-size size>
<data-size size> <factory> <media type> <partition> <root-size size>
<swap-size size>
```

Table 194 describes the request system snapshot command options.

Table 194: CLI Request System Snapshot Command Options

Option	Description
as-primary	Creates a snapshot (as with the request system snapshot command) that can be used to replace the medium in the primary compact flash drive. Using the as-primary option allows you to write a snapshot to a device other than the primary compact flash disk (either the removable compact flash disk or a USB drive) and then use that medium as the primary boot medium. The as-primary option can be used on the removable compact flash or a USB drive. **NOTE:** Once the boot device is created as a primary compact flash drive, it can operate only in a primary compact flash drive slot. This option causes the boot medium to be partitioned.

Option	Description
config-size *size*	Specifies the size of the config partition, in megabytes. The default value is 10 percent of physical memory on the boot medium. The config partition is mounted on /config. The configuration files are stored in this partition. **NOTE:** This option causes the boot medium to be partitioned.
data-size *size*	Specifies the size of the data partition, in megabytes. The default value is 0 MB. The data partition is mounted on /data. This space is not used by the router, and can be used for extra storage. **NOTE:** This option causes the boot medium to be partitioned.
factory	Copies only default files that were loaded on the primary compact flash drive when it was shipped from the factory, plus the rescue configuration if one has been set. **NOTE:** Once the boot medium is created with the factory option, it can operate in only the primary compact flash drive slot.
media *type*	Specifies the boot device the software is copied to: ■ compact-flash—Copies software to the primary compact flash drive. ■ removable-compact-flash—Copies software to the removable compact flash drive. This option is available on J4300 and J6300 Services Routers only. ■ usb—Copies software to the device connected to the USB port. **NOTE:** You cannot copy software to the active boot device.
partition	Partitions the medium. This option is usually necessary for boot devices that do not have software already installed on them.
root-size *size*	Specifies the size of the root partition, in megabytes. The default value is the boot device's physical memory minus the config, data, and swap partitions. The root partition is mounted on / and does not include configuration files. **NOTE:** This option causes the boot medium to be partitioned.
swap-size *size*	Specifies the size of the swap partition, in megabytes. The default value is one-third of the physical memory on a boot medium larger than 128 MB, or 0 MB on a smaller boot device. The swap partition is used for swap files and software failure memory snapshots. Software failure memory snapshots are saved to the boot medium only if it is specified as the dump device. For information about the setting the dump device, see "Configuring a Boot Device to Receive Software Failure Memory Snapshots" on page 511. **NOTE:** This option causes the boot medium to be partitioned.

Copying Software Images to Boot Devices with UNIX

To create a boot device with a UNIX computer:

1. If you are copying a boot image to a compact flash drive, first plug the drive into a PCMCIA adapter or USB card reader.

2. Connect the removable medium—compact flash drive or USB—to the UNIX computer.

3. Determine the device address of the drive that the removable medium was mounted on.

4. Copy the software package to the removable medium by entering the following command:

   dd if= *filename* **of=/dev/r** *device address* **bs=64k**

NOTE: The copy process can take several minutes.

The removable medium is now configured to be installed as a primary boot device on a J-series Services Router. For information about installing the boot device, see “Removing and Installing the Primary Compact Flash Disk” on page 523, “Removing and Installing the Removable Compact Flash Disk” on page 525, or “Removing and Installing the USB Drive” on page 527.

Copying Software Images to Boot Devices with Cygwin

To access a raw device on Windows, you must install Cygwin. Cygwin is a Linux environment for Windows. With Cygwin installed, you can use many standard UNIX utilities. These utilities can be accessed from one of the provided shells or from the Windows command prompt.

To create a boot device with Cygwin on a Windows computer:

1. If you are copying a boot image to a compact flash drive, first plug the drive into a PCMCIA adapter or USB card reader.

2. Connect the removable medium—compact flash drive or USB—to the Windows computer on which you have installed Cygwin.

3. Determine the device address of the drive that the removable medium was mounted on.

4. Copy the software package to the removable medium by entering the following command:

   dd if= *filename* **of=/dev/** *device address* **bs=64k**

NOTE: The copy process can take several minutes.

The removable medium is now configured to be installed as a primary boot device on a J-series Services Router. For information about installing the boot device, see “Removing and Installing the Primary Compact Flash Disk” on page

523, “Removing and Installing the Removable Compact Flash Disk” on page 525, or “Removing and Installing the USB Drive” on page 527.

Configuring a Boot Device to Receive Software Failure Memory Snapshots

You can use the set system dump device CLI command to specify the medium to use for the Services Router to store system software failure memory snapshots. In this way, when the operating system fails, if you have specified a system dump device in the configuration, the operating system preserves a snapshot of the state of the router when it failed.

After you reboot the system, the dump device is checked for a snapshot as part of the operating system boot process. If a snapshot is found, it is written to the crash dump directory on the router (/var/crash). The customer support team can examine this memory snapshot to help determine the cause of the system software failure.

NOTE: If the swap partition on the dump device medium is not large enough for a system memory snapshot, either a partial snapshot or no snapshot is written into the crash dump directory.

The syntax for the set system dump device CLI command is as follows:

```
user@host> set system dump device <compact-flash>
<removable-compact-flash > <usb>
```

Table 195 describes the set system dump device command options.

Table 195: CLI Set System Dump Device Command Options

Option	Description
compact-flash	Uses the primary compact flash as the system software failure memory snapshot device.
removable-compact-flash	Uses the compact flash device on the front of the router (J4300 and J6300 only) as the system software failure memory snapshot device.
usb	Uses the device attached to the USB port as the system software failure memory snapshot device.

Deleting a Rescue Configuration

To delete a rescue configuration using the CLI, issue the following command:

```
user@host> request system configuration rescue delete
```

Alternatively, using J-web, **select configuration > rescue menu > delete rescue configuration** to delete the rescue configuration.

Rebooting or Halting a Services Router with the J-Web Interface

You can use the J-Web interface to schedule a reboot or halt the Services Router.

Figure 99 shows the Reboot page for the router.

Figure 99: Reboot Page

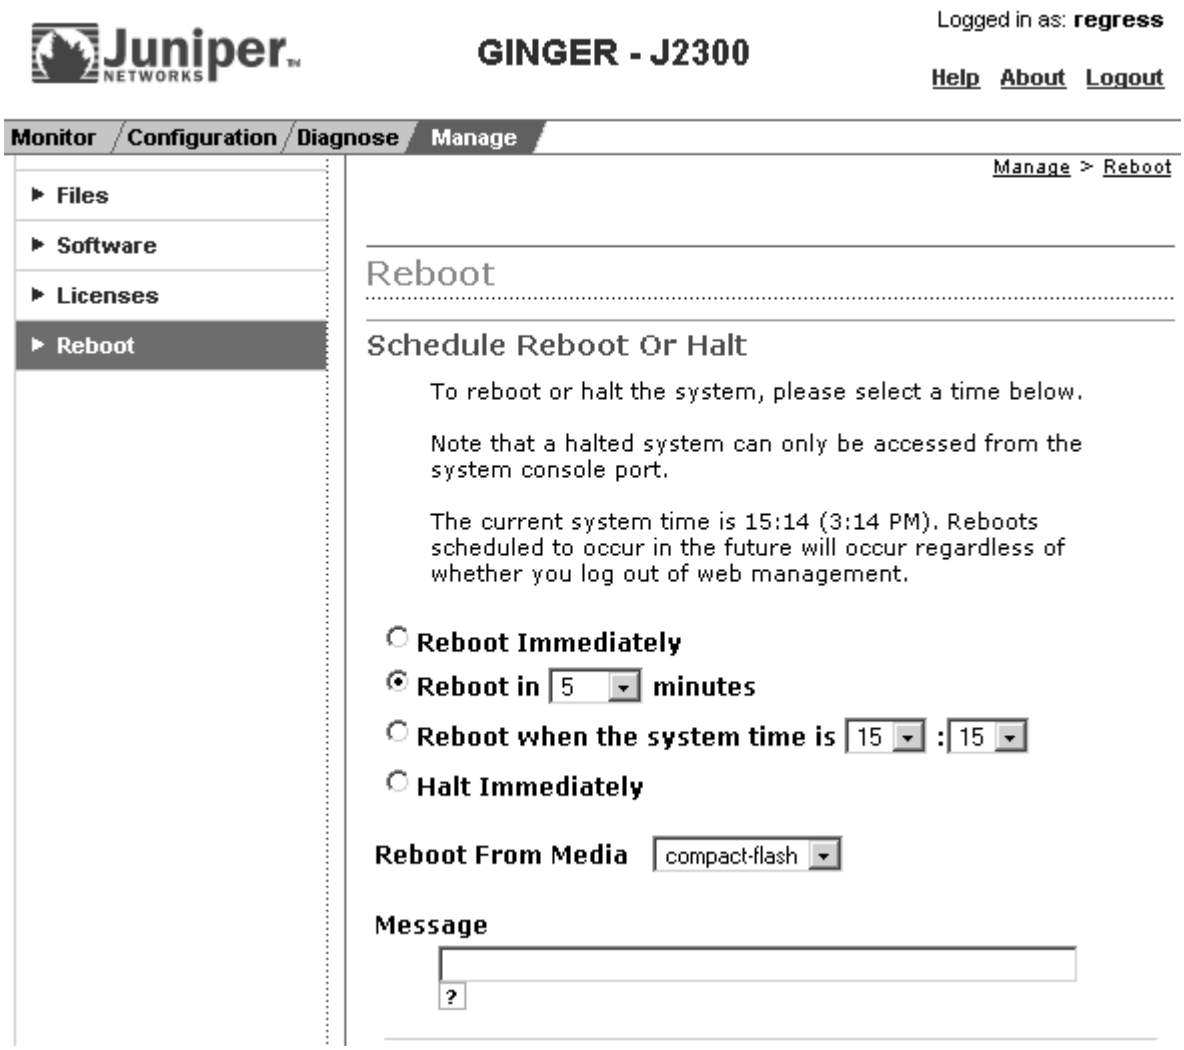

To reboot or halt the router with the J-Web interface:

1. In the J-Web interface, select **Manage > Reboot**.
2. Select one of the following options:
   - **Reboot Immediately**—Reboots the router immediately.

- **Reboot in** *number of* **minutes**—Reboots the router in the number of minutes from now that you specify.
- **Reboot when the system time is** *hour* **:** *minute* —Reboots the router at the absolute time that you specify, on the current day. You must select a 2-digit hour in 24-hour format, and a 2-digit minute.
- **Halt Immediately**—Stops the router software immediately. Once the router software has stopped, you can access the router through the CONSOLE port only.

3. Choose the boot device from the **Reboot from media** drop-down menu:
   - **compact-flash**—Reboots from the primary compact flash drive. This selection is the default choice.
   - **removable-compact-flash**—Reboots from the optional removable compact flash drive. This selection is available on J4300 and J6300 Services Routers only.
   - **usb**—Reboots from the USB drive.
4. (Optional) In the Message box, type a message to be displayed to any users on the router before the reboot occurs.
5. Click **Schedule**. The J-Web interface requests confirmation to perform the reboot or halt.
6. Click **OK** to confirm the operation.
   - If the reboot is scheduled to occur immediately, the router reboots. You cannot access the J-Web interface until the router has restarted and the boot sequence is complete. Once the reboot is complete, refresh the browser window to display the J-Web interface login page.
   - If the reboot is scheduled to occur in the future, the Reboot page displays the time until reboot. You have the option to cancel the request by clicking **Cancel Reboot** on the J-Web interface Reboot page.
   - If the router is halted, all software processes stop and you can access the router through the CONSOLE port only. Reboot the router by pressing any key on the keyboard.

NOTE: If you cannot connect to the router through the CONSOLE port, shut down the router by pressing and holding the power button on the front panel until the POWER ON LED turns off. Once the router has shut down, you can power on the router by pressing the power button again. The POWER ON LED lights during startup and remains steadily green when the router is operating normally.

Rebooting the Services Router with the CLI

You can use the request system reboot CLI command to schedule a reboot of the Services Router:

```
user@host> request system reboot <at time> <in minutes> <media type>
<message “text”>
```

Table 196 describes the request system reboot command options.

Table 196: CLI Request System Reboot Command Options

Option	Description
none	Same as at now (reboots the router immediately).
at *time*	Specifies the time at which to reboot the router. You can specify time in one of the following ways: ■ now—Reboots the router immediately. This is the default. ■ +*minutes*—Reboots the router in the number of minutes from now that you specify. ■ *yymmddhhmm*—Reboots the router at the absolute time on the date you specify. Enter the year, month, day, hour (in 24-hour format), and minute. ■ *hh:mm*—Reboots the router at the absolute time you specify, on the current day. Enter the time in 24-hour format, using a colon (:) to separate hours from minutes.
in *minutes*	Specifies the number of minutes from now to reboot the router. This option is a synonym for the at +*minutes* option.
media *type*	Specifies the boot device to boot the router from: ■ compact-flash—Reboots from the primary compact flash drive. This is the default. ■ removable-compact-flash—Reboots from the optional removable compact flash drive. This option is available on J4300 and J6300 Services Routers only. ■ usb—Reboots from the USB drive.
message *"text"*	Provides a message to display to all system users before the router reboots.

Halting the Services Router with the CLI

You can use the request system halt CLI command to halt the Services Router:

```
user@host> request system halt <at time> <in minutes> <media type>
<message “text”>
```

When the router is halted, all software processes stop and you can access the router through the CONSOLE port only. Reboot the router by pressing any key on the keyboard.

NOTE: If you cannot connect to the router through the CONSOLE port, shut down the router by pressing and holding the power button on the front panel until the POWER ON LED turns off. Once the router has shut down, you can power on the router by pressing the power button again. The POWER ON LED lights during startup and remains steadily green when the router is operating normally.

Table 197 describes the request system halt command options.

Table 197: CLI Request System Halt Command Options

Option	Description
none	Same as at now (stops software processes on the router immediately).
at *time*	Time at which to stop the software processes on the router. You can specify time in one of the following ways: ■ now—Stops the software processes immediately. This is the default. ■ +*minutes* —Stops the software processes in the number of minutes from now that you specify. ■ *yymmddhhmm* —Stops the software processes at the absolute time you specify. Enter the year, month, day, hour (in 24-hour format), and minute. ■ *hh:mm* —Stops the software processes at the absolute time that you specify, on the current day. Enter the time in 24-hour format, using a colon (:) to separate hours from minutes.
in *minutes*	Specifies the number of minutes from now to stop the software processes on the router. This option is a synonym for the at +*minutes* option.
media *type*	Specifies the boot device to boot the router from after the halt: ■ compact-flash—Reboots from the primary compact flash drive. This is the default. ■ removable-compact-flash—Reboots from the optional removable compact flash drive. This option is available on J4300 and J6300 Services Routers only. ■ usb—Reboots from the USB drive.
message *"text"*	Provides a message to display to all system users before the software processes on the router are stopped.

Chapter 25

Replacing and Troubleshooting Hardware Components

Because many of the Services Router's hardware components are field-replaceable units (FRUs), you can remove and replace them yourself. When you need to replace a router component, contact your customer support or sales representative to order the field-replaceable unit (FRU) that contains the component. For instructions, see "Contacting Customer Support and Returning Hardware" on page 603.

This chapter contains the following topics:

Replacing Hardware Components

This section contains the following topics:

Tools and Parts Required

To replace hardware components, you need the tools and parts listed in Table 198.

Table 198: Tools and Parts Required

Tool or Part	Components
Electrostatic bag or antistatic mat	All
Electrostatic discharge (ESD) grounding wrist strap	All
Phillips (+) screwdriver, number 2	■ PIM ■ DRAM ■ Compact flash

Replacing the Console Port Cable

The RJ-45 port labeled CONSOLE on the Services Router's front panel allows you to connect the router to an external management device, such as a laptop or a terminal server. For cable specifications, see "Network Cable Specifications and Connector Pinouts" on page 551.

To replace the console port cable, follow this procedure:

1. Locate an appropriate replacement cable and connector.
2. Plug the Ethernet connector at either end of the cable into the CONSOLE port on the front panel (see Figure 22 and Figure 23).
3. Plug the connector at the other end of the cable into the external management device. If you are connecting to a DB-9 serial port, use the provided RJ-45 to DB-9 serial port adapter.

Replacing a PIM

Physical Interface Modules (PIMs) in J4300 and J6300 Services Routers are field replaceable. The router must be powered off before the PIMs are removed or installed. This section contains the following topics:

Removing a PIM

The PIMs are installed in the front of the Services Router. A PIM weighs less than 1 lb (0.5 kg).

To remove a PIM (see Figure 100):

1. Place an electrostatic bag or antistatic mat on a flat, stable surface to receive the PIM.
2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the Services Router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
3. Press and release the power button to power off the router. Verify that the POWER ON LED blinks and then turns off.
4. Label the cables connected to the PIM so that you can later reconnect each cable to the correct PIM.
5. Disconnect the cables from the PIM.
6. If necessary, arrange the cables to prevent them from dislodging or developing stress points:
   - Secure the cable so that it is not supporting its own weight as it hangs to the floor.
   - Place excess cable out of the way in a neatly coiled loop.
   - Use fasteners to maintain the shape of cable loops.
7. Loosen the captive screws on each side of the PIM faceplate.
8. Grasp the handles on each side of the PIM faceplate and slide the PIM out of the router. Place it in the electrostatic bag or on the antistatic mat.
9. If you are not reinstalling a PIM into the emptied slot, install a blank PIM panel over the slot to maintain proper airflow.

Figure 100: Removing a PIM

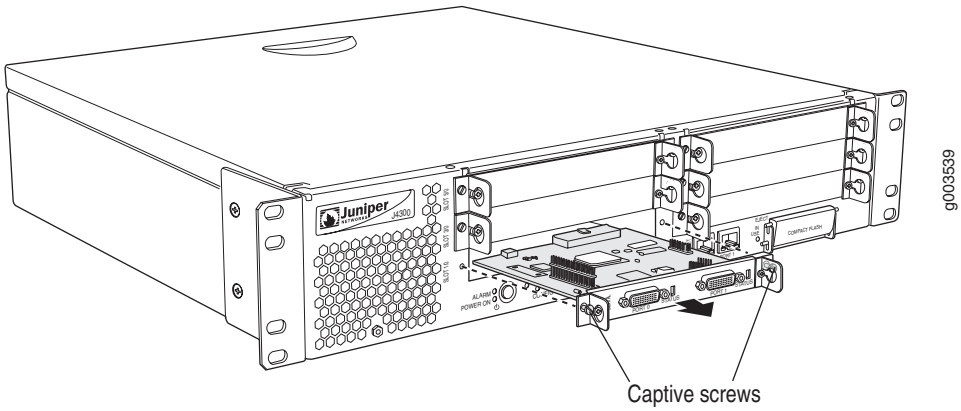

Installing a PIM

To install a PIM (see Figure 101):

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the Services Router is disconnected from earth ground. For more information about ESD, see “Preventing Electrostatic Discharge Damage” on page 567.
2. Press and release the power button to power off the router. Verify that the POWER ON LED blinks and then turns off.
3. Align the notches in the connector at the rear of the PIM with the notches in the PIM slot in the Services Router, and then slide the PIM in until it lodges firmly in the router.

CAUTION: Slide the PIM straight into the slot to avoid damaging the components on the PIM.

4. Tighten the captive screws on each side of the PIM faceplate.
5. Insert the appropriate cables into the cable connectors on the PIM.
6. If necessary, arrange the cables to prevent them from dislodging or developing stress points:

- Secure the cable so that it is not supporting its own weight as it hangs to the floor.
- Place excess cable out of the way in a neatly coiled loop.
- Use fasteners to maintain the shape of cable loops.

7. Press and release the power button to power on the router. Verify that the POWER ON LED lights steadily after you press the power button.
8. Verify that the PIM status LED lights steadily green to confirm that the PIM is online.

   You can also verify correct PIM functioning by issuing the show chassis fpc pic-status command described in the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Figure 101: Installing a PIM

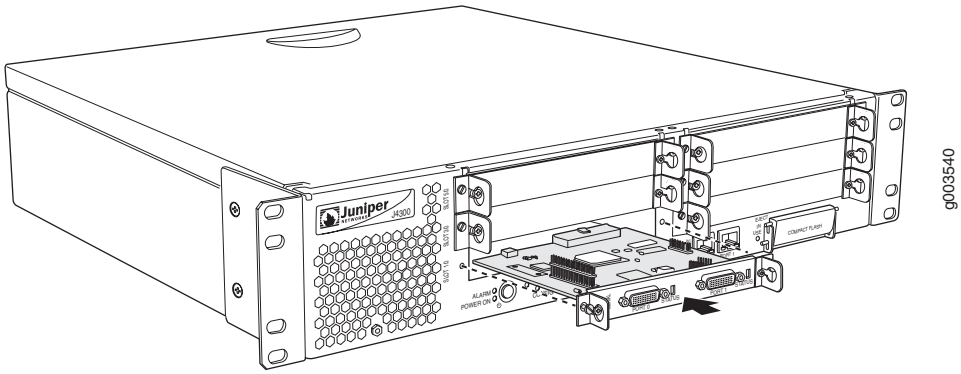

Replacing PIM Cables

Removing and installing PIM cables does not affect Services Router function, except that a PIM does not receive or transmit data while its cable is disconnected. To replace a PIM cable, perform the following procedures:

- "Removing a PIM Cable" on page 522
- "Installing a PIM Cable" on page 522

Removing a PIM Cable

To remove a PIM cable:

1. If you are removing all cables connected to the PIM, issue the following CLI command to take the PIM offline:

   ```
   user@host> request chassis pic fpc-slot fpc-slot pic-slot pim-slot
   offline
   ```

   For example, to take the PIM in slot 4 offline, enter the following command:

   ```
   user@host> request chassis pic fpc-slot 4 pic-slot 0 offline
   ```

   For more information about the command, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

2. Unplug the cable from the cable connector port.

3. Detach the cable from the destination port.

Installing a PIM Cable

To install a PIM cable:

1. Have ready a length of the type of cable used by the PIM. For cable specifications, see "Network Cable Specifications and Connector Pinouts" on page 551.

2. Insert the cable connector into the cable connector port on the PIM faceplate.

3. Arrange the cable as necessary to prevent it from dislodging or developing stress points:

   - Secure the cable so that it is not supporting its own weight as it hangs to the floor.
   - Place excess cable out of the way in a neatly coiled loop.
   - Use fasteners to maintain the shape of cable loops.

4. Insert the other end of the cable into the destination port.

5. Repeat the previous steps for any additional cables.

6. If the PIM is offline (its status LED is steadily red), issue the following CLI command to bring the PIM online:

   ```
   user@host> request chassis pic fpc-slot fpc-slot pic-slot pim-slot
   online
   ```

For example, to bring the PIM in slot 4 online, enter the following command:

```
user@host> request chassis pic fpc-slot 4 pic-slot 0 online
```

For more information about the command, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

7. Verify that the PIM status LED shines steadily green to confirm that the PIM is online.

   You can also verify correct PIM functioning by issuing the show chassis fpc pic-status command described in the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Removing and Installing the Primary Compact Flash Disk

The primary compact flash drive is located in a slot at the rear of the Services Router as shown in Figure 2, Figure 7, and Figure 8. The compact flash disk that you install in the compact flash drive provides primary storage for the router. It can accommodate software images, configuration files, and microcode.

For information about configuring the primary compact flash disk, see "Configuring Boot Devices" on page 508.

To remove and install a primary compact flash disk, perform the following procedures:

- "Removing the Primary Compact Flash Disk" on page 523
- "Installing the Primary Compact Flash Disk" on page 524

Removing the Primary Compact Flash Disk

To remove the primary compact flash disk (see Figure 102):

1. Place an electrostatic bag or antistatic mat on a flat, stable surface.
2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
3. Press and release the power button to power off the router. Wait for the POWER ON LED to turn off.
4. Remove the power cord from the power supply.
5. Loosen the thumbscrew that secures the primary compact flash drive cover on the rear of the chassis.
6. Remove the compact flash drive cover.

7. Gently grasp the compact flash disk, and slide it out of the connector.

8. Place the compact flash disk on the antistatic mat or in the electrostatic bag (see Figure 102).

Figure 102: Removing the Primary Compact Flash Disk

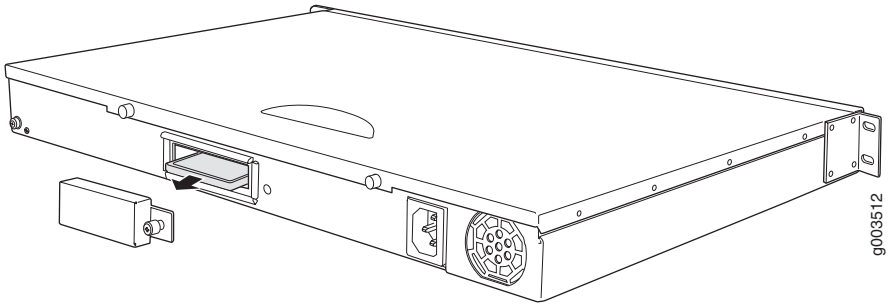

Installing the Primary Compact Flash Disk

To install the primary compact flash disk (see Figure 103):

NOTE: If you plan to boot the Services Router from the primary compact flash disk, you must first configure the primary compact flash disk in another router or with a computer running UNIX or Cygwin. For more information, see "Configuring Boot Devices" on page 508.

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.

2. Press and release the power button to power off the router. Wait for the POWER ON LED to turn off.

3. Remove the power cord from the power supply.

4. Loosen the thumbscrew that secures the primary compact flash drive cover on the rear of the chassis.

5. Remove the compact flash drive cover.

6. Slide the compact flash disk into the connector on the Routing Engine (see Figure 103).

7. Replace the compact flash drive cover.

8. Tighten the thumbscrew that secures the compact flash drive cover to the rear of the chassis.

9. Install the power cord into the power supply.

10. Press and release the power button to power on the router. Verify that the POWER ON LED lights steadily after you press the power button.

Figure 103: Installing the Primary Compact Flash Disk

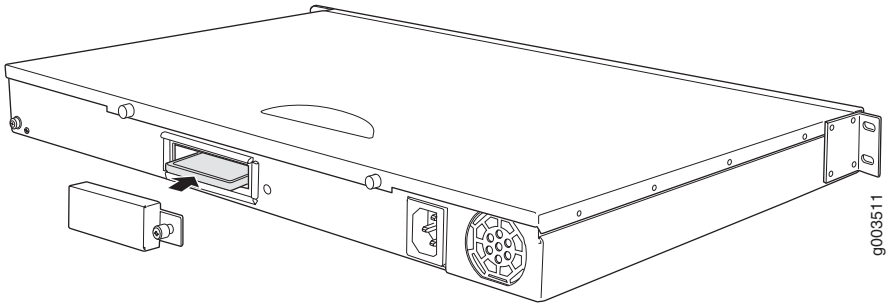

Removing and Installing the Removable Compact Flash Disk

The removable compact flash drive is an optional component on J4300 and J6300 Services Routers. The removable compact flash disk provides secondary storage for the router. It can accommodate software images, configuration files, and microcode. If the primary compact flash disk fails on startup, the router boots from the removable compact flash disk.

For information about configuring the removable compact flash disk, see "Configuring Boot Devices" on page 508.

To remove and install a removable compact flash disk, perform the following procedures:

- "Removing the Removable Compact Flash Disk" on page 525
- "Installing the Removable Compact Flash Disk" on page 527

Removing the Removable Compact Flash Disk

NOTE: Depending on your configuration, the Services Router might not have a backup compact flash drive. If no backup compact flash drive is installed, proceed directly to the next section, "Installing the Removable Compact Flash Disk" on page 527.

The removable compact flash drive is located in a slot on the front panel of the Services Router. To remove the removable compact flash disk (see Figure 104):

1. Place an electrostatic bag or antistatic mat on a flat, stable surface.

2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.

3. Verify the CF REMOVE LED is off.

   If the CF REMOVE LED is on, the router might have booted from the removable compact flash disk.

   To see which device the router used to boot, issue the show system storage command from the CLI. For example:

   ```
   user@host> show system storage

   Filesystem        512-blocks       Used      Avail Capacity Mounted on
   /dev/ad0s1a           218254     175546      40526     81%  /
   ...
   ```

   The boot device is mounted on /. The *primary* compact flash disk is located at ad0. The *removable* compact flash disk is located at ad2. The USB drive is located at usb0. This example shows that the router booted from the primary compact flash disk.

4. If the show system storage output indicates that the router booted from the removable compact flash disk, press and release the power button to power off the router. Wait for the POWER ON LED to turn off before you remove the compact flash drive.

5. Slide the compact flash drive door up to unlatch the door, then tilt the top of the door out (see Figure 104).

6. Eject the removable compact flash disk by pressing the button to the left of the compact flash drive once to unlock the button, and again to eject the compact flash drive.

7. Gently grasp the compact flash disk, and slide it out of the connector.

8. Place the compact flash disk on the antistatic mat or in the electrostatic bag.

Figure 104: Removing the Removable Compact Flash Disk

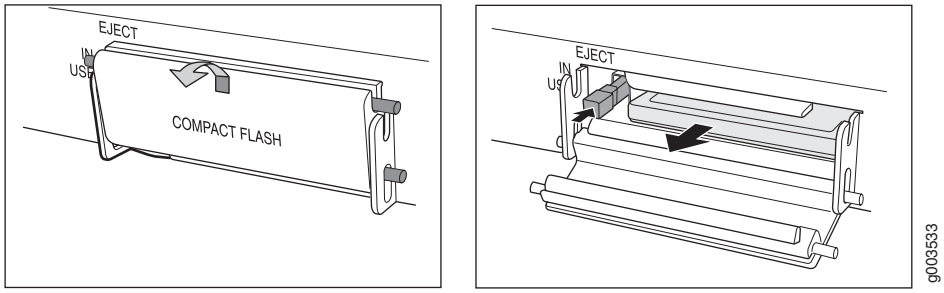

Installing the Removable Compact Flash Disk

To install the removable compact flash disk, follow this procedure (see Figure 105):

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
2. Slide the compact flash door up to unlatch the door, then tilt the top of the door out (see Figure 105).
3. Slide the compact flash disk into the connector on the Routing Engine.
4. Tilt the compact flash door in, and slide it down until it is secured.
5. To configure the removable compact flash disk with the request system snapshot command, see "Configuring Boot Devices with the CLI" on page 508.

Figure 105: Installing the Removable Compact Flash Disk

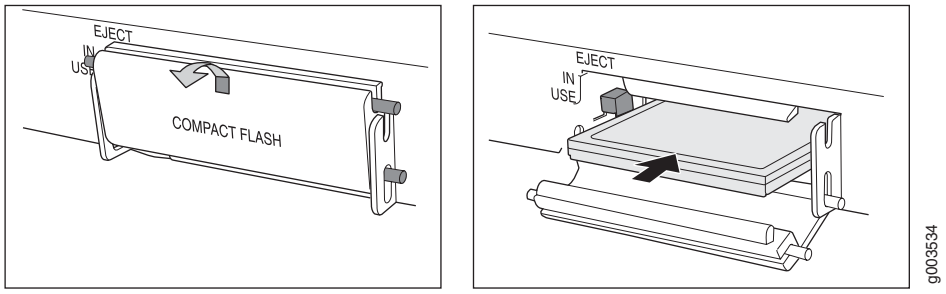

Removing and Installing the USB Drive

The USB drive is an optional component on J-series Services Routers. If installed, the USB drive provides secondary storage for the router. It can accommodate software images, configuration files, and microcode. If the

primary compact flash disk fails on startup, and the removable compact flash disk is not installed or fails, the router boots from the USB drive.

For information about configuring the USB drive, see "Configuring Boot Devices" on page 508.

NOTE: For a list of supported USB drives, see the J-series release notes at http://www.juniper.net.

To remove and install a USB drive, perform the following procedures:

- "Removing the USB Drive" on page 528
- "Installing the USB Drive" on page 529

Removing the USB Drive

NOTE: Depending on your configuration, the Services Router might not have a USB drive. If no USB drive is installed, proceed directly to the next section, "Installing the USB Drive" on page 529.

The USB drive is installed into the USB port on the front panel of the Services Router. To remove the USB drive:

1. Place an electrostatic bag or antistatic mat on a flat, stable surface.
2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
3. Verify that the router did not boot from the USB drive by issuing the show system storage command from the CLI. For example:

   ```
   user@host> show system storage

   Filesystem        512-blocks       Used      Avail Capacity Mounted on
   /dev/ad0s1a           218254     175546      40526     81%  /
   ...
   ```

   The boot device is mounted on /. The primary compact flash disk is located at ad0. The removable compact flash disk is located at ad2. The USB drive is located at usb0. This example shows that the router booted from the primary compact flash disk.

4. If the show system storage output indicates that the router booted from the USB drive, press and release the power button to power off the router. Wait for the POWER ON LED to turn off before you remove the USB drive.

5. Gently grasp the USB drive and slide it out of the USB port.

6. Place the USB drive on the antistatic mat or in the electrostatic bag.

Installing the USB Drive

To install the USB drive:

NOTE: For a list of supported USB drives, see the J-series release notes at http://www.juniper.net.

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.

2. Orient the USB drive with the USB port on the front panel of the router.

3. Insert the USB drive into the USB port. If the USB drive does not easily slide into the port, it might not be oriented correctly. Turn the USB drive upside-down and try again.

4. To configure the USB drive with the request system snapshot command, see "Configuring Boot Devices with the CLI" on page 508.

Removing and Installing DRAM Modules

The DRAM installed on the Routing Engine provides storage for the routing and forwarding tables and for other Routing Engine processes. The design of the Routing Engine allows you to modify the DRAM configuration by adding DIMM memory modules to the Routing Engine board, or removing DIMMs from the board. The Routing Engine contains one or two 168-pin DIMMs.

To modify the DRAM configuration, use the following procedures:

- "Removing a DRAM Module" on page 529
- "Installing a DRAM Module" on page 531

Removing a DRAM Module

NOTE: Depending on your configuration, the Services Router might have an empty DRAM slot. If you are adding a single DIMM to the DRAM configuration, proceed directly to the next section, "Installing a DRAM Module" on page 531.

The DRAM modules are located on the top of the Routing Engine.
To remove a DRAM module:

1. Place an electrostatic bag or antistatic mat on a flat, stable surface.
2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
3. Press and release the power button to power off the router. Wait for the POWER ON LED to turn off.
4. Loosen the thumbscrews at the rear of the chassis that secure the cover to the chassis.
5. Slide the cover off the chassis.
6. To release the DRAM module, press the plastic ejectors on both sides of the module (see Figure 106).
7. Grasp the DRAM module, being careful not to touch any electrical components on the module, and firmly pull it out of the slot on the Routing Engine.
8. Place the DRAM module on the antistatic mat or in the electrostatic bag.

Figure 106: Removing a DRAM Module from the Routing Engine

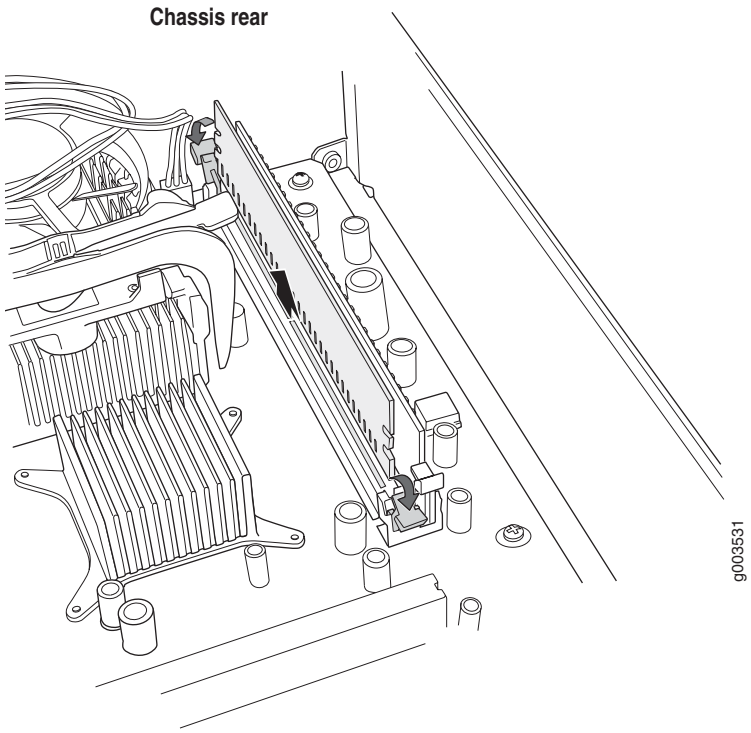

Installing a DRAM Module

To install a DRAM module onto the Routing Engine (see Figure 107):

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
2. Press and release the power button to power off the router. Wait for the POWER ON LED to turn off.
3. Loosen the thumbscrews at the rear of the chassis that secure the cover to the chassis.
4. Slide the cover off the chassis.
5. Remove the DRAM module from its electrostatic bag.
6. To open the empty DRAM slot, press the plastic ejectors on both sides (see Figure 107).
7. Grasp the DRAM module by the edges, being careful not to touch any electrical components.

8. Pressing firmly on both ends, push the module into the slot until the ejectors return completely to the closed position (see Figure 107).

9. Slide the cover onto the chassis.

10. Tighten the thumbscrews at the rear of the chassis that secure the cover to the chassis.

11. Press and release the power button to power on the router. Verify that the POWER ON LED lights steadily after you press the power button.

    You can view the DRAM configuration and verify it was installed correctly by issuing the show chassis routing-engine command, described in the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Figure 107: Installing a DRAM Module

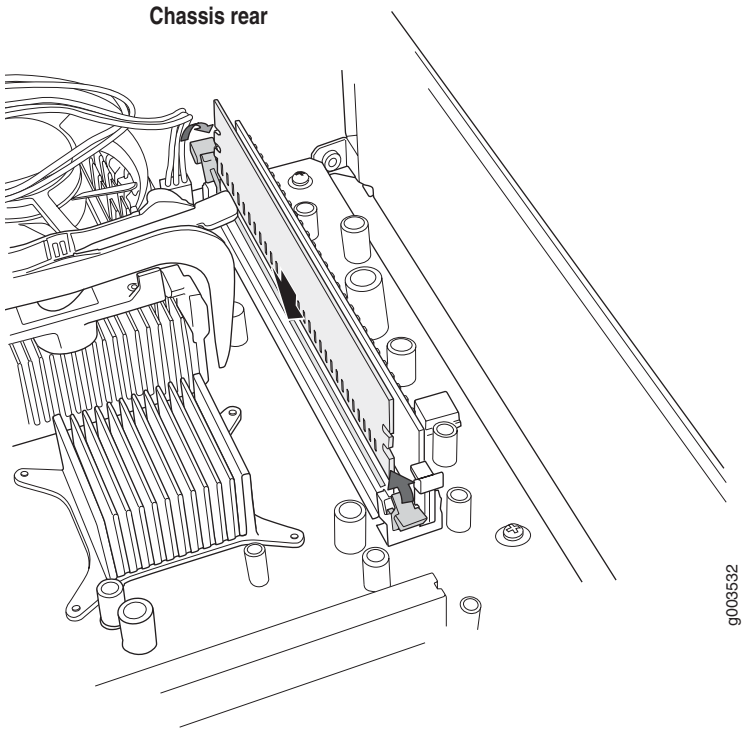

Replacing a Power Supply Cord in a J2300 or J4300 Router

To replace the power cord for an AC power supply:

1. Locate a replacement power cord with the type of plug appropriate for your geographical location (see “AC Power, Connection, and Power Cord Specifications” on page 547).

2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.

3. Press and release the power button to power off the Services Router. Wait for the POWER ON LED to turn off.

4. Unplug the power cord from the power source receptacle.

5. Unplug the power cord from the appliance inlet on the power supply faceplate.

6. Insert the appliance coupler end of the replacement power cord into the appliance inlet on the power supply faceplate.

7. Insert the power cord plug into an AC power source receptacle.

NOTE: The router must be connected to a dedicated AC power feed. For information about connecting to AC power sources, see "Connecting Power to the Services Router" on page 43.

8. Press and release the power button to power on the router. Verify that the POWER ON LED lights steadily after you press the power button.

9. Verify that the power cord does not block access to router components or drape where people might trip on it.

Replacing Power System Components in a J6300 Router

The J6300 Services Router has one or two load-sharing AC power supplies (see Figure 8), located at the right rear of the chassis. Each AC power supply provides power to all components in the router. The AC power supplies are fully redundant. If one power supply fails or is removed, the remaining power supply instantly assumes the entire electrical load. One power supply can provide full power for as long as the router is operational.

Each J6300 power supply is hot-insertable and hot-removable. To replace a power supply in a J6300 router, use the following procedures:

- "Removing a Power Supply in a J6300 Router" on page 534
- "Installing a Power Supply in a J6300 Router" on page 535
- "Replacing a Power Supply Cord in a J6300 Router" on page 536

Removing a Power Supply in a J6300 Router

The power supplies are located at the right rear of the chassis. A power supply weighs 2.4 lb (1.1 kg).

CAUTION: Do not leave a power supply slot empty for more than a short time while the Services Router is operational. The power supply or a blank power supply panel must remain in the chassis for proper airflow.

To remove a power supply from a J6300 Services Router:

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.
2. Unplug the power cord from the power source receptacle.
3. Unplug the power cord from the appliance inlet on the power supply faceplate.
4. Slide the red ejector tab on the power supply faceplate to the right and hold it in place, to unlock the power supply.
5. Grasp the handle on the power supply faceplate, and pull firmly to start removing the power supply. Slide it halfway out of the chassis (see Figure 108).
6. Place one hand underneath the power supply to support it and slide it completely out of the chassis.
7. If you are not reinstalling a power supply into the emptied slot, install a blank power supply panel over the slot.

Figure 108: Removing a Power Supply

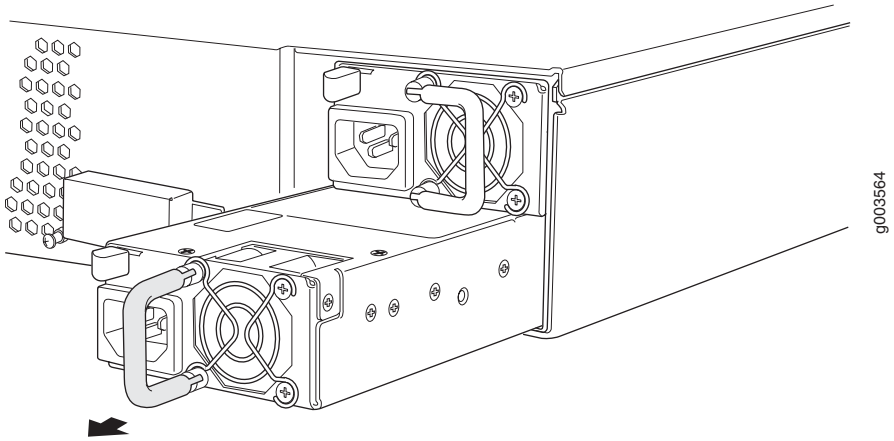

Installing a Power Supply in a J6300 Router

To install a power supply in a J6300 Services Router (see Figure 109):

1. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see “Preventing Electrostatic Discharge Damage” on page 567.
2. Using both hands, slide the power supply into the chassis until you feel resistance.
3. Firmly push the power supply into the chassis until it comes to a stop. Make sure that the power supply faceplate is flush with any adjacent power supply faceplate.
4. Insert the appliance coupler end of a power cord into the appliance inlet on the power supply faceplate.
5. Insert the power cord plug into an AC power source receptacle.

NOTE: Each power supply must be connected to a dedicated AC power feed. For information about connecting to AC power sources, see “Connecting Power to the Services Router” on page 43.

6. Verify that the power cord does not block access to router components or drape where people might trip on it.

Figure 109: Installing an AC Power Supply

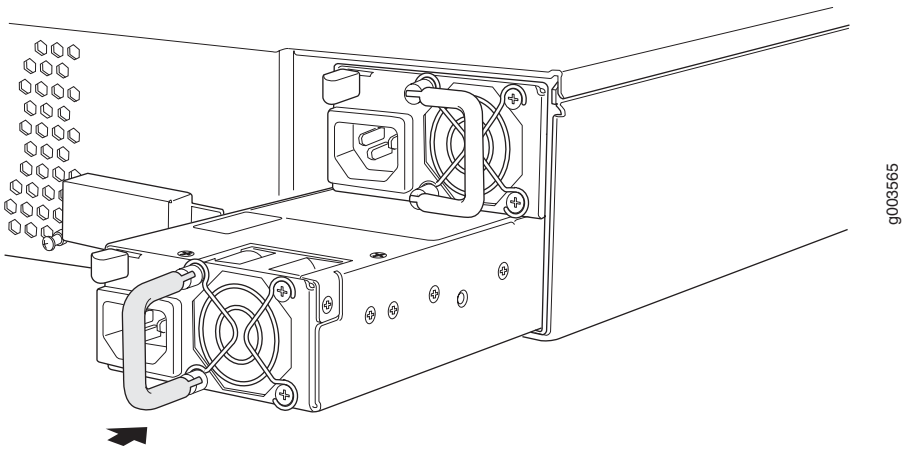

Replacing a Power Supply Cord in a J6300 Router

To replace the power cord for a redundant power supply:

1. Locate a replacement power cord with the type of plug appropriate for your geographical location (see “AC Power, Connection, and Power Cord Specifications” on page 547).
2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see “Preventing Electrostatic Discharge Damage” on page 567.
3. Unplug the power cord from the power source receptacle.
4. Unplug the power cord from the appliance inlet on the power supply faceplate.
5. Insert the appliance coupler end of the replacement power cord into the appliance inlet on the power supply faceplate.
6. Insert the power cord plug into an AC power source receptacle.

NOTE: Each power supply must be connected to a dedicated AC power feed. For information about connecting to AC power sources, see “Connecting Power to the Services Router” on page 43.

7. Verify that the power cord does not block access to Services Router components or drape where people might trip on it.

Troubleshooting Hardware Components

This section provides an overview of the resources you can use to troubleshoot hardware problems on the Services Router:

Chassis Alarm Conditions

When the Routing Engine detects an alarm condition, it lights the yellow (amber) ALARM LED on the front panel as appropriate. To view a more detailed description of the alarm cause, issue the show chassis alarms CLI command:

NOTE: The ALARM LED on the Services Router is a single color alarm regardless of the severity of the alarm condition (critical, major, or minor). When an alarm condition triggers the LED, you see the yellow light turn on.

```
user@host> show chassis alarms
```

Table 199 describes alarms that can occur for a chassis component such as the Routing Engine or a Physical Interface Module (PIM).

Table 199: Chassis Alarm Conditions

Component	Alarm Conditions	Remedy	Alarm Severity
Alternative boot media	The Services Router boots from an alternative boot device—the removable compact flash disk or the USB drive. Typically, the router boots from the primary compact flash disk. If you configured your router to boot from an alternative boot device, ignore this alarm condition.	If you did not configure the router to boot from an alternative boot device, contact JTAC. (See "Contacting the Juniper Networks Technical Assistance Center" on page 538.)	Yellow (critical)
PIM	A PIM has failed. When a PIM fails, it attempts to reboot. If the Routing Engine detects that a PIM is rebooting too often, it shuts down the PIM.	Replace the failed PIM. (See "Replacing a PIM" on page 518.)	Red (warning)

Component	Alarm Conditions	Remedy	Alarm Severity
Routing Engine	An error occurred during the process of reading or writing compact flash.	Reformat the compact flash and install a bootable image. (See "Performing Software Upgrades and Reboots" on page 501.) If this remedy fails, you must replace the failed Routing Engine. To contact JTAC, see "Contacting the Juniper Networks Technical Assistance Center" on page 538.	Yellow (critical)
	Routing Engine temperature is too warm.	■ Check the room temperature. (See "Router Environmental Tolerances" on page 543.) ■ Check the air flow. (See "General Site Guidelines" on page 541.) ■ Check the fans. (See "J2300 Cooling System" on page 15 or "J4300 and J6300 Cooling System" on page 27.) If you must replace a fan or the Routing Engine, contact JTAC. (See "Contacting the Juniper Networks Technical Assistance Center" on page 538.)	Yellow (critical)
	Routing Engine temperature is too hot.		Red (warning)
	Routing Engine fan has failed.	Replace the failed fan. To contact JTAC, see "Contacting the Juniper Networks Technical Assistance Center" on page 538.	Red (warning)

Contacting the Juniper Networks Technical Assistance Center

If you need assistance while troubleshooting a Services Router, open a support case using the Case Manager link at http://www.juniper.net/support/, or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).

Part 10

J-series Requirements and Specifications

Chapter 26

Preparing for Router Installation

This chapter describes how to prepare for installation of a J-series Services Router. It discusses the following topics:

General Site Guidelines

The following precautions help you plan an acceptable operating environment for your Services Router and avoid environmentally caused equipment failures:

- For the cooling system to function properly, the airflow around the chassis must be unrestricted. Allow at least 6 in. (15.2 cm) of clearance between the front and back of the chassis and adjacent equipment. Ensure that there is adequate circulation in the installation location.
- Follow ESD procedures described in “Preventing Electrostatic Discharge Damage” on page 567, to avoid damaging equipment. Static discharge can cause components to fail completely or intermittently over time.
- Install blank PIM panels in empty slots, to prevent any interruption or reduction in the flow of air across internal components.

Desktop and Wall Mounting Requirements

The J2300 Services Router can be installed on a desktop or wall. When choosing a location, allow at least 6 in. (15.2 cm) of clearance between the front and back of the chassis and adjacent equipment or walls.

If you are mounting the J2300 router on a wall, use wall screws or wall anchors capable of supporting the full weight of the chassis, up to 12 lb (5.4 kg). If possible, install the wall anchors into wall studs, which provide added support for the chassis.

Rack Requirements

All J-series Services Routers can be installed in a rack. J4300 and J6300 Services Routers must be installed in a rack. Many types of racks are acceptable, including front-mount racks, four-post (telco) racks, and center-mount racks.

The following sections describe rack requirements:

- Rack Size and Strength on page 542
- Spacing of Mounting Holes on page 543
- Connection to Building Structure on page 543

Rack Size and Strength

The Services Router is designed for installation in a rack that complies with either of the following standards:

- A 19-in. rack as defined in Cabinets, Racks, Panels, and Associated Equipment (document number EIA-310-D) published by the Electronics Industry Association (http://www.eia.org)
- A 600-mm rack as defined in the four-part *Equipment Engineering (EE); European telecommunications standard for equipment practice* (document numbers ETS 300 119-1 through 119-4) published by the European Telecommunications Standards Institute (http://www.etsi.org)

  The horizontal spacing between the rails in a rack that complies with this standard is usually wider than the router's mounting ears, which measure 19 in. (48.2 cm) from outer edge to outer edge. Use approved wing devices to narrow the opening between the rails as required.

The rack rails must be spaced widely enough to accommodate the router chassis's external dimensions:

- A J2300 chassis is 1.75 in. (4.4 cm) high, 12.37 in. (31.4 cm) deep, and 17.25 in. (43.8 cm) wide.
- A J4300 or J6300 chassis is 3.5 in. (8.9 cm) high, 19 in. (48.3 cm) deep, and 17 in. (43.2 cm) wide.

The outer edges of the mounting ears extend the width of either chassis to 19 in. (48.2 cm), and the front of the chassis extends approximately 0.5 in. (1.27 cm) beyond the mounting ears. The spacing of rails and adjacent racks must also allow for the clearances around the router and rack. (See "General Site Guidelines" on page 541.)

CAUTION: If you are mounting the router in a cabinet, be sure that ventilation is sufficient to prevent overheating.

If a front-mount rack is used, we recommend supporting the back of the router with a shelf or other structure.

The J2300 chassis height of 1.75 in. (4.4 cm) equals 1 U. The J4300 and J6300 chassis height of 3.5 in. (8.9 cm) equals 2 U. Each *U* is a standard rack unit defined in Cabinets, Racks, Panels, and Associated Equipment (document number EIA-310-D) published by the Electronics Industry Association.

Spacing of Mounting Holes

The mounting holes in the mounting brackets provided with the J2300 Services Router chassis are spaced 1.25 in. (3.2 cm) apart, measured from the center of each hole.

The mounting holes in the mounting brackets attached to the J4300 and J6300 chassis are spaced in two groups. The space between the holes in each group is 0.6 in. (1.5 cm) apart, measured from the center of each hole. The space between the two groups is 1.75 in. (4.4 cm) apart, measured from the center of the lower hole in the top group to the upper hole in the bottom group.

Connection to Building Structure

Always secure the rack to the structure of the building. If your geographical area is subject to earthquakes, bolt the rack to the floor. For maximum stability, also secure the rack to ceiling brackets. For more information, see "Rack-Mounting Requirements and Warnings" on page 578.

Router Environmental Tolerances

Table 200 specifies the environmental conditions required for normal Services Router operation. In addition, the site must be as dust-free as possible. Dust can clog air intake vents, reducing cooling system efficiency. Check vents frequently, cleaning them as necessary.

Table 200: Router Environmental Tolerances

Description	Value
Altitude	No performance degradation to 10,000 ft (3048 m)
Relative humidity	Normal operation ensured in relative humidity range of 5% to 90%, noncondensing
Temperature	Normal operation ensured in temperature range of 32°F (0°C) to 104°F (40°C) Non-operating storage temperature in shipping carton: –40°F (–40°C) to 158°F (70°C)
Seismic	Designed to meet Telcordia Technologies Zone 4 earthquake requirements
Maximum thermal output	J2300: 1638 BTU/hour (480 W) J4300: 2457 BTU/hour (720 W) J6300: 2457 BTU/hour (720 W)

Fire Safety Requirements

In the event of a fire emergency involving Services Routers and other network equipment, the safety of people is the primary concern. Establish procedures for protecting people in the event of a fire emergency, provide safety training, and properly provision fire-control equipment and fire extinguishers.

In addition, establish procedures to protect your equipment in the event of a fire emergency. Juniper Networks products must be installed in an environment suitable for electronic equipment. We recommend that fire suppression equipment be available in the event of a fire in the vicinity of the equipment, and that all local fire, safety, and electrical codes and ordinances be observed when you are installing and operating your equipment.

Fire Suppression

In the event of an electrical hazard or an electrical fire, first unplug the power cord. (For shutdown instructions, see “Powering a Services Router On and Off” on page 44.)

Then, use a Type C fire extinguisher, which uses noncorrosive fire retardants, to extinguish the fire. For more information about fire extinguishers, see “Fire Suppression Equipment” on page 544.

Fire Suppression Equipment

Type C fire extinguishers, which use noncorrosive fire retardants such as carbon dioxide (CO_2) and Halotron, are most effective for suppressing electrical fires. Type C fire extinguishers displace the oxygen from the point of combustion to

eliminate the fire. For extinguishing fire on or around equipment that draws air from the environment for cooling, use this type of inert oxygen displacement extinguisher instead of an extinguisher that leave residues on equipment.

Do not use multipurpose Type ABC chemical fire extinguishers (dry chemical fire extinguishers) near Juniper Networks equipment. The primary ingredient in these fire extinguishers is monoammonium phosphate, which is very sticky and difficult to clean. In addition, in minute amounts of moisture, monoammonium phosphate can become highly corrosive and corrodes most metals.

NOTE: To keep warranties effective, do not use a dry chemical fire extinguisher to control a fire at or near a Juniper Networks router. If a dry chemical fire extinguisher is used, the unit is no longer eligible for coverage under a service agreement.

Any equipment in a room in which a chemical fire extinguisher has been discharged is subject to premature failure and unreliable operation. The equipment is considered to be irreparably damaged.

We recommend that you dispose of any irreparably damaged equipment in an environmentally responsible manner.

Power Guidelines, Requirements, and Specifications

All Services Routers use AC power. For information about each router's power system, see "J2300 Power System" on page 15, "J4300 Power System" on page 26, and "J6300 Power System" on page 26.

For site wiring and power system guidelines, requirements, and specifications, see the following sections:

- Site Electrical Wiring Guidelines on page 545
- Router Power Requirements on page 546
- AC Power, Connection, and Power Cord Specifications on page 547

Site Electrical Wiring Guidelines

When planning the electrical wiring at your site, consider the factors discussed in the following sections.

Signaling Limitations

Improperly installed wires can emit radio interference. In addition, the potential for damage from lightning strikes increases if wires exceed recommended distances, or if wires pass between buildings. The electromagnetic pulse (EMP) caused by lightning can damage unshielded conductors and destroy electronic

devices. If your site has previously experienced such problems, you might want to consult experts in electrical surge suppression and shielding.

Radio Frequency Interference

You can reduce or eliminate the emission of radio frequency interference (RFI) from your site wiring by using twisted-pair cable with a good distribution of grounding conductors. If you must exceed the recommended distances, use a high-quality twisted-pair cable with one ground conductor for each data signal when applicable.

Electromagnetic Compatibility

If your site is susceptible to problems with electromagnetic compatibility (EMC), particularly from lightning or radio transmitters, you might want to seek expert advice. Strong sources of electromagnetic interference (EMI) can destroy the signal drivers and receivers in the router and conduct power surges over the lines into the equipment, resulting in an electrical hazard. It is particularly important to provide a properly grounded and shielded environment and to use electrical surge-suppression devices.

CAUTION: To comply with intrabuilding lightning/surge requirements, intrabuilding wiring must be shielded, and the shield for the wiring must be grounded at both ends.

Router Power Requirements

Table 201 lists the power system electrical specifications for the J2300 Services Router.

Table 201: Power System Electrical Specifications for the J2300 Services Router

Item	Specification
AC input voltage	Operating range: 100 to 240 VAC
AC input line frequency	47 to 63 Hz
AC system current rating	4 to 2 A

Table 202 lists the power system electrical specifications for the J4300 Services Router.

Table 202: Power System Electrical Specifications for the J4300 Services Router

Item	Specification
AC input voltage	Operating range: 100 to 240 VAC
AC input line frequency	47 to 63 Hz
AC system current rating	6 to 3 A

Table 203 lists the power system electrical specifications for the J6300 Services Router.

Table 203: Power System Electrical Specifications J6300 Services Router

Item	Specification
AC input voltage	Operating range: 100 to 240 VAC
AC input line frequency	47 to 63 Hz
AC system current rating	6 to 3 A

AC Power, Connection, and Power Cord Specifications

Detachable AC power cords, each 2.5 m (approximately 8 ft) long, are supplied with the Services Router. The appliance coupler at the female end of the cord inserts into the appliance inlet on the faceplate of the AC power supply. The coupler is type C19 as described by International Electrotechnical Commission (IEC) standard 60320. The plug at the male end of the power cord fits into the power source receptacle that is standard for your geographical location.

NOTE: In North America, AC power cords must not exceed 4.5 m (approximately 14.75 ft) in. length, to comply with National Electrical Code (NEC) Sections 400-8 (NFPA 75, 5-2.2) and 210-52, and Canadian Electrical Code (CEC) Section 4-010(3). The cords supplied with the router are in compliance.

Table 204 lists power cord specifications and Figure 110 illustrates the plug on the AC power cord provided for each country or region.

Table 204: AC Power Cord Specifications

Country	Electrical Specifications	Plug Standards
Australia	250 VAC, 10 A, 50 Hz	AS/NZ 3112–1993
China	250 VAC, 10 A, 50 Hz	GB2099.1 1996 and GB1002 1996 (CH1-10P)
Europe (except Italy and United Kingdom)	250 VAC, 10 A, 50 Hz	CEE (7) VII

Country	Electrical Specifications	Plug Standards
Italy	250 VAC, 10 A, 50 Hz	CEI 23–16/VII
Japan	125 VAC, 12 A, 50 Hz or 60 Hz	JIS 8303
North America	125 VAC, 10 A, 60 Hz	NEMA 5-15
United Kingdom	250 VAC, 10 A, 50 Hz	BS 1363A

Figure 110: AC Plug Types

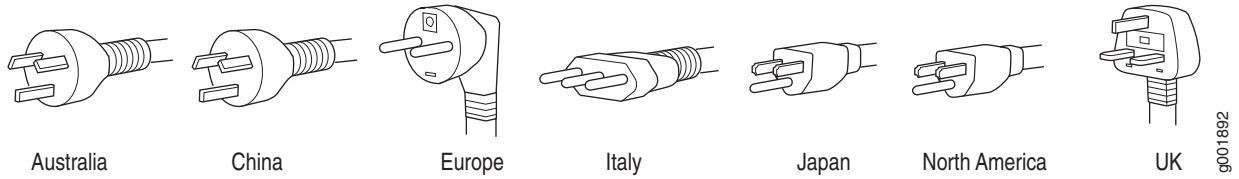

NOTE: Power cords and cables must not block access to router components or drape where people might trip on them.

For information about the AC power supply, see “J2300 Power System” on page 15, “J4300 Power System” on page 26, or “J6300 Power System” on page 26.

To connect the power cord during initial installation, see “Connecting Power to the Services Router” on page 43.

To replace the AC power cord, see “Replacing a Power Supply Cord in a J2300 or J4300 Router” on page 532 or “Replacing a Power Supply Cord in a J6300 Router” on page 536.

Network Cable Specifications

The Services Router supports interfaces that use various kinds of network cable. For information about the type of cable used by each interface, see “Network Cable Specifications and Connector Pinouts” on page 551.

Site Preparation Checklist

The checklist in Table 205 summarizes the tasks you need to perform when preparing a site for Services Router installation.

Table 205: Site Preparation Checklist

Item or Task	Performed By	Date	Notes
Verify that environmental factors such as temperature and humidity do not exceed router tolerances.			
Measure the distances between external power sources and the router installation site.			
Select the type of rack.			
Plan the rack location, including required space clearances.			
Secure the rack to the floor and the building structure.			
Acquire appropriate cables and connectors.			

Chapter 27

Network Cable Specifications and Connector Pinouts

The network interfaces supported on the router accept different kinds of network cable.

Serial PIM Cable Specifications

The 2-port serial PIM uses the cables and connectors summarized in Table 206. Pinouts are detailed in Table 207 through Table 216.

Table 206: 2-Port Serial PIM Cables and Connectors

Name	Connector	Connector Hardware	End-to-End Conductors	Pinouts
RS-232 DTE	DB-25 male	4-40 threaded jackscrews	13	Table 207
RS-232 DCE	DB-25 female	4-40 threaded jacknuts	13	Table 208
RS-422/449 (EIA-449) DTE	DC-37 (DB-37) male	4-40 threaded jackscrews	25	Table 209
RS-422/449 (EIA-449) DCE	DC-37 (DB-37) female	4-40 threaded jacknuts	25	Table 210
EIA-530A DTE	DB-25 male	4-40 threaded jackscrews	23	Table 211
EIA-530A DCE	DB-25 female	4-40 threaded jacknuts	22	Table 212
V.35 DTE	M/34 male	Standard (Normally included with M/34 connector shell)	18	Table 213

Name	Connector	Connector Hardware	End-to-End Conductors	Pinouts
V.35 DCE	M/34 female	Standard (Normally included with M/34 connector shell)	18	Table 214
X.21 DTE	DB-15 male	M3 threaded jackscrews	13	Table 215
X.21 DCE	DB-15 female	M3 threaded jacknuts	13	Table 216

RS-232 DTE Cable Pinout

Table 207: RS-232 DTE Cable Pinout

LFH-60 Pin	DB-25 Pin	LFH-60 Pairing	Description
15	1	–	Frame Ground
60	2	–	Transmit Data
1	3	–	Receive Data
48	4	–	Request to Send
37	5	–	Clear to Send
9	6	–	Data Set Ready
57	7	–	Signal Ground
13	8	–	Data Carrier Detect
56	15	–	Transmit Clock
5	17	–	Receive Clock
41	18	–	Local Loopback
33	20	–	Data Terminal Ready
52	24	–	Terminal Clock
22 to 21	–	–	–
18 to 17	–	–	–

RS-232 DCE Cable Pinout

Table 208: RS-232 DCE Cable Pinout

LFH-60 Pin	DB-25 Pin	LFH-60 Pairing	Description
15	1	–	Frame Ground
1	2	–	Transmit Data
60	3	–	Receive Data
37	4	–	Request to Send

LFH-60 Pin	DB-25 Pin	LFH-60 Pairing	Description
48	5	–	Clear to Send
33	6	–	Data Set Ready
57	7	–	Signal Ground
13	8	–	Data Carrier Detect
56	15	–	Transmit Clock
52	17	–	Receive Clock
45	18	–	Local Loopback
9	20	–	Data Terminal Ready
5	24	–	Terminal Clock
22 to 21	–	–	–

RS-422/449 (EIA-449) DTE Cable Pinout

Table 209: RS-422/449 (EIA-449) DTE Cable Pinout

LFH-60 Pin	DC-37 (DB-37) Pin	LFH-60 Pairing	Description
15	1	–	Shield Ground
60	4	59	Send Data (A)
56	5	55	Send Timing (A)
1	6	2	Receive Data (A)
48	7	47	Request to Send (A)
5	8	6	Receive Timing (A)
37	9	38	Clear to Send (A)
41	10	–	Local Loopback
9	11	10	Data Mode (A)
33	12	34	Terminal Ready (A)
13	13	14	Receive Ready (A)
52	17	51	Terminal Timing (A)
36	19	–	Signal Ground
4	20	–	Receive Common
59	22	60	Send Data (B)
55	23	56	Send Timing (B)
2	24	1	Receive Data (B)
47	25	48	Request to Send (B)
6	26	5	Receive Timing (B)
38	27	37	Clear to Send (B)

LFH-60 Pin	DC-37 (DB-37) Pin	LFH-60 Pairing	Description
10	29	9	Data Mode (B)
34	30	33	Terminal Ready (B)
14	31	13	Receiver Ready (B)
51	35	52	Terminal Timing (B)
57	37	–	Send Common
26 to 25	–	–	–
18 to 17	–	–	–

RS-422/449 (EIA-449) DCE Cable Pinout

Table 210: RS-422/449 (EIA-449) DCE Cable Pinout

LFH-60 Pin	DC-37 (DB-37) Pin	LFH-60 Pairing	Description
15	1	–	Shield Ground
1	4	2	Send Data (A)
56	5	55	Send Timing (A)
60	6	59	Receive Data (A)
37	7	38	Request to Send (A)
52	8	51	Receive Timing (A)
48	9	47	Clear to Send (A)
45	10	–	Local Loopback
33	11	34	Data Mode (A)
9	12	10	Terminal Ready (A)
13	13	14	Receive Ready (A)
5	17	6	Terminal Timing (A)
36	19	–	Signal Ground
4	20	–	Receive Common
2	22	1	Send Data (B)
55	23	56	Send Timing (B)
59	24	60	Receive Data (B)
38	25	37	Request to Send (B)
51	26	52	Receive Timing (B)
47	27	48	Clear to Send (B)
34	29	33	Data Mode (B)
10	30	9	Terminal Ready (B)
14	31	13	Receiver Ready (B)

LFH-60 Pin	DC-37 (DB-37) Pin	LFH-60 Pairing	Description
6	35	5	Terminal Timing (B)
57	37	–	Send Common
26 to 25	–	–	–

EIA-530A DTE Cable Pinout

Table 211: EIA-530A DTE Cable Pinout

LFH-60 Pin	DB-25 Pin	LFH-60 Pairing	Description
15	1	–	Shield Ground
60	2	59	Transmit Data (A)
1	3	2	Receive Data (A)
48	4	47	Request to Send (A)
37	5	38	Clear to Send (A)
9	6	–	Data Set Ready (A)
57	7	–	Signal Ground
13	8	14	Received Line Signal Detector (A)
6	9	5	Receive Clock (B)
14	10	13	Received Line Signal Detector (B)
51	11	52	Terminal Timing (B)
55	12	56	Transmit Clock (B)
38	13	37	Clear to Send (B)
59	14	60	Transmit Data (B)
56	15	55	Transmit Clock (A)
2	16	1	Receive Data (B)
5	17	6	Receive Clock (A)
41	18	–	Local Loopback
47	19	48	Request to Send (B)
33	20	–	Data Terminal Ready (A)
4	23	–	Signal Ground
52	24	51	Terminal Timing (A)
26 to 25	–	–	–
30 to 29	–	–	–
18 to 17	–	–	–

EIA-530A DCE Cable Pinout

Table 212: EIA-530A DCE Cable Pinout

LFH-60 Pin	DB-25 Pin	LFH-60 Pairing	Description
15	1	–	Shield Ground
1	2	2	Transmit Data (A)
60	3	59	Receive Data (A)
37	4	38	Request to Send (A)
48	5	47	Clear to Send (A)
33	6	–	Data Set Ready (A)
57	7	–	Signal Ground
13	8	14	Received Line Signal Detector (A)
51	9	52	Receive Clock (B)
14	10	13	Received Line Signal Detector (B)
6	11	5	Terminal Timing (B)
55	12	56	Transmit Clock (B)
47	13	48	Clear to Send (B)
2	14	1	Transmit Data (B)
56	15	55	Transmit Clock (A)
59	16	60	Receive Data (B)
52	17	51	Receive Clock (A)
45	18	–	Local Loopback
38	19	37	Request to Send (B)
9	20	–	Data Terminal Ready (A)
4	23	–	Signal Ground
5	24	6	Terminal Timing (A)
26 to 25	–	–	–
30 to 29	–	–	–

V.35 DTE Cable Pinout

Table 213: V.35 DTE Cable Pinout

LFH-60 Pin	M/34 Pin	LFH-60 Pairing	Description
15	A	–	Frame Ground
57	B	–	Signal Ground

LFH-60 Pin	M/34 Pin	LFH-60 Pairing	Description
48	C	–	Request to Send
37	D	–	Clear to Send
9	E	–	Data Set Ready
13	F	–	Received Line Signal Detector
33	H	–	Data Terminal Ready
41	K	–	Test Mode
60	P	59	Transmit Data (A)
1	R	2	Receive Data (A)
59	S	60	Transmit Data (B)
2	T	1	Receive Data (B)
52	U	51	Terminal Timing (A)
5	V	6	Receive Timing (A)
51	W	52	Terminal Timing (B)
6	X	5	Receive Timing (B)
56	Y	55	Transmit Timing (A)
55	AA	56	Transmit Timing (B)
22 to 21	–	–	–
26 to 25	–	–	–
18 to 17	–	–	–

V.35 DCE Cable Pinout

Table 214: V.35 DCE Cable Pinout

LFH-60 Pin	M/34 Pin	LFH-60 Pairing	Description
15	A	–	Frame Ground
57	B	–	Signal Ground
37	C	–	Request to Send
48	D	–	Clear to Send
33	E	–	Data Set Ready
13	F	–	Received Line Signal Detector
9	H	–	Data Terminal Ready
45	K	–	Test Mode
1	P	2	Transmit Data (A)
60	R	59	Receive Data (A)
2	S	1	Transmit Data (B)

LFH-60 Pin	M/34 Pin	LFH-60 Pairing	Description
59	T	60	Receive Data (B)
5	U	6	Terminal Timing (A)
52	V	51	Receive Timing (A)
6	W	5	Terminal Timing (B)
51	X	52	Receive Timing (B)
56	Y	55	Transmit Timing (A)
55	AA	56	Transmit Timing (B)
22 to 21	–	–	–
26 to 25	–	–	–

X.21 DTE Cable Pinout

Table 215: X.21 DTE Cable Pinout

LFH-60 Pin	DB-15 Pin	LFH-60 Pairing	Description
15	1	–	Shield Ground
60	2	59	Transmit Data (A)
48	3	47	Control (A)
1	4	2	Receive (A)
37	5	38	Indicate (A)
5	6	6	Signal Element Timing (A)
57	8	–	Signal Ground
59	9	60	Transmit Data (B)
47	10	48	Control (B)
2	11	1	Receive (B)
38	12	37	Indicate (B)
6	13	5	Signal Element Timing (B)
30 to 29	–	–	–
18 to 17	–	–	–

X.21 DCE Cable Pinout

Table 216: X.21 DCE Cable Pinout

LFH-60 Pin	DB-15 Pin	LFH-60 Pairing	Description
15	1	–	Shield Ground
1	2	2	Transmit Data (A)
37	3	38	Control (A)
60	4	59	Receive (A)
48	5	47	Indicate (A)
52	6	51	Signal Element Timing (A)
57	8	–	Signal Ground
2	9	1	Transmit Data (B)
38	10	37	Control (B)
59	11	60	Receive (B)
47	12	48	Indicate (B)
51	13	52	Signal Element Timing (B)
30 to 29	–	–	–

RJ-45 Connector Pinouts for the Routing Engine (Ethernet) Port

Table 217 describes the RJ-45 connector pinout information.

Table 217: RJ-45 Connector Pinout

Pin	Signal
1	TX +
2	TX-
3	RX +
4	Termination network
5	Termination network
6	RX-
7	Termination network
8	Termination network

DB-9 Connector Pinouts for the Console Port

Table 218 describes the DB-9 connector pinouts.

Table 218: DB-9 Connector Pinout

Pin	Signal	Direction	Description
1	DCD	<–	Carrier Detect
2	RxD	<–	Receive Data
3	TxD	–>	Transmit Data
4	DTR	–>	Data Terminal Ready
5	Ground	—	Signal Ground
6	DSR	<–	Data Set Ready
7	RTS	–>	Request To Send
8	CTS	<–	Clear To Send
9	RING	<–	Ring Indicator

E1 and T1 RJ-48 Cable Pinouts

The E1 and T1 PIMs use an RJ-48 cable, which is not supplied with the PIM.

CAUTION: To maintain agency approvals, use only a properly constructed, shielded cable.

Table 219, Table 220, Table 221, and Table 222 describe the RJ-48 connector pinouts.

Table 219: RJ-48 Connector to RJ-48 Connector (Straight) Pinout

RJ-48 Pin (on T1/E1 PIM) (Data Numbering Form)	RJ-48 Pin (Data Numbering Form)	Signal
1	1	RX, Ring, –
2	2	RX, Tip, +
4	4	TX, Ring, –
5	5	TX, Tip, +
3	3	Shield/Return/Ground
6	6	Shield/Return/Ground
7	No connect	No connect
8	No connect	No connect

Table 220: RJ-48 Connector to RJ-48 Connector (Crossover) Pinout

RJ-48 Pin (on T1/E1 PIM) (Data numbering form)	RJ-48 Pin (Data numbering form)	Signal
1	4	RX/Ring/– <––> TX/Ring/–
2	5	RX/Tip/+ <––> TX/Tip/+
4	1	TX/Ring/– <––> RX/Ring/–
5	2	TX/Tip/+ <––> RX/Tip/+
3	3	Shield/Return/Ground
6	6	Shield/Return/Ground
7	No connect	No connect
8	No connect	No connect

Table 221: RJ-48 Connector to DB-15 Connector (Straight) Pinout

RJ-48 Pin (on T1/E1 PIM) (Data numbering form)	DB-15 Pin (Data numbering form)	Signal
1	11	RX/Ring/– <––> RX/Ring/–
2	3	RX/Tip/+ <––> RX/Tip/+
4	9	TX/Ring/– <––> TX/Ring/–
5	1	TX/Tip/+ <––> TX/Tip/+
3	4	Shield/Return/Ground
6	2	Shield/Return/Ground
7	No connect	No connect
8	No connect	No connect
9	No connect	No connect
10	No connect	No connect
11	No connect	No connect
12	No connect	No connect
13	No connect	No connect
14	No connect	No connect
15	No connect	No connect

Table 222: RJ-48 Connector to DB-15 Connector (Crossover) Pinout

RJ-48 Pin (on T1/E1 PIM) (Data numbering form)	DB-15 Pin (Data numbering form)	Signal
1	9	RX/Ring/– <––> TX/Ring/–
2	1	RX/Tip/+ <––> TX/Tip/+
4	11	TX/Ring/– <––> RX/Ring/–
5	3	TX/Tip/+ <––> RX/Tip/+
3	4	Shield/Return/Ground
6	2	Shield/Return/Ground
7	No connect	No connect
8	No connect	No connect
9	No connect	No connect
10	No connect	No connect
11	No connect	No connect
12	No connect	No connect
13	No connect	No connect
14	No connect	No connect
15	No connect	No connect

Chapter 28

Safety and Regulatory Compliance Information

To install and use the Services Router safely, follow proper safety procedures. This chapter discusses the following safety and regulatory compliance information:

Definition of Safety Warning Levels

This manual uses the following three levels of safety warnings:

NOTE: You might find this information helpful in a particular situation, or might otherwise overlook it.

CAUTION: You need to observe the specified guidelines to avoid minor injury or discomfort to you, or severe damage to the Services Router.

WARNING: This symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents.

WARNING: Waarschuwing Dit waarschuwingssymbool betekent gevaar. U verkeert in een situatie die lichamelijk letsel kan veroorzaken. Voordat u aan enige apparatuur gaat werken, dient u zich bewust te zijn van de bij elektrische schakelingen betrokken risico's en dient u op de hoogte te zijn van standaard maatregelen om ongelukken te voorkomen.

WARNING: Varoitus Tämä varoitusmerkki merkitsee vaaraa. Olet tilanteessa, joka voi johtaa ruumiinvammaan. Ennen kuin työskentelet minkään laitteiston parissa, ota selvää sähkökytkentöihin liittyvistä vaaroista ja tavanomaisista onnettomuuksien ehkäisykeinoista.

WARNING: Attention Ce symbole d'avertissement indique un danger. Vous vous trouvez dans une situation pouvant causer des blessures ou des dommages corporels. Avant de travailler sur un équipement, soyez conscient des dangers posés par les circuits électriques et familiarisez-vous avec les procédures couramment utilisées pour éviter les accidents.

WARNING: Warnung Dieses Warnsymbol bedeutet Gefahr. Sie befinden sich in einer Situation, die zu einer Körperverletzung führen könnte. Bevor Sie mit der Arbeit an irgendeinem Gerät beginnen, seien Sie sich der mit elektrischen Stromkreisen verbundenen Gefahren und der Standardpraktiken zur Vermeidung von Unfällen bewußt.

WARNING: Avvertenza Questo simbolo di avvertenza indica un pericolo. La situazione potrebbe causare infortuni alle persone. Prima di lavorare su qualsiasi apparecchiatura, occorre conoscere i pericoli relativi ai circuiti elettrici ed essere al corrente delle pratiche standard per la prevenzione di incidenti.

WARNING: Advarsel Dette varselsymbolet betyr fare. Du befinner deg i en situasjon som kan føre til personskade. Før du utfører arbeid på utstyr, må du vare oppmerksom på de faremomentene som elektriske kretser innebærer, samt gjøre deg kjent med vanlig praksis når det gjelder å unngå ulykker.

WARNING: Aviso Este símbolo de aviso indica perigo. Encontra-se numa situação que lhe poderá causar danos físicos. Antes de começar a trabalhar com qualquer equipamento, familiarize-se com os perigos relacionados com circuitos eléctricos, e com quaisquer práticas comuns que possam prevenir possíveis acidentes.

WARNING: ¡Atención! Este símbolo de aviso significa peligro. Existe riesgo para su integridad física. Antes de manipular cualquier equipo, considerar los riesgos que entraña la corriente eléctrica y familiarizarse con los procedimientos estándar de prevención de accidentes.

WARNING: Varning! Denna varningssymbol signalerar fara. Du befinner dig i en situation som kan leda till personskada. Innan du utför arbete på någon utrustning måste du vara medveten om farorna med elkretsar och känna till vanligt förfarande för att förebygga skador.

Safety Guidelines and Warnings

This section lists safety guidelines and warnings for installing, operating, and maintaining the Services Router.

General Safety Guidelines and Warnings

The following guidelines help ensure your safety and protect the Services Router from damage. The list of guidelines might not address all

potentially hazardous situations in your working environment, so be alert and exercise good judgment at all times.

- Perform only the procedures explicitly described in this manual. Make sure that only authorized service personnel perform other system services.
- Keep the area around the chassis clear and free from dust before, during, and after installation.
- Keep tools away from areas where people could trip over them while walking.
- Do not wear loose clothing or jewelry, such as rings, bracelets, or chains, which could become caught in the chassis.
- Wear safety glasses if you are working under any conditions that could be hazardous to your eyes.
- Do not perform any actions that create a potential hazard to people or make the equipment unsafe.
- Never attempt to lift an object that is too heavy for one person to handle.
- Never install or manipulate wiring during electrical storms.
- Never install electrical jacks in wet locations unless the jacks are specifically designed for wet environments.
- Operate the Services Router only when it is properly grounded.
- The separate protective earthing terminal provided on this product shall be permanently connected to earth.
- Replace fuses only with fuses of the same type and rating.
- Do not open or remove chassis covers or sheet metal parts unless instructions are provided in this manual. Such an action could cause severe electrical shock.
- Do not push or force any objects through any opening in the chassis frame. Such an action could result in electrical shock or fire.
- Avoid spilling liquid onto the Services Router chassis or onto any Services Router component. Such an action could cause electrical shock or damage the Services Router.
- Avoid touching uninsulated electrical wires or terminals that have not been disconnected from their power source. Such an action could cause electrical shock.

In addition, observe the warnings and guidelines in the following sections.

Qualified Personnel Warning

WARNING: Only trained and qualified personnel should install or replace the Services Router.

Waarschuwing Installatie en reparaties mogen uitsluitend door getraind en bevoegd personeel uitgevoerd worden.

Varoitus Ainoastaan koulutettu ja pätevä henkilökunta saa asentaa tai vaihtaa tämän laitteen.

Attention Tout installation ou remplacement de l'appareil doit être réalisé par du personnel qualifié et compétent.

Warnung Gerät nur von geschultem, qualifiziertem Personal installieren oder auswechseln lassen.

WARNING: Avvertenza Solo personale addestrato e qualificato deve essere autorizzato ad installare o sostituire questo apparecchio.

Advarsel Kun kvalifisert personell med riktig opplæring bør montere eller bytte ut dette utstyret.

Aviso Este equipamento deverá ser instalado ou substituído apenas por pessoal devidamente treinado e qualificado.

¡Atención! Estos equipos deben ser instalados y reemplazados exclusivamente por personal técnico adecuadamente preparado y capacitado.

Varning! Denna utrustning ska endast installeras och bytas ut av utbildad och kvalificerad personal.

Preventing Electrostatic Discharge Damage

Many Services Router hardware components are sensitive to damage from static electricity. Some components can be impaired by voltages as low as 30 V. You can easily generate potentially damaging static voltages whenever you handle plastic or foam packing material or if you move components across plastic or carpets. Observe the following guidelines to minimize the potential for electrostatic discharge (ESD) damage, which can cause intermittent or complete component failures:

- Always use an ESD wrist strap or ankle strap, and make sure that it is in direct contact with your skin.

CAUTION: For safety, periodically check the resistance value of the ESD strap. The measurement should be in the range of 1 to 10 Mohms.

- When handling any component that is removed from the chassis, make sure the equipment end of your ESD strap is attached to one of the electrostatic discharge points on the chassis, which are shown in Figure 1 and Figure 2 for the J2300 chassis and in Figure 6 and Figure 7 for the J4300 chassis and J6300 chassis.
- Avoid contact between the component and your clothing. ESD voltages emitted from clothing can still damage components.
- When removing or installing a component, always place it component-side up on an antistatic surface, in an antistatic card rack, or in an electrostatic bag (see Figure 111). If you are returning a component, place it in an electrostatic bag before packing it.

Figure 111: Place a Component into an Electrostatic Bag

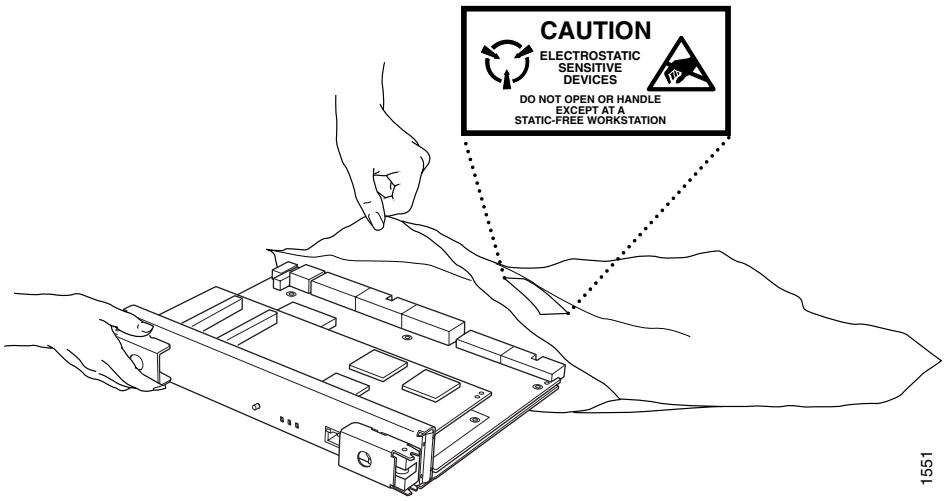

Electrical Safety Guidelines and Warnings

When working on equipment powered by electricity, follow the guidelines described in the following sections.

General Electrical Safety Guidelines

- Install the Services Router in compliance with the following local, national, or international electrical codes:
  - United States—National Fire Protection Association (NFPA 70), United States National Electrical Code.
  - Canada—Canadian Electrical Code, Part 1, CSA C22.1.
  - Other countries—International Electromechanical Commission (IEC) 60364, Part 1 through Part 7.
  - Evaluated to the TN power system.
- Locate the emergency power-off switch for the room in which you are working so that if an electrical accident occurs, you can quickly turn off the power.
- Do not work alone if potentially hazardous conditions exist anywhere in your workspace.
- Never assume that power is disconnected from a circuit. Always check the circuit before starting to work.
- Carefully look for possible hazards in your work area, such as moist floors, ungrounded power extension cords, and missing safety grounds.
- Operate the Services Router within marked electrical ratings and product usage instructions.
- For the Services Router and peripheral equipment to function safely and correctly, use the cables and connectors specified for the attached peripheral equipment, and make certain they are in good condition.

Many Services Router components can be removed and replaced without powering down or disconnecting power to the Services Router, as detailed in elsewhere in this manual. Never install equipment if it appears damaged.

AC Power Electrical Safety Guidelines

The following electrical safety guidelines apply to AC-powered routers:

- AC-powered routers are shipped with a three-wire electrical cord with a grounding-type plug that fits only a grounding-type power outlet. Do not

circumvent this safety feature. Equipment grounding should comply with local and national electrical codes.

- You must provide an external circuit breaker rated minimum 15 A in the building installation.
- The power cord serves as the main disconnecting device. The socket outlet must be near the router and be easily accessible.
- The cores in the mains lead are colored in accordance with the following code:
  - Green and yellow—Earth
  - Blue—Neutral
  - Brown—Live
- When a router is equipped with two AC power supplies, both power cords (one for each power supply) must be unplugged to completely disconnect power to the router.
- Note the following warnings printed on the AC power supply faceplate:
  - To completely de-energize the system disconnect maximum of 2 power cordsets.
  - Apparaten skall anslutas till jordat uttag när den ansluts till ett nätverk. [Swedish]

Grounded Equipment Warning

WARNING: The router is intended to be grounded. Ensure that the router is connected to earth ground during normal use.

Waarschuwing Deze apparatuur hoort geaard te worden Zorg dat de host-computer tijdens normaal gebruik met aarde is verbonden.

Varoitus Tämä laitteisto on tarkoitettu maadoitettavaksi. Varmista, että isäntälaite on yhdistetty maahan normaalikäytön aikana.

Attention Cet équipement doit être relié à la terre. S'assurer que l'appareil hôte est relié à la terre lors de l'utilisation normale.

Warnung Dieses Gerät muß geerdet werden. Stellen Sie sicher, daß das Host-Gerät während des normalen Betriebs an Erde gelegt ist.

WARNING: Avvertenza Questa apparecchiatura deve essere collegata a massa. Accertarsi che il dispositivo host sia collegato alla massa di terra durante il normale utilizzo.

Advarsel Dette utstyret skal jordes. Forviss deg om vertsterminalen er jordet ved normalt bruk.

Aviso Este equipamento deverá estar ligado à terra. Certifique-se que o host se encontra ligado à terra durante a sua utilização normal.

¡Atención! Este equipo debe conectarse a tierra. Asegurarse de que el equipo principal esté conectado a tierra durante el uso normal.

Varning! Denna utrustning är avsedd att jordas. Se till att värdenheten är jordad vid normal användning.

Warning Statement for Norway and Sweden

WARNING: The equipment must be connected to an earthed mains socket-outlet.

Advarsel Apparatet skal kobles til en jordet stikkontakt.

Varning! Apparaten skall anslutas till jordat nätuttag.

In Case of Electrical Accident

If an electrical accident results in an injury, take the following actions in this order:

1. Use caution. Be aware of potentially hazardous conditions that could cause further injury.
2. Disconnect power from the Services Router.
3. If possible, send another person to get medical aid. Otherwise, assess the condition of the victim, then call for help.

Multiple Power Supplies Disconnection Warning

WARNING: The J6300 Services Router has more than one power supply connection. All connections must be removed completely to remove power from the unit completely.

WARNING: Waarschuwing Deze J6300 eenheid heeft meer dan één stroomtoevoerverbinding; alle verbindingen moeten volledig worden verwijderd om de stroom van deze eenheid volledig te verwijderen.

WARNING: Varoitus Tässä laitteessa on useampia virtalähdekytkentöjä. Kaikki kytkennät on irrotettava kokonaan, jotta virta poistettaisiin täysin laitteesta.

WARNING: Attention Cette J6300 unité est équipée de plusieurs raccordements d'alimentation. Pour supprimer tout courant électrique de l'unité, tous les cordons d'alimentation doivent être débranchés.

WARNING: Warnung Diese J6300 Einheit verfügt über mehr als einen Stromanschluß; um Strom gänzlich von der Einheit fernzuhalten, müssen alle Stromzufuhren abgetrennt sein.

WARNING: Avvertenza Questa J6300 unità ha più di una connessione per alimentatore elettrico; tutte le connessioni devono essere completamente rimosse per togliere l'elettricità dall'unità.

WARNING: Advarsel Denne J6300 enheten har mer enn én strømtilkobling. Alle tilkoblinger må kobles helt fra for å eliminere strøm fra enheten.

WARNING: Aviso Este J6300 dispositivo possui mais do que uma conexão de fonte de alimentação de energia; para poder remover a fonte de alimentação de energia, deverão ser desconectadas todas as conexões existentes.

WARNING: ¡Atención! Esta J6300 unidad tiene más de una conexión de suministros de alimentación; para eliminar la alimentación por completo, deben desconectarse completamente todas las conexiones.

WARNING: Varning! Denna J6300 enhet har mer än en strömförsörjningsanslutning; alla anslutningar måste vara helt avlägsnade innan strömtillförseln till enheten är fullständigt bruten.

Power Disconnection Warning

WARNING: Before working on the router or near power supplies, unplug the power cord from an AC router.

WARNING: Waarschuwing Voordat u aan een frame of in de nabijheid van voedingen werkt, dient u bij wisselstroom toestellen de stekker van het netsnoer uit het stopcontact te halen.

WARNING: Varoitus Kytke irti vaihtovirtalaitteiden virtajohto, ennen kuin teet mitään asennuspohjalle tai työskentelet virtalähteiden läheisyydessä.

WARNING: Attention Avant de travailler sur un châssis ou à proximité d'une alimentation électrique, débrancher le cordon d'alimentation des unités en courant alternatif.

WARNING: Warnung Bevor Sie an einem Chassis oder in der Nähe von Netzgeräten arbeiten, ziehen Sie bei Wechselstromeinheiten das Netzkabel ab bzw.

WARNING: Avvertenza Prima di lavorare su un telaio o intorno ad alimentatori, scollegare il cavo di alimentazione sulle unità CA.

WARNING: Advarsel Før det utføres arbeid på kabinettet eller det arbeides i nærheten av strømforsyningsenheter, skal strømledningen trekkes ut på vekselstrømsenheter.

WARNING: Aviso Antes de trabalhar num chassis, ou antes de trabalhar perto de unidades de fornecimento de energia, desligue o cabo de alimentação nas unidades de corrente alternada.

WARNING: ¡Atención! Antes de manipular el chasis de un equipo o trabajar cerca de una fuente de alimentación, desenchufar el cable de alimentación en los equipos de corriente alterna (CA).

WARNING: Varning! Innan du arbetar med ett chassi eller nära strömförsörjningsenheter skall du för växelströmsenheter dra ur nätsladden.

TN Power Warning

WARNING: The router is designed to work with a TN power system.

WARNING: Waarschuwing Het apparaat is ontworpen om te functioneren met TN energiesystemen.

WARNING: Varoitus Koje on suunniteltu toimimaan TN-sähkövoimajärjestelmien yhteydessä.

WARNING: Attention Ce dispositif a été conçu pour fonctionner avec des systèmes d'alimentation TN.

WARNING: Warnung Das Gerät ist für die Verwendung mit TN-Stromsystemen ausgelegt.

WARNING: Avvertenza Il dispositivo è stato progettato per l'uso con sistemi di alimentazione TN.

WARNING: Advarsel Utstyret er utfomet til bruk med TN-strømsystemer.

WARNING: Aviso O dispositivo foi criado para operar com sistemas de corrente TN.

WARNING: ¡Atención! El equipo está diseñado para trabajar con sistemas de alimentación tipo TN.

WARNING: Varning! Enheten är konstruerad för användning tillsammans med elkraftssystem av TN-typ.

Telecommunication Line Cord Warning

WARNING: To reduce the risk of fire, use only No. 26 AWG or larger UL-listed or CSA-certified telecommunication line cord.

WARNING: Waarschuwing Om brandgevaar te reduceren, dient slechts telecommunicatielijnsnoer nr. 26 AWG of groter gebruikt te worden.

WARNING: Varoitus Tulipalovaaran vähentämiseksi käytä ainoastaan nro 26 AWG- tai paksumpaa tietoliikennejohdinta.

WARNING: Attention Pour réduire les risques d'incendie, n'utiliser que des cordons de lignes de télécommunications de type AWG nº 26 ou plus larges.

WARNING: Warnung Zur Reduzierung der Feuergefahr eine Fernmeldeleitungsschnur der Größe 26 AWG oder größer verwenden.

WARNING: Avvertenza Per ridurre il rischio di incendio, usare solo un cavo per linea di telecomunicazioni di sezione 0,12 mm2 (26 AWG) o maggiore.

WARNING: Advarsel Bruk kun AWG nr. 26 eller telekommunikasjonsledninger med større dimensjon for å redusere faren for brann.

WARNING: Aviso Para reduzir o risco de incêndio, utilize apenas terminais de fio de telecomunicações Nº. 26 AWG ou superiores.

WARNING: ¡Atención! Para reducir el riesgo de incendios, usar sólo líneas de telecomunicaciones de calibre No. 26 AWG o más gruesas.

WARNING: Varning! För att minska brandrisken skall endast Nr. 26 AWG eller större telekommunikationsledning användas.

Installation Safety Guidelines and Warnings

Observe the following guidelines and warnings before and during Services Router installation.

Chassis Lifting Guidelines

The weight of a fully configured chassis is approximately 12 lbs (5.4 kg) for a J2300 Services Router, 21 lbs (9.5 kg) for a J4300 Services Router, and 24 lb (10.9 kg) for a J6300 Services Router. Observe the following guidelines for lifting and moving a Services Router:

- Before moving the Services Router, read the guidelines in “Preparing for Router Installation” on page 541 to verify that the intended site meets the specified power, environmental, and clearance requirements.
- Before lifting or moving the Services Router, disconnect all external cables.
- As when lifting any heavy object, lift most of the weight with your legs rather than your back. Keep your knees bent and your back relatively straight and avoid twisting your body as you lift. Balance the load evenly and be sure that your footing is solid.

Installation Instructions Warning

WARNING: Read the installation instructions before you connect the router to a power source.

Waarschuwing Raadpleeg de installatie-aanwijzingen voordat u het systeem met de voeding verbindt.

Varoitus Lue asennusohjeet ennen järjestelmän yhdistämistä virtalähteeseen.

Attention Avant de brancher le système sur la source d'alimentation, consulter les directives d'installation.

Warnung Lesen Sie die Installationsanweisungen, bevor Sie das System an die Stromquelle anschließen.

WARNING: Avvertenza Consultare le istruzioni di installazione prima di collegare il sistema all'alimentatore.

Advarsel Les installasjonsinstruksjonene før systemet kobles til strømkilden.

Aviso Leia as instruções de instalação antes de ligar o sistema à sua fonte de energia.

¡Atención! Ver las instrucciones de instalación antes de conectar el sistema a la red de alimentación.

Varning! Läs installationsanvisningarna innan du kopplar systemet till dess strömförsörjningsenhet.

Rack-Mounting Requirements and Warnings

Ensure that the equipment rack into which the Services Router is installed is evenly and securely supported, to avoid the hazardous condition that could result from uneven mechanical loading.

WARNING: To prevent bodily injury when mounting or servicing the router in a rack, take the following precautions to ensure that the system remains stable. The following directives help maintain your safety:

- The router must be installed into a rack that is secured to the building structure.
- The router should be mounted at the bottom of the rack if it is the only unit in the rack.
- When mounting the router in a partially filled rack, load the rack from the bottom to the top with the heaviest component at the bottom of the rack.
- If the rack is provided with stabilizing devices, install the stabilizers before mounting or servicing the router in the rack.

WARNING: Waarschuwing Om lichamelijk letsel te voorkomen wanneer u dit toestel in een rek monteert of het daar een servicebeurt geeft, moet u speciale voorzorgsmaatregelen nemen om ervoor te zorgen dat het toestel stabiel blijft. De onderstaande richtlijnen worden verstrekt om uw veiligheid te verzekeren:

- De Juniper Networks router moet in een stellage worden geïnstalleerd die aan een bouwsel is verankerd.
- Dit toestel dient onderaan in het rek gemonteerd te worden als het toestel het enige in het rek is.
- Wanneer u dit toestel in een gedeeltelijk gevuld rek monteert, dient u het rek van onderen naar boven te laden met het zwaarste onderdeel onderaan in het rek.
- Als het rek voorzien is van stabiliseringshulpmiddelen, dient u de stabilisatoren te monteren voordat u het toestel in het rek monteert of het daar een servicebeurt geeft.

WARNING: Varoitus Kun laite asetetaan telineeseen tai huolletaan sen ollessa telineessä, on noudatettava erityisiä varotoimia järjestelmän vakavuuden

säilyttämiseksi, jotta vältytään loukkaantumiselta. Noudata seuraavia turvallisuusohjeita:

- Juniper Networks router on asennettava telineeseen, joka on kiinnitetty rakennukseen.
- Jos telineessä ei ole muita laitteita, aseta laite telineen alaosaan.
- Jos laite asetetaan osaksi täytettyyn telineeseen, aloita kuormittaminen sen alaosasta kaikkein raskaimmalla esineellä ja siirry sitten sen yläosaan.
- Jos telinettä varten on vakaimet, asenna ne ennen laitteen asettamista telineeseen tai sen huoltamista siinä.

WARNING: Attention Pour éviter toute blessure corporelle pendant les opérations de montage ou de réparation de cette unité en casier, il convient de prendre des précautions spéciales afin de maintenir la stabilité du système. Les directives ci-dessous sont destinées à assurer la protection du personnel:

- Le rack sur lequel est monté le Juniper Networks router doit être fixé à la structure du bâtiment.
- Si cette unité constitue la seule unité montée en casier, elle doit être placée dans le bas.
- Si cette unité est montée dans un casier partiellement rempli, charger le casier de bas en haut en plaçant l'élément le plus lourd dans le bas.
- Si le casier est équipé de dispositifs stabilisateurs, installer les stabilisateurs avant de monter ou de réparer l'unité en casier.

WARNING: Warnung Zur Vermeidung von Körperverletzung beim Anbringen oder Warten dieser Einheit in einem Gestell müssen Sie besondere Vorkehrungen treffen,

um sicherzustellen, daß das System stabil bleibt. Die folgenden Richtlinien sollen zur Gewährleistung Ihrer Sicherheit dienen:

- Der Juniper Networks router muß in einem Gestell installiert werden, das in der Gebäudestruktur verankert ist.
- Wenn diese Einheit die einzige im Gestell ist, sollte sie unten im Gestell angebracht werden.
- Bei Anbringung dieser Einheit in einem zum Teil gefüllten Gestell ist das Gestell von unten nach oben zu laden, wobei das schwerste Bauteil unten im Gestell anzubringen ist.
- Wird das Gestell mit Stabilisierungszubehör geliefert, sind zuerst die Stabilisatoren zu installieren, bevor Sie die Einheit im Gestell anbringen oder sie warten.

WARNING: Avvertenza Per evitare infortuni fisici durante il montaggio o la manutenzione di questa unità in un supporto, occorre osservare speciali precauzioni per garantire che il sistema rimanga stabile. Le seguenti direttive vengono fornite per garantire la sicurezza personale:

- Il Juniper Networks router deve essere installato in un telaio, il quale deve essere fissato alla struttura dell'edificio.
- Questa unità deve venire montata sul fondo del supporto, se si tratta dell'unica unità da montare nel supporto.
- Quando questa unità viene montata in un supporto parzialmente pieno, caricare il supporto dal basso all'alto, con il componente più pesante sistemato sul fondo del supporto.
- Se il supporto è dotato di dispositivi stabilizzanti, installare tali dispositivi prima di montare o di procedere alla manutenzione dell'unità nel supporto.

WARNING: Advarsel Unngå fysiske skader under montering eller reparasjonsarbeid på denne enheten når den befinner seg i et kabinett. Vær nøye med at systemet er stabilt. Følgende retningslinjer er gitt for å verne om sikkerheten:

- Juniper Networks router må installeres i et stativ som er forankret til bygningsstrukturen.
- Denne enheten bør monteres nederst i kabinettet hvis dette er den eneste enheten i kabinettet.
- Ved montering av denne enheten i et kabinett som er delvis fylt, skal kabinettet lastes fra bunnen og opp med den tyngste komponenten nederst i kabinettet.
- Hvis kabinettet er utstyrt med stabiliseringsutstyr, skal stabilisatorene installeres før montering eller utføring av reparasjonsarbeid på enheten i kabinettet.

WARNING: Aviso Para se prevenir contra danos corporais ao montar ou reparar esta unidade numa estante, deverá tomar precauções especiais para se certificar de que o sistema possui um suporte estável. As seguintes directrizes ajudá-lo-ão a efectuar o seu trabalho com segurança:

- O Juniper Networks router deverá ser instalado numa prateleira fixa à estrutura do edificio.
- Esta unidade deverá ser montada na parte inferior da estante, caso seja esta a única unidade a ser montada.
- Ao montar esta unidade numa estante parcialmente ocupada, coloque os itens mais pesados na parte inferior da estante, arrumando-os de baixo para cima.
- Se a estante possuir um dispositivo de estabilização, instale-o antes de montar ou reparar a unidade.

WARNING: ¡Atención! Para evitar lesiones durante el montaje de este equipo sobre un bastidor, o posteriormente durante su mantenimiento, se debe poner

mucho cuidado en que el sistema quede bien estable. Para garantizar su seguridad, proceda según las siguientes instrucciones:

- El Juniper Networks router debe instalarse en un bastidor fijado a la estructura del edificio.
- Colocar el equipo en la parte inferior del bastidor, cuando sea la única unidad en el mismo.
- Cuando este equipo se vaya a instalar en un bastidor parcialmente ocupado, comenzar la instalación desde la parte inferior hacia la superior colocando el equipo más pesado en la parte inferior.
- Si el bastidor dispone de dispositivos estabilizadores, instalar éstos antes de montar o proceder al mantenimiento del equipo instalado en el bastidor.

WARNING: Varning! För att undvika kroppsskada när du installerar eller utför underhållsarbete på denna enhet på en ställning måste du vidta särskilda försiktighetsåtgärder för att försäkra dig om att systemet står stadigt. Följande riktlinjer ges för att trygga din säkerhet:

- Juniper Networks router måste installeras i en ställning som är förankrad i byggnadens struktur.
- Om denna enhet är den enda enheten på ställningen skall den installeras längst ned på ställningen.
- Om denna enhet installeras på en delvis fylld ställning skall ställningen fyllas nedifrån och upp, med de tyngsta enheterna längst ned på ställningen.
- Om ställningen är försedd med stabiliseringsdon skall dessa monteras fast innan enheten installeras eller underhålls på ställningen.

Ramp Warning

WARNING: When installing the router, do not use a ramp inclined at more than 10 degrees.

Waarschuwing Gebruik een oprijplaat niet onder een hoek van meer dan 10 graden.

Varoitus Älä käytä sellaista kaltevaa pintaa, jonka kaltevuus ylittää 10 astetta.

Attention Ne pas utiliser une rampe dont l’inclinaison est supérieure à 10 degrés.

Warnung Keine Rampen mit einer Neigung von mehr als 10 Grad verwenden.

WARNING: Avvertenza Non usare una rampa con pendenza superiore a 10 gradi.

Advarsel Bruk aldri en rampe som heller mer enn 10 grader.

Aviso Não utilize uma rampa com uma inclinação superior a 10 graus.

¡Atención! No usar una rampa inclinada más de 10 grados

Varning! Använd inte ramp med en lutning på mer än 10 grader.

Laser and LED Safety Guidelines and Warnings

Single-mode Physical Interface Modules (PIMs) are equipped with laser transmitters, which are considered a Class 1 Laser Product by the U.S. Food and Drug Administration, and are evaluated as a Class 1 Laser Product per EN 60825–1 + A11 + A2 requirements.

Observe the following guidelines and warnings.

General Laser Safety Guidelines

When working around PIMs, observe the following safety guidelines to prevent eye injury:

- Do not look into unterminated ports or at fibers that connect to unknown sources.
- Do not examine unterminated optical ports with optical instruments.
- Avoid direct exposure to the beam.

WARNING: Unterminated optical connectors can emit invisible laser radiation. The lens in the human eye focuses all the laser power on the retina, so focusing the eye directly on a laser source—even a low-power laser—could permanently damage the eye.

Class 1 Laser Product Warning

WARNING: Class 1 laser product.

Waarschuwing Klasse-1 laser produkt.

Varoitus Luokan 1 lasertuote.

Attention Produit laser de classe I.

Warnung Laserprodukt der Klasse 1.

WARNING: Avvertenza Prodotto laser di Classe 1.

Advarsel Laserprodukt av klasse 1.

Aviso Produto laser de classe 1.

¡Atención! Producto láser Clase I.

Varning! Laserprodukt av klass 1.

Class 1 LED Product Warning

WARNING: Class 1 LED product.

Waarschuwing Klasse 1 LED-product.

Varoitus Luokan 1 valodiodituote.

Attention Alarme de produit LED Class I.

Warnung Class 1 LED-Produktwarnung.

WARNING: Avvertenza Avvertenza prodotto LED di Classe 1.

Advarsel LED-produkt i klasse 1.

Aviso Produto de classe 1 com LED.

¡Atención! Aviso sobre producto LED de Clase 1.

Varning! Lysdiodprodukt av klass 1.

Laser Beam Warning

WARNING: Do not stare into the laser beam or view it directly with optical instruments.

WARNING: Waarschuwing Niet in de straal staren of hem rechtstreeks bekijken met optische instrumenten.

WARNING: Varoitus Älä katso säteeseen äläkä tarkastele sitä suoraan optisen laitteen avulla.

WARNING: Attention Ne pas fixer le faisceau des yeux, ni l'observer directement à l'aide d'instruments optiques.

WARNING: Warnung Nicht direkt in den Strahl blicken und ihn nicht direkt mit optischen Geräten prüfen.

WARNING: Avvertenza Non fissare il raggio con gli occhi né usare strumenti ottici per osservarlo direttamente.

WARNING: Advarsel Stirr eller se ikke direkte p strlen med optiske instrumenter.

WARNING: Aviso Não olhe fixamente para o raio, nem olhe para ele directamente com instrumentos ópticos.

WARNING: ¡Atención! No mirar fijamente el haz ni observarlo directamente con instrumentos ópticos.

WARNING: Varning! Rikta inte blicken in mot strålen och titta inte direkt på den genom optiska instrument.

Radiation from Open Port Apertures Warning

WARNING: Because invisible radiation may be emitted from the aperture of the port when no fiber cable is connected, avoid exposure to radiation and do not stare into open apertures.

WARNING: Waarschuwing Aangezien onzichtbare straling vanuit de opening van de poort kan komen als er geen fiberkabel aangesloten is, dient blootstelling aan straling en het kijken in open openingen vermeden te worden.

WARNING: Varoitus Koska portin aukosta voi emittoitua näkymätöntä säteilyä, kun kuitukaapelia ei ole kytkettynä, vältä säteilylle altistumista äläkä katso avoimiin aukkoihin.

WARNING: Attention Des radiations invisibles à l'il nu pouvant traverser l'ouverture du port lorsqu'aucun câble en fibre optique n'y est connecté, il est recommandé de ne pas regarder fixement l'intérieur de ces ouvertures.

WARNING: Warnung Aus der Port-Öffnung können unsichtbare Strahlen emittieren, wenn kein Glasfaserkabel angeschlossen ist. Vermeiden Sie es, sich den Strahlungen auszusetzen, und starren Sie nicht in die Öffnungen!

WARNING: Avvertenza Quando i cavi in fibra non sono inseriti, radiazioni invisibili possono essere emesse attraverso l'apertura della porta. Evitate di esporvi alle radiazioni e non guardate direttamente nelle aperture.

WARNING: Advarsel Unngå utsettelse for stråling, og stirr ikke inn i åpninger som er åpne, fordi usynlig stråling kan emiteres fra portens åpning når det ikke er tilkoblet en fiberkabel.

WARNING: Aviso Dada a possibilidade de emissão de radiação invisível através do orifício da via de acesso, quando esta não tiver nenhum cabo de fibra conectado, deverá evitar a exposição à radiação e não deverá olhar fixamente para orifícios que se encontrarem a descoberto.

WARNING: ¡Atención! Debido a que la apertura del puerto puede emitir radiación invisible cuando no existe un cable de fibra conectado, evite mirar directamente a las aperturas para no exponerse a la radiación.

WARNING: Varning! Osynlig strålning kan avges från en portöppning utan ansluten fiberkabel och du bör därför undvika att bli utsatt för strålning genom att inte stirra in i oskyddade öppningar.

Maintenance and Operational Safety Guidelines and Warnings

As you maintain the Services Router, observe the following guidelines and warnings.

Battery Handling Warning

WARNING: Replacing the battery incorrectly might result in an explosion. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

WARNING: Waarschuwing Er is ontploffingsgevaar als de batterij verkeerd vervangen wordt. Vervang de batterij slechts met hetzelfde of een equivalent type dat door de fabrikant aanbevolen is. Gebruikte batterijen dienen overeenkomstig fabrieksvoorschriften weggeworpen te worden.

WARNING: Varoitus Räjähdyksen vaara, jos akku on vaihdettu väärään akkuun. Käytä vaihtamiseen ainoastaan saman- tai vastaavantyyppistä akkua, joka on valmistajan suosittelema. Hävitä käytetyt akut valmistajan ohjeiden mukaan.

WARNING: Attention Danger d'explosion si la pile n'est pas remplacée correctement. Ne la remplacer que par une pile de type semblable ou équivalent, recommandée par le fabricant. Jeter les piles usagées conformément aux instructions du fabricant.

WARNING: Warnung Bei Einsetzen einer falschen Batterie besteht Explosionsgefahr. Ersetzen Sie die Batterie nur durch den gleichen oder vom Hersteller empfohlenen Batterietyp. Entsorgen Sie die benutzten Batterien nach den Anweisungen des Herstellers.

WARNING: Advarsel Det kan være fare for eksplosjon hvis batteriet skiftes på feil måte. Skift kun med samme eller tilsvarende type som er anbefalt av produsenten. Kasser brukte batterier i henhold til produsentens instruksjoner.

WARNING: Avvertenza Pericolo di esplosione se la batteria non è installata correttamente. Sostituire solo con una di tipo uguale o equivalente, consigliata dal produttore. Eliminare le batterie usate secondo le istruzioni del produttore.

WARNING: Aviso Existe perigo de explosão se a bateria for substituída incorrectamente. Substitua a bateria por uma bateria igual ou de um tipo equivalente recomendado pelo fabricante. Destrua as baterias usadas conforme as instruções do fabricante.

WARNING: ¡Atención! Existe peligro de explosión si la batería se reemplaza de manera incorrecta. Reemplazar la batería exclusivamente con el mismo tipo o el equivalente recomendado por el fabricante. Desechar las baterías gastadas según las instrucciones del fabricante.

WARNING: Varning! Explosionsfara vid felaktigt batteribyte. Ersätt endast batteriet med samma batterityp som rekommenderas av tillverkaren eller motsvarande. Följ tillverkarens anvisningar vid kassering av använda batterier.

Jewelry Removal Warning

WARNING: Before working on equipment that is connected to power lines, remove jewelry, including rings, necklaces, and watches. Metal objects heat up when connected to power and ground and can cause serious burns or weld the metal object to the terminals.

WARNING: Waarschuwing Alvorens aan apparatuur te werken die met elektrische leidingen is verbonden, sieraden (inclusief ringen, kettingen en horloges) verwijderen. Metalen voorwerpen worden warm wanneer ze met stroom en aarde zijn verbonden, en kunnen ernstige brandwonden veroorzaken of het metalen voorwerp aan de aansluitklemmen lassen.

WARNING: Varoitus Ennen kuin työskentelet voimavirtajohtoihin kytkettyjen laitteiden parissa, ota pois kaikki korut (sormukset, kaulakorut ja kellot mukaan lukien). Metalliesineet kuumenevat, kun ne ovat yhteydessä sähkövirran ja maan kanssa, ja ne voivat aiheuttaa vakavia palovammoja tai hitsata metalliesineet kiinni liitäntänapoihin.

WARNING: Attention Avant d'accéder à cet équipement connecté aux lignes électriques, ôter tout bijou (anneaux, colliers et montres compris). Lorsqu'ils sont branchés à l'alimentation et reliés à la terre, les objets métalliques chauffent, ce qui peut provoquer des blessures graves ou souder l'objet métallique aux bornes.

WARNING: Warnung Vor der Arbeit an Geräten, die an das Netz angeschlossen sind, jeglichen Schmuck (einschließlich Ringe, Ketten und Uhren) abnehmen. Metallgegenstände erhitzen sich, wenn sie an das Netz und die Erde angeschlossen werden, und können schwere Verbrennungen verursachen oder an die Anschlußklemmen angeschweißt werden.

WARNING: Avvertenza Prima di intervenire su apparecchiature collegate alle linee di alimentazione, togliersi qualsiasi monile (inclusi anelli, collane, braccialetti ed orologi). Gli oggetti metallici si riscaldano quando sono collegati tra punti di alimentazione e massa: possono causare ustioni gravi oppure il metallo può saldarsi ai terminali.

WARNING: Advarsel Fjern alle smykker (inkludert ringer, halskjeder og klokker) før du skal arbeide på utstyr som er koblet til kraftledninger. Metallgjenstander som er koblet til kraftledninger og jord blir svært varme og kan forårsake alvorlige brannskader eller smelte fast til polene.

WARNING: Aviso Antes de trabalhar em equipamento que esteja ligado a linhas de corrente, retire todas as jóias que estiver a usar (incluindo anéis, fios e relógios). Os objectos metálicos aquecerão em contacto com a corrente e em contacto com

a ligação à terra, podendo causar queimaduras graves ou ficarem soldados aos terminais.

WARNING: ¡Atención! Antes de operar sobre equipos conectados a líneas de alimentación, quitarse las joyas (incluidos anillos, collares y relojes). Los objetos de metal se calientan cuando se conectan a la alimentación y a tierra, lo que puede ocasionar quemaduras graves o que los objetos metálicos queden soldados a los bornes.

WARNING: Varning! Tag av alla smycken (inklusive ringar, halsband och armbandsur) innan du arbetar på utrustning som är kopplad till kraftledningar. Metallobjekt hettas upp när de kopplas ihop med ström och jord och kan förorsaka allvarliga brännskador; metallobjekt kan också sammansvetsas med kontakterna.

Lightning Activity Warning

WARNING: Do not work on the system or connect or disconnect cables during periods of lightning activity.

WARNING: Waarschuwing Tijdens onweer dat gepaard gaat met bliksem, dient u niet aan het systeem te werken of kabels aan te sluiten of te ontkoppelen.

WARNING: Varoitus Älä työskentele järjestelmän parissa äläkä yhdistä tai irrota kaapeleita ukkosilmalla.

WARNING: Attention Ne pas travailler sur le système ni brancher ou débrancher les câbles pendant un orage.

WARNING: Warnung Arbeiten Sie nicht am System und schließen Sie keine Kabel an bzw. trennen Sie keine ab, wenn es gewittert.

WARNING: Avvertenza Non lavorare sul sistema o collegare oppure scollegare i cavi durante un temporale con fulmini.

WARNING: Advarsel Utfør aldri arbeid på systemet, eller koble kabler til eller fra systemet når det tordner eller lyner.

WARNING: Aviso Não trabalhe no sistema ou ligue e desligue cabos durante períodos de mau tempo (trovoada).

WARNING: ¡Atención! No operar el sistema ni conectar o desconectar cables durante el transcurso de descargas eléctricas en la atmósfera.

WARNING: Varning! Vid åska skall du aldrig utföra arbete på systemet eller ansluta eller koppla loss kablar.

Operating Temperature Warning

WARNING: To prevent the router from overheating, do not operate it in an area that exceeds the maximum recommended ambient temperature of 104°F (40°C). To prevent airflow restriction, allow at least 6 inches (15.2 cm) of clearance around the ventilation openings.

WARNING: Waarschuwing Om te voorkomen dat welke router van de Juniper Networks router dan ook oververhit raakt, dient u deze niet te bedienen op een plaats waar de maximale aanbevolen omgevingstemperatuur van 40°C wordt overschreden. Om te voorkomen dat de luchtstroom wordt beperkt, dient er minstens 15,2 cm speling rond de ventilatie-openingen te zijn.

WARNING: Varoitus Ettei Juniper Networks router-sarjan reititin ylikuumentuisi, sitä ei saa käyttää tilassa, jonka lämpötila ylittää korkeimman suositellun ympäristölämpötilan 40°C. Ettei ilmanvaihto estyisi, tuuletusaukkojen ympärille on jätettävä ainakin 15,2 cm tilaa.

WARNING: Attention Pour éviter toute surchauffe des routeurs de la gamme Juniper Networks router, ne l'utilisez pas dans une zone où la température ambiante est supérieure à 40°C. Pour permettre un flot d'air constant, dégagez un espace d'au moins 15,2 cm autour des ouvertures de ventilations.

WARNING: Warnung Um einen Router der router vor Überhitzung zu schützen, darf dieser nicht in einer Gegend betrieben werden, in der die Umgebungstemperatur das empfohlene Maximum von 40°C überschreitet. Um Lüftungsverschluß zu verhindern, achten Sie darauf, daß mindestens 15,2 cm lichter Raum um die Lüftungsöffnungen herum frei bleibt.

WARNING: Avvertenza Per evitare il surriscaldamento dei router, non adoperateli in un locale che ecceda la temperatura ambientale massima di 40°C. Per evitare che la circolazione dell'aria sia impedita, lasciate uno spazio di almeno 15.2 cm di fronte alle aperture delle ventole.

WARNING: Advarsel Unngå overoppheting av eventuelle rutere i Juniper Networks router Disse skal ikke brukes på steder der den anbefalte maksimale omgivelsestemperaturen overstiger 40°C (104°F). Sørg for at klaringen rundt lufteåpningene er minst 15,2 cm (6 tommer) for å forhindre nedsatt luftsirkulasjon.

WARNING: Aviso Para evitar o sobreaquecimento do encaminhador Juniper Networks router, não utilize este equipamento numa área que exceda a temperatura máxima recomendada de 40°C. Para evitar a restrição à circulação de ar, deixe pelo menos um espaço de 15,2 cm à volta das aberturas de ventilação.

WARNING: ¡Atención! Para impedir que un encaminador de la serie Juniper Networks router se recaliente, no lo haga funcionar en un área en la que se supere la temperatura ambiente máxima recomendada de 40°C. Para impedir la restricción de la entrada de aire, deje un espacio mínimo de 15,2 cm alrededor de las aperturas para ventilación.

WARNING: Varning! Förhindra att en Juniper Networks router överhettas genom att inte använda den i ett område där den maximalt rekommenderade omgivningstemperaturen på 40°C överskrids. Förhindra att luftcirkulationen inskränks genom att se till att det finns fritt utrymme på minst 15,2 cm omkring ventilationsöppningarna.

Product Disposal Warning

WARNING: Disposal of this product must be handled according to all national laws and regulations.

WARNING: Waarschuwing Dit produkt dient volgens alle landelijke wetten en voorschriften te worden afgedankt.

WARNING: Varoitus Tämän tuotteen lopullisesta hävittämisestä tulee huolehtia kaikkia valtakunnallisia lakeja ja säännöksiä noudattaen.

WARNING: Attention La mise au rebut définitive de ce produit doit être effectuée conformément à toutes les lois et réglementations en vigueur.

WARNING: Warnung Dieses Produkt muß den geltenden Gesetzen und Vorschriften entsprechend entsorgt werden.

WARNING: Avvertenza L'eliminazione finale di questo prodotto deve essere eseguita osservando le normative italiane vigenti in materia

WARNING: Advarsel Endelig disponering av dette produktet må skje i henhold til nasjonale lover og forskrifter.

WARNING: Aviso A descartagem final deste produto deverá ser efectuada de acordo com os regulamentos e a legislação nacional.

WARNING: ¡Atención! El desecho final de este producto debe realizarse según todas las leyes y regulaciones nacionales

WARNING: Varning! Slutlig kassering av denna produkt bör skötas i enlighet med landets alla lagar och föreskrifter.

Agency Approvals

The Services Router complies with the following standards:

- Safety
  - CAN/CSA-22.2 No. 60950–1–03–UL 60950–1 Safety of Information Technology Equipment
  - EN 60950–1 Safety of Information Technology Equipment
  - EN 60825-1 Safety of Laser Products - Part 1: Equipment Classification, Requirements and User's Guide
- EMC
  - AS/NZS 3548 Class B (Australia/New Zealand)
  - EN 55022 Class B Emissions (Europe)
  - FCC Part 15 Class B (USA)
  - VCCI Class B (Japan)
  - FCC Part 68
  - Industry Canada CS-03
- Immunity
  - EN 61000-3-2 Power Line Harmonics
  - EN 61000-3-3 Voltage Fluctuations and Flicker
  - EN 61000-4-2 ESD
  - EN 61000-4-3 Radiated Immunity
  - EN 61000-4-4 EFT
  - EN 61000-4-5 Surge
  - EN 61000-4-6 Low Frequency Common Immunity
  - EN 61000-4-11 Voltage Dips and Sags
- ETSI
  - ETSI EN-300386-2 Telecommunication Network Equipment. Electromagnetic Compatibility Requirements

Compliance Statements for EMC Requirements

Canada

This Class B digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

The Industry Canada label identifies certified equipment. This certification means that the equipment meets certain telecommunications network protective, operational, and safety requirements. Industry Canada does not guarantee the equipment will operate to the users' satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local telecommunications company. The equipment must also be installed using an acceptable method of connection. In some cases, the inside wiring associated with a single line individual service may be extended by means of a certified connector assembly. The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations.

Repairs to certified equipment should be made by an authorized Canadian maintenance facility designated by the supplier. Any repairs or alterations made by the user to this equipment, or equipment malfunctions, may give the telecommunications company cause to request the user to disconnect the equipment.

CAUTION: Users should not attempt to make electrical ground connections by themselves, but should contact the appropriate inspection authority or an electrician, as appropriate.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines, and internal metallic water pipe system, if present, are connected together. This precaution may be particularly important in rural areas.

Japan

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＢ情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。

The preceding translates as follows:

This is a Class B product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this product is used near a radio or television receiver in a domestic environment, it may cause radio interference. Install and use the equipment according to the instruction manual.

Taiwan

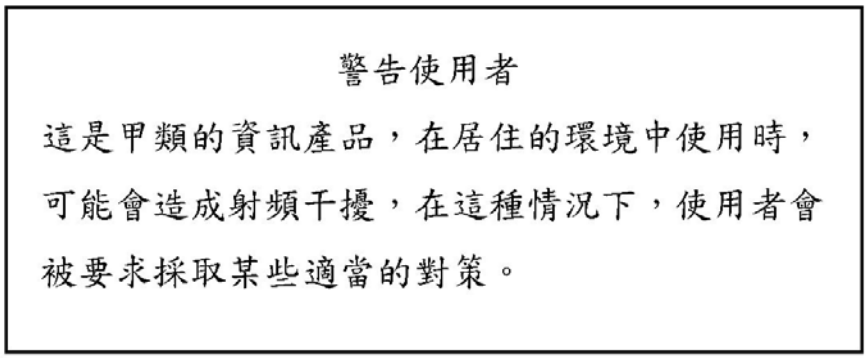

United States

The Services Router has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

FCC Part 15 Statement

This equipment has been tested and found to comply with the limits for a Class B digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is

encouraged to try and correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and the receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or TV technician for help.

FCC Part 68 Statement

This equipment complies with Part 68 of the Federal Communications Commission (FCC) rules. On the product is a label that contains the FCC registration number for this device. If requested, this information must be provided to the telephone company.

This equipment is designed to be connected to the telephone network or premises wiring using a compatible modular jack which is Part 68 compliant. See installation instructions for details.

If this device causes harm to the telephone network, the telephone company will notify you in advance that temporary discontinuance of service may be required. The telephone company may request that you disconnect the equipment until the problem is resolved. The telephone company may make changes in its facilities, equipment, operations or procedures that could affect the operation of this equipment. If this happens, the telephone company will provide advance notice in order for you to make necessary modifications to maintain uninterrupted service.

If trouble is experienced with this equipment or for repair or warranty information, please follow the applicable procedures explained in the “Technical Support” section of this manual.

- FCC Registration Number—See label on product.
- Required Connector (USOC)—RJ-48C
- Service Order Code (SOC)—6.ON

Part 11

Customer Support and Product Return

Chapter 29

Contacting Customer Support and Returning Hardware

This chapter describes how to return the Services Router or individual components to Juniper Networks for repair or replacement. It contains the following topics:

Locating Component Serial Numbers

Before contacting Juniper Networks to request a Return Materials Authorization (RMA), you must find the serial number on the router or component. To list the router components and their serial numbers, enter the following command-line interface (CLI) command:

```
    user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN000192AB        J4300
Midplane         REV 02.04 710-010001  CORE99563
System IO        REV 02.03 710-010003  CORE100885        P12/P45 System IO board
Routing Engine   RevX2.6  750-010005   IWGS40735451      RE-J.2
FPC 0                                                   FPC
  PIC 0                                                 2x FE
```

NOTE: In the show chassis hardware output, PIMs are identified as PICs.

Most components also have a small rectangular serial number ID label (see Figure 112 through Figure 114) attached to the component body.

Figure 112: J2300 Serial Number ID Label

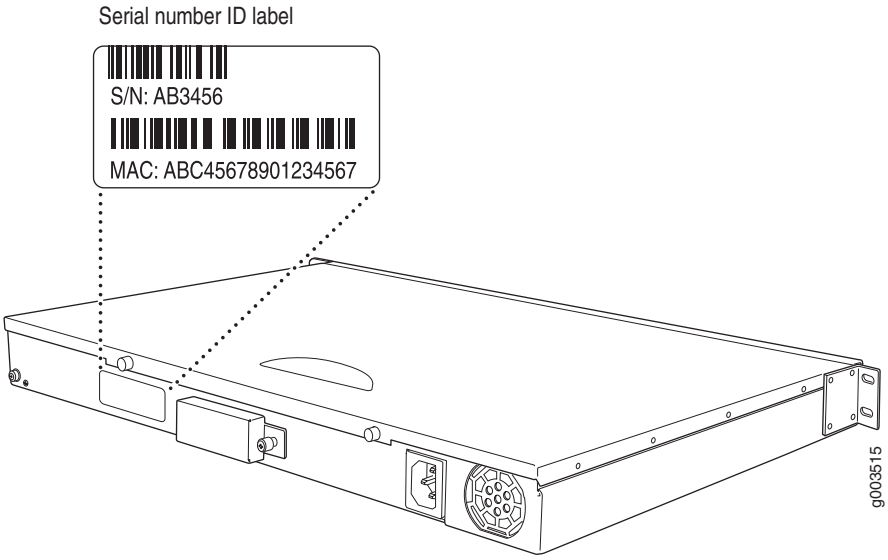

Figure 113: J4300 Serial Number ID Label

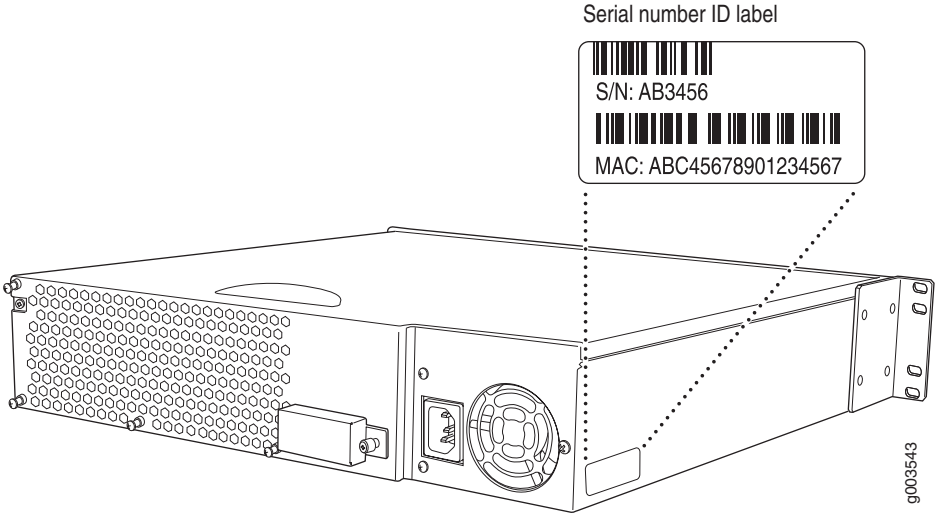

Figure 114: J6300 Serial Number ID Label

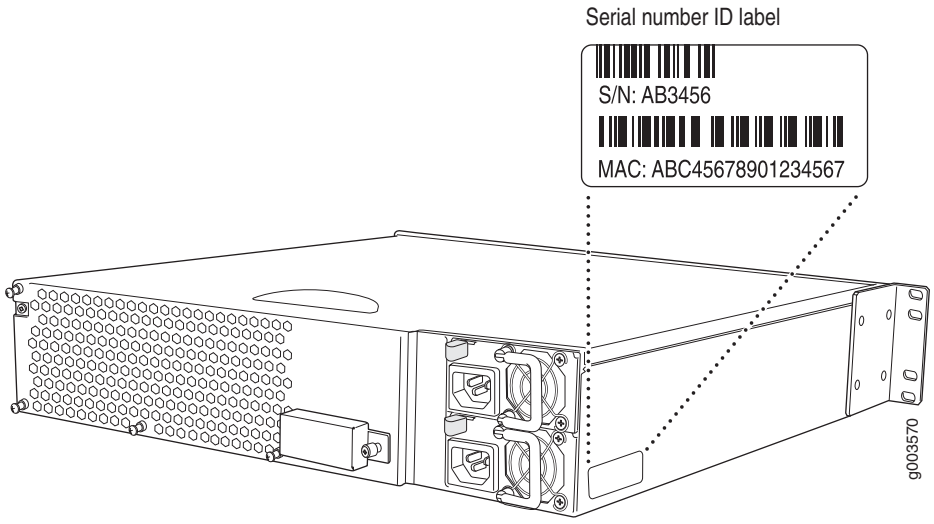

The following sections describe the label location on each type of component:

- PIM Serial Number Label on page 605
- J6300 Power Supply Serial Number Labels on page 605

PIM Serial Number Label

The PIMs installed in the J4300 and J6300 Services Routers are field-replaceable. Each PIM has a unique serial number. The serial number label is located on the right side of the PIM, when the PIM is horizontally oriented (as it would be installed in the router). The exact location may be slightly different on different PIMs, depending on the placement of components on the PIM board.

J6300 Power Supply Serial Number Labels

The power supplies installed in the J6300 Services Router are field-replaceable. Each power supply has a unique serial number. The serial number label is located on the top of the AC power supply.

Contacting Customer Support

After you have located the serial numbers of the components you need to return, contact Juniper Networks Technical Assistance Center (JTAC) in one of the following ways.

You can contact JTAC 24 hours a day, seven days a week.

- On the Web, using the Case Manager link at http://www.juniper.net/support/
- By telephone:

  From the US and Canada: 1-888-314-JTAC

  From all other locations: 1-408-745-9500

If contacting JTAC by telephone, enter your 11–digit case number followed by the pound (#) key if this is an existing case, or press the star (*) key to be routed to the next available support engineer.

Information You Might Need to Supply to JTAC

When requesting support from JTAC by telephone, be prepared to provide the following information:

- Your existing case number, if you have one
- Details of the failure or problem
- Type of activity being performed on the router when the problem occurred
- Configuration data displayed by one or more show commands

Return Procedure

If the problem cannot be resolved by the JTAC technician, an RMA number is issued. This number is used to track the returned material at the factory and to return repaired or new components to the customer as needed.

NOTE: Do not return any component to Juniper Networks unless you have first obtained an RMA number. Juniper Networks reserves the right to refuse shipments that do not have an RMA. Refused shipments will be returned to the customer via collect freight.

For more information about return and repair policies, see the customer support Web page at http:/www./juniper.net/support/guidelines.html.

For product problems or technical support issues, open a support case using the Case Manager link at http://www.juniper.net/support/, or call 1–888–314–JTAC (within the United States) or 1–408–745–9500 (outside the United States).

When you need to return a component, follow this procedure:

1. Determine the part number and serial number of the component. For instructions, see “Locating Component Serial Numbers” on page 603.

2. Obtain a Return Materials Authorization (RMA) number from the Juniper Networks Technical Assistance Center (JTAC). You can send e-mail or telephone as described above.

3. Provide the following information in your e-mail message or during the telephone call:

   - Part number and serial number of component
   - Your name, organization name, telephone number, and fax number
   - Description of the failure

4. The support representative validates your request and issues an RMA number for return of the component.

5. Pack the router or component for shipment, as described in "Packing a Router or Component for Shipment" on page 607.

Packing a Router or Component for Shipment

This section contains the following topics:

Tools and Parts Required

To remove components from the router or the router from a rack, you need the following tools and parts:

- Blank panels to cover empty slots
- Electrostatic bag or antistatic mat, for each component
- Electrostatic discharge (ESD) grounding wrist strap
- Phillips (+) screwdrivers, numbers 1 and 2

Packing the Services Router for Shipment

To pack the router for shipment, follow this procedure:

1. Retrieve the shipping carton and packing materials in which the router was originally shipped. If you do not have these materials, contact your Juniper Networks representative about approved packaging materials.

2. Attach an electrostatic discharge (ESD) grounding strap to your bare wrist and connect the strap to the ESD point on the chassis, or to an outside ESD point if the router is disconnected from earth ground. For more information about ESD, see "Preventing Electrostatic Discharge Damage" on page 567.

3. On the console or other management device connected to the master Routing Engine, enter CLI operational mode and issue the following command to shut down the router software.

   ```
   user@host> request system halt
   ```

   Wait until a message appears on the console confirming that the operating system has halted. For more information about the command, see "Halting the Services Router with the CLI" on page 514.

4. Shut down power to the router by pressing the power button on the front panel of the router.

5. Disconnect power from the router. For instructions, see "Replacing a Power Supply Cord in a J2300 or J4300 Router" on page 532 or "Replacing a Power Supply Cord in a J6300 Router" on page 536.

6. Remove the cables that connect to all external devices. For instructions, see "Removing a PIM Cable" on page 522.

7. Remove all field-replaceable units (FRUs) from the router.

8. If the router is installed on a wall or rack, have one person support the weight of the router, while another person unscrews and removes the mounting screws.

9. Place the router in the shipping carton.

10. Cover the router with an ESD bag, and place the packing foam on top of and around the router.

11. Replace the accessory box on top of the packing foam.

12. Securely tape the box closed.

13. Write the RMA number on the exterior of the box to ensure proper tracking.

Packing Components for Shipment

To pack and ship individual components, follow these guidelines:

- When you return components, make sure they are adequately protected with packing materials and packed so that the pieces are prevented from moving around inside the carton.
- Use the original shipping materials if they are available.
- Place individual boards in electrostatic bags.
- Write the RMA number on the exterior of the box to ensure proper tracking.

CAUTION: Do not stack any of the router components.

Part 12

Indexes

Index

Symbols

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

S

T

U

V

W

X

Y