JUNOS™ Internet Software
for J-series™, M-series™, and T-series™
Routing Platforms

System Basics Configuration Guide

Release 7.4

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, CA 94089
USA
408-745-2000

www.juniper.net

Part Number: 530-014073-01, Revision 1

JUNOS Internet Software for J-series, M-series, and T-series Routing Platforms System Basics Configuration Guide, Release 7.4
Writing: Lisa Kelly
Editing: Stella Hackell
Illustration: Nathaniel Woodward
Cover design: Edmonds Design

Revision History
14 September 2005—Revision 1

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

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Year 2000 Notice

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

Software License

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

For complete product documentation, please see the Juniper Networks Web site at www.juniper.net/techpubs.

End User License Agreement

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

1. The Parties. The parties to this Agreement are Juniper Networks, Inc. and its subsidiaries (collectively "Juniper"), and the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. The Software. In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, and updates and releases of such software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller.

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. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. 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, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use the 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 for Customer's internal business purposes.

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

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

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

10. Taxes. All license fees for the Software are exclusive of taxes, withholdings, duties, or levies (collectively "Taxes"). Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software.

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

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

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

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

15. Miscellaneous. This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

Part 4 System Management

Part 5 Configuration Groups

Part 6 Access

Part 7 Security Services

Part 8 Router Chassis

Part 9 Indexes

Table of Contents

Chapter 9 **Summary of CLI Environment Commands** 285

Chapter 10 **Summary of CLI Configuration Mode Commands** 293

Part 3 Software Installation and Upgrade

Part 4 System Management

Chapter 22 **Configuring User Access** 401

Chapter 23 **Configuring Time** 417

Chapter 24 **Configuring System Log Messages** 427

Chapter 26 Security Configuration Example 493

Chapter 27 Summary of System Management Configuration Statements 523

Part 5 Configuration Groups

Part 6 Access

Chapter 30 **Summary of Access Configuration Statements** **683**

Part 8 Router Chassis

Chapter 34 Router Chassis Configuration Guidelines 799

Part 9 Indexes

About This Guide

This preface provides the following guidelines for using the *JUNOS Internet Software for J-series, M-series, and T-series Routing Platforms System Basics Configuration Guide* and related Juniper Networks, Inc., technical documents:

- Objectives on page xxxi
- Audience on page xxxii
- Using the Indexes on page xxxii
- Documentation Conventions on page xxxiii
- Related Juniper Networks Documentation on page xxxiv
- Documentation Feedback on page xxxvii
- Requesting Support on page xxxvii

Objectives

This guide provides an overview of the JUNOS software and describes how to install and upgrade the software. This manual also describes how to configure system management functions and how to configure the chassis, including user accounts, passwords, and redundancy.

NOTE: This guide documents Release 7.4 of the JUNOS Internet software. For additional information about the JUNOS software—either corrections to or information that might have been omitted from this guide—see the software release notes at http://www.juniper.net/.

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks J-series, M-series, or T-series routing platform.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

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

Using the Indexes

This guide contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines*, refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

Documentation Conventions

Table 1 defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
☞	Informational note	Indicates important features or instructions.
⚠	Caution	Indicates a situation that might result in loss of data or hardware damage.

Table 2 defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

`Fixed-width typeface`	Represents output on the terminal screen.	`user@host>` **`show chassis alarms`**
		`No alarms currently active`

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*
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*]

Convention	Element	Example

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

Table 3 lists the software and hardware guides and release notes for Juniper Networks J-series, M-series, and T-series routing platforms and describes the contents of each document. Table 4 lists the books included in the *Network Operations Guide* series.

Table 3: Technical Documentation for J-series, M-series, and T-series Routing Platforms

Document	Description
JUNOS Internet Software for J-series, M-series, and T-series Routing Platforms Configuration Guides	
Class of Service	Provides an overview of the class-of-service (CoS) functions of the JUNOS software and describes how to configure CoS features, including configuring multiple forwarding classes for transmitting packets, defining which packets are placed into each output queue, scheduling the transmission service level for each queue, and managing congestion through the random early detection (RED) algorithm.
Feature Guide	Provides a detailed explanation and configuration examples for several of the most complex features in the JUNOS software.
JUNOS-FIPS	(M-series and T-series routing platforms only) Provides an overview of JUNOS-FIPS 140-2 concepts and describes how to install and configure the JUNOS-FIPS software. Describes FIPS-related commands and how to configure, authorize, and zeroize the Adaptive Services (AS) II FIPS Physical Interface Card (PIC).
MPLS Applications	Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols.
Multicast Protocols	Provides an overview of multicast concepts and describes how to configure multicast routing protocols.
Network Interfaces	Provides an overview of the network interface functions of the JUNOS software and describes how to configure the network interfaces on the routing platform.
Network Management	Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP and accounting options.
Policy Framework	Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, forwarding options, and cflowd.

Document	Description
Routing Protocols	Provides an overview of routing concepts and describes how to configure routing, routing instances, and unicast routing protocols.
Services Interfaces	Provides an overview of the services interfaces functions of the JUNOS software and describes how to configure the services interfaces on the routing platform.
System Basics	Provides an overview of the JUNOS software and describes how to install and upgrade the software. This guide also describes how to configure system management functions and how to configure the chassis, including user accounts, passwords, and redundancy.
VPNs	Provides an overview and describes how to configure Layer 2 and Layer 3 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 2 circuits. Provides configuration examples.
JUNOS References	
Interfaces Command Reference	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot interfaces.
Routing Protocols and Policies Command Reference	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot routing protocols and policies, including firewall filters.
System Basics and Services Command Reference	Describes the JUNOS software operational mode commands you use to monitor and troubleshoot system basics, including commands for real-time monitoring and route (or path) tracing, system software management, and chassis management. Also describes commands for monitoring and troubleshooting services such as CoS, IP Security (IPSec), stateful firewalls, flow collection, and flow monitoring.
System Log Messages Reference	Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message.
J-Web User Guide	
J-Web Interface User Guide	Describes how to use the J-Web GUI to configure, monitor, and manage Juniper Networks routing platforms.
JUNOS API and Scripting Documentation	
JUNOScript API Guide	Describes how to use the JUNOScript application programming interface (API) to monitor and configure Juniper Networks routing platforms.
JUNOScript API Configuration Reference	Provides a reference page for the configuration tags in the JUNOScript API.
JUNOScript API Operational Reference	Provides a reference page for the operational tags in the JUNOScript API.
JUNOS Configuration Scripting Guide	Provides an overview, instructions for using, and examples of the commit script feature of the JUNOS software. This guide explains how to enforce custom configuration rules defined in scripts that run at commit time. This guide also explains how to use commit script macros to provide simplified aliases for frequently used configuration statements.
JUNOS Comprehensive Index and Glossary	
Comprehensive Index and Glossary	Provides a complete index of all JUNOS software books and the *JUNOScript API Guide.* Also provides a comprehensive glossary.
JUNOScope Documentation	
JUNOScope Software User Guide	Describes the JUNOScope software GUI, how to install and administer the software, and how to use the software to manage routing platform configuration files and monitor routing platform operations.

Document	Description
J-series Services Router Documentation	
J-series Services Router Getting Started Guide	Provides an overview, basic instructions, and specifications for J-series Services Routers. The guide explains how to prepare your site for installation, unpack and install the router and its components, install licenses, and establish basic connectivity.
J-series Services Router Configuration Guide	Explains how to configure the interfaces on J-series Services Routers for basic IP routing with standard routing protocols. The guide also shows how to configure VPNs, configure and manage multicast networks, and apply routing techniques such as policies, firewall filters, IPSec tunnels, and service classification for safer, more efficient routing.
J-series Services Router Administration Guide	Shows how to manage users and operations, monitor network performance, upgrade software, and diagnose common problems on J-series Services Routers.
M-series and T-series Hardware Documentation	
Hardware Guide	Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide.
PIC Guide	Describes the routing platform PICs. Each platform has its own PIC guide.
Release Notes	
JUNOS Release Notes	Summarize new features and known problems for a particular software release, provide corrections and updates to published JUNOS and JUNOScript manuals, provide information that might have been omitted from the manuals, and describe upgrade and downgrade procedures.
Hardware Release Notes	Describe the available documentation for the routing platform and the supported PICs, and summarize known problems with the hardware and accompanying software. Each platform has its own release notes.
JUNOScope Software Release Notes	Contain corrections and updates to the published JUNOScope manual, provide information that might have been omitted from the manual, and describe upgrade and downgrade procedures.
J-series Services Router Release Notes	Briefly describe the J-series Services Router features, identify known hardware problems, and provide upgrade and downgrade instructions.

Table 4: JUNOS Internet Software Network Operations Guides

Book	Description
JUNOS Internet Software for M-series and T-series Routing Platforms Network Operations Guides	
Baseline	Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling JUNOS software, gathering basic system management information, verifying your network topology, and searching log messages.
Interfaces	Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms.
MPLS	Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routers in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network.

Book	Description
MPLS Log Reference	Describes MPLS status and error messages that appear in the output of the show mpls lsp extensive command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network.
Hardware	Describes tasks for monitoring M-series and T-series routing platforms.

Documentation Feedback

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

- Document name
- Document part number
- Page number
- Software release version

Requesting 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 (from the United States, Canada, or Mexico) or 1-408-745-9500 (from elsewhere).

Part 1

Overview

Chapter 1
JUNOS Software Overview

The JUNOS software runs on the router’s Routing Engine. It consists of software processes that support Internet routing protocols, control the router’s interfaces and the router chassis itself, and allow router system management. All these processes run on top of a kernel that enables communication among all the processes and has a direct link to the Packet Forwarding Engine software. You use the JUNOS software to configure the routing protocols that should run on the router and to configure properties of the router’s interfaces. Afterward, you use the JUNOS software to monitor the router and to troubleshoot protocol and network connectivity problems. For more information about monitoring the router and troubleshooting problems, see the *JUNOS System Basics and Services Command Reference*, the *JUNOS Interfaces Command Reference*, and the *JUNOS Routing Protocols and Policies Command Reference*.

This chapter discusses the following topics:

Routing Engine Software Components

The Routing Engine software consists of several software processes that control router functionality and a kernel that provides the communication among all the processes (see Figure 1 on page 45). This section describes the Routing Engine components:

- Routing Protocol Process on page 4
- VPNs on page 9
- Interface Process on page 9
- Chassis Process on page 9
- SNMP and MIB II Processes on page 10
- Management Process on page 10
- Routing Engine Kernel on page 10

For information about Routing Engine software components and Routing Engine functions in a routing matrix, see the *TX Matrix Platform Hardware Guide*.

Routing Protocol Process

The routing protocol process controls the routing protocols that run on the router. It starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols into common tables. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements routing policy, which allows you to control the routing information that is transferred between the routing protocols and the routing table. Using routing policy, you can filter routing information so that only some of it is transferred, and you also can set properties associated with the routes.

This section discusses the following topics:

- IPv4 Routing Protocols on page 5
- IPv6 Routing Protocols on page 6
- Routing and Forwarding Tables on page 7
- Routing Policy on page 8

IPv4 Routing Protocols

The JUNOS software implements full IP routing functionality, providing support for IP version 4 (IPv4). The routing protocols are fully interoperable with existing IP routing protocols, and they have been developed to provide the scale and control necessary for the Internet core.

The software provides the following routing and Multiprotocol Label Switching (MPLS) applications protocols:

- Unicast routing protocols:
  - BGP—Border Gateway Protocol, version 4, is an exterior gateway protocol (EGP) that guarantees loop-free exchange of routing information between routing domains (also called autonomous systems). BGP, in conjunction with JUNOS routing policy, provides a system of administrative checks and balances that can be used to implement peering and transit agreements.
  - ICMP—Internet Control Message Protocol router discovery allows hosts to discover the addresses of operational routers on the subnet.
  - IS-IS—Intermediate System-to-Intermediate System is a link-state interior gateway protocol (IGP) for IP networks that uses the shortest-path-first (SPF) algorithm, which also is referred to as the Dijkstra algorithm, to determine routes. The JUNOS IS-IS software is a new and complete implementation of the protocol, addressing issues of scale, convergence, and resilience.
  - OSPF—Open Shortest Path First, version 2, is an IGP that was developed for IP networks by the Internet Engineering Task Force (IETF). OSPF is a link-state protocol that makes routing decisions based on the SPF algorithm. The JUNOS OSPF software is a new and complete implementation of the protocol, addressing issues of scale, convergence, and resilience.
  - RIP—Routing Information Protocol, version 2, is an IGP for IP networks based on the Bellman-Ford algorithm. RIP is a distance-vector protocol. RIP dynamically routes packets between a subscriber and a service provider without the subscriber having to configure BGP or to participate in the service provider's IGP discovery process.
- Multicast routing protocols:
  - DVMRP—Distance Vector Multicast Routing Protocol is a dense-mode (flood-and-prune) multicast routing protocol.
  - IGMP—Internet Group Management Protocol, versions 1 and 2, is used to manage membership in multicast groups.
  - MSDP—Multicast Source Discovery Protocol allows multiple Protocol Independent Multicast (PIM) sparse mode domains to be joined. A rendezvous point (RP) in a PIM sparse mode domain has a peer relationship with an RP in another domain, enabling it to discover multicast sources from other domains.

  - PIM sparse mode and dense mode—Protocol-Independent Multicast is a multicast routing protocol. PIM sparse mode routes to multicast groups that might span wide-area and interdomain internets. PIM dense mode is a flood-and-prune protocol.

  - SAP/SDP—Session Announcement Protocol and Session Description Protocol handle conference session announcements.

- MPLS applications protocols:

  - LDP—The Label Distribution Protocol provides a mechanism for distributing labels in nontraffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data-link layer switched paths. LSPs created by LDP can also traverse LSPs created by the Resource Reservation Protocol (RSVP).

  - MPLS—Multiprotocol Label Switching, formerly known as tag switching, allows you to manually or dynamically configure LSPs through a network. It lets you direct traffic through particular paths rather than rely on the IGP's least-cost algorithm to choose a path.

  - RSVP—The Resource Reservation Protocol, version 1, provides a mechanism for engineering network traffic patterns that is independent of the shortest path decided upon by a routing protocol. RSVP itself is not a routing protocol; it operates with current and future unicast and multicast routing protocols. The primary purpose of the JUNOS RSVP software is to support dynamic signaling for MPLS LSPs.

IPv6 Routing Protocols

The JUNOS software implements IP routing functionality, providing support for IP version 6 (IPv6). The routing protocols have been developed to provide the scale and control necessary for the Internet core.

The software supports the following unicast routing protocols:

- BGP—Border Gateway Protocol version 4, is an EGP that guarantees loop-free exchange of routing information between routing domains (also called autonomous systems). BGP, in conjunction with JUNOS routing policies, provides a system of administrative checks and balances that can be used to implement peering and transit agreements.

- ICMP—Internet Control Message Protocol router discovery allows hosts to discover the addresses of operational routers on the subnet.

- IS-IS—Intermediate System-to-Intermediate System is a link-state IGP for IP networks that uses the SPF algorithm, which also is referred to as the Dijkstra algorithm, to determine routes. The JUNOS software supports a new and complete implementation of the protocol, addressing issues of scale, convergence, and resilience.

- OSPF version 3 (OSPFv3) supports IPv6. The fundamental mechanisms of OSPF such as flooding, designated router (DR) election, area-based topologies, and the SPF calculations remain unchanged. Some differences exist either due to changes in protocol semantics between IPv4 and IPv6, or to handle the increased address size of IPv6.

- RIP—Routing Information Protocol version 2 is an IGP for IP networks based on the Bellman-Ford algorithm. RIP is a distance-vector protocol. RIP dynamically routes packets between a subscriber and a service provider without the subscriber having to configure BGP or to participate in the service provider's IGP discovery process.

Routing and Forwarding Tables

A major function of the JUNOS routing protocol process is to maintain the Routing Engine's routing tables and from these tables determine the active routes to network destinations. The routing protocol process then installs these routes into the Routing Engine's forwarding table. The JUNOS kernel then copies this forwarding table to the Packet Forwarding Engine. Refer to Figure 1 on page 45 for an illustration of the interrelationships between the routing and forwarding tables.

The routing protocol process maintains multiple routing tables. By default, it maintains the following three routing tables. You can configure additional routing tables to suit your requirements.

- Unicast routing table—Stores routing information for all unicast routing protocols running on the router. BGP, IS-IS, OSPF, and RIP all store their routing information in this routing table. You can configure additional routes, such as static routes, to be included in this routing table. BGP, IS-IS, OSPF, and RIP use the routes in this routing table when advertising routing information to their neighbors.

- Multicast routing table (cache)—Stores routing information for all the running multicast protocols. DVMRP and PIM both store their routing information in this routing table, and you can configure additional routes to be included in this routing table.

- MPLS routing table—Stores MPLS path and label information.

With each routing table, the routing protocol process uses the collected routing information to determine active routes to network destinations.

For unicast routes, the routing protocol process determines active routes by choosing the most preferred route, which is the route with the lowest preference value. By default, the route's preference value is simply a function of how the routing protocol process learned about the route. You can modify the default preference value using routing policy and with software configuration parameters.

For multicast traffic, the routing protocol process determines active routes based on traffic flow and other parameters specified by the multicast routing protocol algorithms. The routing protocol process then installs one or more active routes to each network destination into the Routing Engine's forwarding table.

Routing Policy

By default, all routing protocols place their routes into the routing table. When advertising routes, the routing protocols by default advertise only a limited set of routes from the routing table. Specifically, each routing protocol exports only the active routes that were learned by that protocol. In addition, the IGPs (IS-IS, OSPF, and RIP) export the direct (interface) routes for the interfaces on which the protocol is explicitly configured.

You can control the routes that a protocol places into each table and the routes from that table that the protocol advertises. You do this by defining one or more routing policies and then applying them to the specific routing protocol.

Routing policies applied when the routing protocol places routes into the routing table are referred to as *import policies* because the routes are being imported into the routing table. Policies applied when the routing protocol is advertising routes that are in the routing table are referred to as *export policies* because the routes are being exported from the routing table. In other words, the terms *import* and *export* are used with respect to the routing table.

Routing policy allows you to control (filter) which routes a routing protocol imports into the routing table and which routes a routing protocol exports from the routing table. Routing policy also allows you to set the information associated with a route as it is being imported into or exported from the routing table. Filtering imported routes allows you to control the routes used to determine active routes. Filtering routes being exported from the routing table allows you to control the routes that a protocol advertises to its neighbors.

You implement routing policy by defining policies. A policy specifies the conditions to use to match a route and the action to perform on the route when a match occurs. For example, when a routing table imports routing information from a routing protocol, a routing policy might modify the route’s preference, mark the route with a color to identify it and allow it to be manipulated at a later time, or prevent the route from even being installed in a routing table. When exporting routes from a routing table into a routing protocol, a policy might assign metric values, modify the BGP community information, tag the route with additional information, or prevent the route from being exported altogether. You also can define policies for redistributing the routes learned from one protocol into another protocol.

VPNs

The JUNOS software supports several types of virtual private networks (VPNs):

- Layer 2 VPNs—A Layer 2 VPN links a set of sites sharing common routing information, and whose connectivity is controlled by a collection of policies. A Layer 2 VPN is not aware of routes within a customer's network. It simply provides private links between a customer's sites over the service provider's existing public Internet backbone.

- Layer 3 VPNs—A Layer 3 VPN links a set of sites that share common routing information, and whose connectivity is controlled by a collection of policies. A Layer 3 VPN is aware of routes within a customer's network, requiring more configuration on the part of the service provider than a Layer 2 VPN. The sites that make up a Layer 3 VPN are connected over a service provider's existing public Internet backbone.

- Inter-provider VPNs—An inter-provider VPN supplies connectivity between two VPNs in separate autonomous systems (ASs). This functionality could be used by a VPN customer with connections to several various Internet service providers (ISPs), or different connections to the same ISP in various geographic regions.

- Carrier-of-carrier VPNs—Carrier-of-carrier VPNs allow a VPN service provider to supply VPN service to a customer who is also a service provider. The latter service provider supplies Internet or VPN service to an end customer.

Interface Process

The JUNOS interface process allows you to configure and control the physical interface devices and logical interfaces present in a router. You can configure various interface properties such as the interface location (that is, which slot the Flexible PIC Concentrator [FPC] is installed in and which location on the FPC the Physical Interface Card [PIC] is installed in), the interface encapsulation, and interface-specific properties. You can configure the interfaces that currently are present in the router, as well as interfaces that currently are not present but that you may be adding at a future time.

The JUNOS interface process communicates, through the JUNOS kernel, with the interface process in the Packet Forwarding Engine, thus enabling the JUNOS software to track the status and condition of the router's interfaces.

Chassis Process

The JUNOS chassis process allows you to configure and control the properties of the router, including conditions that trigger alarms. The chassis daemon (chassisd) on the Routing Engine communicates directly with its peer processes running on the Packet Forwarding Engine.

SNMP and MIB II Processes

The JUNOS software supports the Simple Network Management Protocol (SNMP), which helps administrators monitor the state of a router. The software supports SNMP version 1, version 2 (also known as version 2c, or v2c), and version 3 (SNMPv3). The JUNOS implementation of SNMP does not include any of the security features that were originally included the IETF SNMP drafts but were later dropped because of the inability to standardize on a particular method. The SNMP software is controlled by the JUNOS SNMP and Management Information Base II (MIB II) processes, which consist of an SNMP master agent and various subagents. For information about SNMP, see the *JUNOS Network Management Configuration Guide*.

Management Process

Within the JUNOS software, a process-controlling process starts and monitors all the other software processes. It also starts the command-line interface (CLI), which is the primary tool you use to control and monitor the JUNOS software. This management process starts all the software processes and the CLI when the router boots. If a software process terminates, the management process attempts to restart it.

Routing Engine Kernel

The Routing Engine kernel provides the underlying infrastructure for all JUNOS software processes. In addition, it provides the link between the routing tables and the Routing Engine’s forwarding table. It is also responsible for all communication with the Packet Forwarding Engine, which includes keeping the Packet Forwarding Engine’s copy of the forwarding table synchronized with the master copy in the Routing Engine.

Software Installation Overview

The JUNOS software is preinstalled on the router. Once the router is powered on, it is ready to be configured. The primary copy of the software is installed on a nonrotating flash drive. Two backup copies are included, one on the router's rotating hard disk and a second on the removable media (either an LS-120 floppy disk [a 120-MB disk] or a PC card) that is shipped with the router.

When the router boots, it first attempts to start the software image from the removable media if one is installed in the router. If this fails, the router next tries the flash drive, then finally the hard disk. Normally, you want the router to boot from the flash drive.

To upgrade the software, you copy a set of software images over the network to the router's flash drive using SCP or another similar utility. The JUNOS software set consists of three images, one for the software processes, a second for the kernel, and the third for the Packet Forwarding Engine. You normally upgrade all images simultaneously.

User Interfaces

You can use two user interfaces to monitor, configure, troubleshoot, and manage the router—the JUNOS command-line interface (CLI) and the J-Web interface.

The CLI is a straightforward command interface. You type commands on a single line, and the commands are executed when you press the Enter key. The CLI provides command help and command completion. It also provides Emacs-style keyboard sequences that allow you to move around on a command line and scroll through a buffer that contains recently executed commands.

The J-Web graphical user interface (GUI) allows you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser with Hypertext Transfer Protocol (HTTP) or HTTP over Secure Sockets Layer (HTTPS) enabled. The J-Web interface provides access to configuration statements supported by the router, so you can configure it without using the CLI. The J-Web graphical user interface comes standard on J-series Services Routers and is an optional software package that you can install on M-series and T-series routers.

Ports for External Devices

The router provides three ports on the craft interface for connecting external management devices to the Routing Engine and the JUNOS software:

- Console port—Connects a system console using an RS-232 serial cable.
- Auxiliary port—Connects a laptop or modem using an RS-232 serial cable.
- Ethernet management port—Connects the Routing Engine to a management LAN (or any other device that plugs into an Ethernet connection) for out-of-band management of the router. The Ethernet port is 10/100 megabits-per-second (Mbps) autosensing and requires an RJ-45 connector.

Software Configuration Overview

To configure the JUNOS software, you specify a hierarchy of configuration statements that define the preferred software properties. You can configure all properties of the JUNOS software, including interfaces, general routing information, routing protocols, and user access, as well as some system hardware properties. After you have created a candidate configuration, you commit the configuration to be evaluated and activated by the JUNOS software.

This section discusses the following topics:

- Methods of Configuring the Software on page 12
- Configuring the Software on page 12
- Activating a Configuration on page 13

Methods of Configuring the Software

The following are some basic ways you can configure the JUNOS software:

- Create the configuration for the router interactively, working in the CLI on the router.
- Use the optional J-Web interface to configure the router.
- Load an ASCII file containing a router configuration that you created earlier, either on this system or on another system. You can then activate and run the configuration file as is, or you can edit it using the CLI and then activate it.

Configuring the Software

When you initially boot up a router, the system prompts you to log in. Log in as the user "root" (with no password) and configure a password for the user "root." Then configure the router's name, domain name, and the Internet address of at least one interface on the router.

After completing this initial minimal configuration, you can configure software properties. If you configure the software interactively using the CLI, you enter software configuration statements to create a candidate configuration that contains a hierarchy of statements. At any hierarchy level, you generally can enter statements in any order. While you are configuring the software, you can display all or portions of the candidate configuration, and you can insert or delete statements. Any changes you make affect only the candidate configuration, not the active configuration that is running on the router.

The configuration hierarchy logically groups related functions, which results in configuration statements that have a regular, consistent syntax. For example, you configure routing protocols, routing policies, interfaces, and SNMP management in their own separate portions of the configuration hierarchy.

At each level of the hierarchy, you can display a list of the statements available at that level, along with short descriptions of the statements' functions. To have the CLI complete the statement name if it is unambiguous or to provide a list of possible completions, you can type a partial statement name followed by a space or tab.

More than one user can edit a router's configuration simultaneously. All changes made by all users are visible to everyone editing the configuration.

Activating a Configuration

To have a candidate configuration take effect, you commit the changes. At this point, 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.

The CLI always maintains a copy of previously committed versions of the software configuration. If you need to return to a previous configuration, you can do this from within the CLI.

Configuration Commit Scripts

You can create and use scripts that run at commit time to enforce custom configuration rules. If a configuration breaks the custom rules, the script can generate actions that the JUNOS software performs. These actions include:

- Generating custom error messages
- Generating custom warning messages
- Generating custom system log messages
- Making changes to the configuration

Configuration commit scripts also enable you to create macros, which expand simplified custom aliases for frequently used configuration statements into standard JUNOS configuration statements. Commit scripts are written in Extensible Stylesheet Language Transformations (XSLT). For more information, see the *JUNOS Configuration Scripting Guide*.

Using Software Monitoring Tools

The primary method of monitoring and troubleshooting the software, routing protocols, network connectivity, and the router hardware is to enter commands from the CLI. The CLI enables you to display information in the routing tables and routing protocol-specific data, and to check network connectivity using ping and traceroute commands.

The JUNOS software includes SNMP software, which allows you to manage routers. The SNMP software consists of an SNMP master agent and a MIB II agent, and supports MIB II SNMP version 1 traps and version 2 notifications, SNMP version 1 Get and GetNext requests, and version 2 GetBulk requests.

The software also supports tracing and logging operations so that you can track events that occur in the router—both normal router operations and error conditions—and track the packets that are generated by or pass through the router. Logging operations use a syslog-like mechanism to record systemwide, high-level operations, such as interfaces' going up or down and users' logging in to or out of the router. Tracing operations record more detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions.

Router Security

Router security consists of three major elements: physical security of the router, operating system security, and security that can be effected through configuration. Physical security involves restricting access to the router. Exploits that can easily be prevented from remote locations are extremely difficult or impossible to prevent if an attacker can gain access to the router's management port or console. The inherent security of the JUNOS operating system also plays an important role in router security. The JUNOS software is extremely stable and robust. The JUNOS software also provides features to protect against attacks, allowing you to configure the router to minimize vulnerabilities.

In designing your router configuration, you can increase router security by *hardening* the configuration, using the JUNOS features to apply sound security policies. In this way, virtually any router configuration should be capable of secure operation. Likewise, misconfiguring the JUNOS software can increase router vulnerability.

This section discusses some JUNOS software features available to improve router security:

- JUNOS Default Settings on page 15
- Router Access on page 16
- User Authentication on page 16
- Specifying Plain-Text Passwords on page 17
- Routing Protocol Security Features on page 18
- Firewall Filters on page 19
- Auditing for Security on page 19

JUNOS Default Settings

Immediately after installation and configuration of a root account password, the JUNOS software presents a hardened target by virtue of its default software settings. The following are some common router security weaknesses that the JUNOS software addresses in the default software settings:

- The JUNOS software does not forward directed broadcast messages. Directed broadcast services send ping requests from a spoofed source address to a broadcast address and can be used to attack other Internet users. For example, if broadcast ping messages were allowed on the 200.0.0.0/24 network, a single ping request could result in up to 254 responses, all aimed at the supposed source of the ping. The result would be that the source actually becomes the victim of a denial-of-service (DoS) attack.

- Only console access to the router is enabled by default. Remote management access to the router and all management access protocols, including telnet, FTP, and SSH (Secure Shell), are disabled by default.

- The JUNOS software does not support the SNMP set capability for editing configuration data. While the software does support the SNMP set capability for monitoring and troubleshooting the network, this support exposes no known security issues. (You can configure the software to disable this SNMP set capability.)

- The JUNOS software ignores martian addresses that contain the following prefixes: 0.0.0.0/8, 127.0.0.0/8, 128.0.0.0/16, 191.255.0.0/16, 192.0.0.0/24, 223.255.55.0/24, and 240.0.0.0/4. Martian addresses are reserved host or network addresses about which all routing information should be ignored.

Router Access

When you first install the JUNOS software, all remote access to the router is disabled, thereby ensuring that remote access is possible only if deliberately enabled by an authorized user. You can establish remote communication with a router in one of the following ways:

- Out-of-band management—Allows connection to the router through an interface dedicated to router management. Juniper Networks routers support out-of-band management with a dedicated management Ethernet interface (fxp0), as well as EIA-232 console and auxiliary ports. The management Ethernet interface connects directly to the Routing Engine. No transit traffic is allowed through this interface, providing complete separation of customer and management traffic and ensuring that congestion or failures in the transit network do not affect the management of the router.

- Inband management—Allows connection to the routers using the same interfaces through which customer traffic flows. While this approach is simple and requires no dedicated management resources, it has some disadvantages:

  - Management flows and transit traffic flows are mixed together. Any attack traffic that is mixed with the normal traffic can affect the communication with the router.

  - The links between the router might not be totally trustworthy, leading to the possibility of wiretapping and replay attacks.

For management access to the router, the standard ways to communicate with the router from a remote console are with telnet and the Secure Shell (SSH). SSH provides secure encrypted communications and is therefore useful for inband router management. telnet provides unencrypted, and therefore less secure, access to the router.

User Authentication

On a route, you can create local user login accounts to control who can log into the router and the access privileges they have. A password, either an ssh key or a Message Digest 5 (MD5) password, is associated with each login account. To define access privileges, you create login classes in to which you group users with similar jobs or job functions. You use these classes to explicitly define what commands their users are and are not allowed to issue while logged in to the router.

The management of multiple routers by many different personnel can create a user account management problem. One solution is to use a central authentication service to simplify account management, creating and deleting user accounts only on a single, central server. A central authentication system also simplifies the use of one-time password systems such as SecureID, which offer protection against password sniffing and password replay attacks (attacks in which someone uses a captured password to pose as a router administrator).

The JUNOS software supports two protocols for central authentication of users on multiple routers—Remote Authentication Dial In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+). 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.

The JUNOS software also supports the following:

- Internet Protocol Security (IPSec). IPSec architecture provides a security suite for the IPv4 and IPv6 network layers. The suite provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. In addition to IPSec, the JUNOS software also supports the Internet Key Exchange (IKE), which defines mechanisms for key generation and exchange, and manages security associations (SAs). For more information about IPSec, see the *JUNOS Services Interfaces Configuration Guide* and “Security Services” on page 709.

- MD5 authentication of MSDP peering sessions. This authentication provides protection against spoofed packets being introduced into a peering session. For more information about SNMPv3, see the *JUNOS Multicast Protocols Configuration Guide*.

- SNMPv3 authentication and encryption. SNMPv3 uses the user-based security model (USM) for message security and the view-based access control model (VACM) for access control. USM specifies authentication and encryption. VACM specifies access-control rules. For more information about SNMPv3, see the *JUNOS Network Management Configuration Guide*.

Specifying Plain-Text Passwords

The JUNOS software has special requirements when you create plain-text passwords on a routing platform. The default requirements are shown in Table 5.

Table 5: Special Requirements for Plain-Text Passwords

JUNOS	JUNOS-FIPS
The password must be between 6 and 128 characters long.	FIPS passwords must be between 10 and 20 characters in length
You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.	You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
Valid passwords must contain at least one change of case or character class.	Valid passwords must use at least three of the five defined character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters).

You can change the requirements for plain-text passwords. For more information, see “Configuring Special Requirements for Plain-Text Passwords” on page 381.

You can include the plain-text-password statement at the following hierarchy levels.

- [edit system diag-port-authentication]
- [edit system pic-console-authentication]
- [edit system root-authentication]
- [edit system login user *username* authentication]

Table 6 lists error messages that appear when you enter an invalid plain-text password.

Table 6: Plain-Text Password Error Messages

Error message	Problem with password
The minimum password length is *number.* (*number* is the default length configured.)	Too few characters; for example, abC.
Require additional changes of case, numbers or punctuation	Does not include the required changes of case, numbers or special characters; for example, abcdefg.
Passwords are not equal; aborting	Does not match the original password.

For more information about how to create plain-text passwords, see "Configuring the Root Password" on page 379, "Configuring User Accounts" on page 413, and "Configuring the Password on the Diagnostics Port" on page 482.

Routing Protocol Security Features

The main task of a router is to forward user traffic toward its intended destination based on the information in the router's routing and forwarding tables. You can configure routing policies that define the flows of routing information through the network, controlling which routes the routing protocols place in the routing tables and which routes they advertise from the tables. You can also use routing policies to change specific route characteristics, change the BGP route flap-damping values, perform per-packet load balancing, and enable class of service (CoS).

Attackers can send forged protocol packets to a router with the intent of changing or corrupting the contents of its routing table or other databases, which in turn could degrade the functionality of the router. To prevent such attacks, you must ensure that routers form routing protocol peering or neighboring relationships with trusted peers. One way to do this is by authenticating routing protocol messages. The JUNOS BGP, IS-IS, OSPF, RIP, and RSVP protocols support HMAC-MD5 authentication, which uses a secret key combined with the data being protected to compute a hash. When the protocols send messages, the computed hash is transmitted with the data. The receiver uses the matching key to validate the message hash.

The JUNOS software supports the IPSec security suite for the IPv4 and IPv6 network layers. The suite provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. The JUNOS software also supports IKE, which defines mechanisms for key generation and exchange, and manages SAs.

Firewall Filters

Firewall filters allow you to control packets transiting the router to a network destination and packets destined for and sent by the router. You can configure firewall filters to control which data packets are accepted on and transmitted from the physical interfaces, and which local packets are transmitted from the physical interfaces and to the Routing Engine. Firewall filters provide a means of protecting your router from excessive traffic. Firewall filters that control local packets can also protect your router from external aggressions, such as DoS attacks.

To protect the Routing Engine, you can configure a firewall filter only on the router's loopback interface. Adding or modifying filters for each interface on the router is not necessary. You can design firewall filters to protect against ICMP and Transmission Control Protocol (TCP) connection request (SYN) floods and to rate-limit traffic being sent to the Routing Engine.

Auditing for Security

The JUNOS software logs significant events that occur on the router and within the network. Although the logging of events and actions does not increase router security, you can use the system logs to monitor the effectiveness of your security policies and router configurations. You can also use the logs when reacting to a continued and deliberate attack as a means of identifying the source address, router, or port of the attacker's traffic. You can configure the logging of different levels of events, from only critical events to all events, including informational events. You can then inspect the contents of the system log files either in real time or at a later time.

Debugging and troubleshooting is much easier when the timestamps in the system log files of all routers are synchronized, because events that span the network might be correlated with synchronous entries in multiple logs. The JUNOS software supports the Network Time Protocol (NTP), which you can enable on the router to synchronize the system clocks of routers and other networking equipment. By default, NTP operates in an unauthenticated mode. You can configure various types of authentication, including an HMAC-MD5 scheme.

FIPS 140-2 Security Compliance

For advanced network security, a special version of JUNOS, called JUNOS-FIPS 140-2, is available. JUNOS-FIPS 140-2 provides customers with software tools to configure a network of Juniper Networks routers in a FIPS environment. FIPS support includes:

- Upgrade package to convert JUNOS to JUNOS-FIPS 140-2
- Revised installation and configuration procedures
- Enhanced password creation and encryption
- Enforced security for remote access
- FIPS user roles (Crypto Officer, User, and Maintenance)
- FIPS-specific system logging and error messages
- IPSec configuration for Routing Engine–to–Routing Engine communications

Where appropriate, distinctive JUNOS-FIPS features are noted in this book. For detailed guidelines on how installation and configuration procedures differ between JUNOS and JUNOS-FIPS 140-2, see the *JUNOS-FIPS Configuration Guide*

Supported Software Standards

This section lists the standards supported by the JUNOS software:

- Supported Internet RFCs and Drafts on page 21
- Supported ISO Standards on page 40
- Supported SDH and SONET Standards on page 40
- Other Supported Standards on page 41

To access Internet RFCs and drafts, go to the IETF Web site: http://www.ietf.org.

Supported Internet RFCs and Drafts

This section lists the supported Internet RFCs and drafts:

- Asynchronous Transfer Mode (ATM) on page 22
- BGP on page 22
- Challenge Handshake Authentication Protocol (CHAP) on page 23
- Dynamic Host Configuration Protocol (DHCP) on page 23
- Firewall Filters on page 24
- Frame Relay on page 24
- Generalized MPLS (GMPLS) on page 24
- Generalized Routing Encapsulation (GRE) and IP-IP Encapsulation on page 25
- Integrated Local Management Interface (ILMI) on page 25
- IP Multicast on page 25
- IPSec and IKE on page 26
- IPv6 on page 27
- IS-IS on page 28
- LDP on page 29
- Link Management Protocol (LMP) on page 29
- Layer 2 Tunneling Protocol (L2TP) on page 29
- MIBs on page 30
- MPLS on page 34
- Network Address Translation (NAT) on page 35
- OSPF on page 35
- Point-to-Point Protocol (PPP) on page 36
- RIP on page 36
- RSVP on page 37
- Secure Sockets Layer (SSL) on page 37
- TCP/IPv4 on page 38

- Voice Services on page 39
- VPNs on page 39

Asynchronous Transfer Mode (ATM)

- RFC 1483, *Multiprotocol Encapsulation over ATM Adaptation Layer 5* (routed protocol data units only)
- RFC 2225, *Classical IP and ARP over ATM* (responses only)
- RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5* (routed protocol data units and Ethernet bridged protocol data units only)
- Internet draft draft-martini-l2circuit-encap-mpls-07.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks* (expires December 2004)

  The JUNOS software has the following exceptions:

  - A packet with a sequence number of 0 is treated as out of sequence.
  - Any packet which does not have the next incremental sequence number is considered out of sequence.
  - When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.
- Internet draft draft-martini-l2circuit-trans-mpls-09.txt, *Transport of Layer 2 Frames Over MPLS*

BGP

- RFC 1771, *A Border Gateway Protocol 4 (BGP-4)*
- RFC 1772, *Application of the Border Gateway Protocol in the Internet*
- RFC 1965, *Autonomous System Confederations for BGP*
- RFC 1966, *BGP Route Reflection—An Alternative to Full-Mesh IBGP*
- RFC 1997, *BGP Communities Attribute*
- RFC 2270, *Using a Dedicated AS for Sites Homed to a Single Provider*
- RFC 2283, *Multiprotocol Extensions for BGP-4*
- RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*
- RFC 2439, *BGP Route Flap Damping*
- RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*
- RFC 2796, *BGP Route Reflection*

- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3065, *Autonomous System Confederations for BGP*
- RFC 3107, *Carrying Label Information in BGP-4*
- Internet draft draft-ramachandra-bgp-ext-communities-09.txt, *BGP Extended Communities Attribut*e
- Internet draft draft-ietf-l3vpn-rfc2547bis-03.txt, *BGP/MPLS IP VPNs*
- Internet draft draft-kato-bgp-ipv6-link-local-00.txt, *BGP4+ Peering Using IPv6 Link-local Address*
- Internet draft draft-ietf-idr-cap-neg-01.txt, *Capabilities Negotiation with BGP4* (expires February 1998)
- Internet draft draft-ietf-idr-restart-10.txt, *Graceful Restart Mechanism for BGP* (expires December 2004)
- Internet draft draft-ietf-mpls-bgp-mpls-restart-03.txt *Graceful Restart Mechanism for BGP with MPLS* (expires August 2004)
- Internet draft draft-ietf-ngtrans-bgp-tunnel-04.txt, *Connecting IPv6 Islands across IPv4 Clouds with BGP* (only multiprotocol-BPG [MP-BGP] over IPv4 approach) (expires July 2002)

Challenge Handshake Authentication Protocol (CHAP)

- RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

Dynamic Host Configuration Protocol (DHCP)

- RFC 1001, *Protocol standard for a NetBIOS service on a TCP/UDP transport: Concepts and methods*
- RFC 1002, *Protocol standard for a NetBIOS service on a TCP/UDP transport: Detailed specifications*
- RFC 1035, *Domain names - implementation and specification*
- RFC 1534, *Interoperation Between DHCP and BOOTP*
- RFC 1700, *Assigned Numbers*
- RFC 2131, *Dynamic Host Configuration Protocol* (The JUNOS software does not support DHCP over virtual local area networks [VLAN]-tagged interfaces.)
- RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
- RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6).* (The JUNOS software supports IPv4 address assignment but not IPv6 address assignment.)
- RFC 3397, *Dynamic Host Configuration Protocol (DHCP) Domain Search Option*

Firewall Filters

- RFC 2474, *Definition of the Differentiated Services (DS) Field*
- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB*
- RFC 2598, *An Expedited Forwarding PHB*

Frame Relay

- RFC 1490, *Multiprotocol Interconnect over Frame Relay*

Generalized MPLS (GMPLS)

- RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS)-Signaling Functional Description.* The JUNOS software supports only the following areas:
  - Generalized label request (only bandwidth encoding)
  - Generalized label (only suggested label)
  - Bidirectional LSPs (only upstream label)
  - Control channel separation
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, “Fault Handling”)
- Internet draft draft-ietf-ccamp-gmpls-sonet-sdh-08.txt, *GMPLS Extensions for SONET and SDH Control* (only SUKLM labels and SONET traffic parameters)
- Internet draft draft-ietf-mpls-generalized-rsvp-te-06.txt, *Generalized MPLS Signaling - RSVP-TE Extensions.* The JUNOS software supports only the following areas:
  - Generalized label request object
  - Generalized label object (only suggested labeled type)
  - Bidirectional LSPs (only upstream label)
  - Control channel separation (only IF-ID Hop object and IF-ID ErrSpec object)
  - New addressing for Path and PathTear messages
- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-09.txt, *OSPF Extensions in Support of Generalized MPLS* (interface switching only)
- Internet draft draft-ietf-mpls-bundle-04.txt, *Link Bundling in MPLS Traffic Engineering*
- Internet draft draft-ietf-ccamp-gmpls-routing-06.txt, *Routing Extensions in Support of Generalized MPLS*

Generalized Routing Encapsulation (GRE) and IP-IP Encapsulation

- RFC 1701, *Generic Routing Encapsulation (GRE)*
- RFC 1702, *Generic Routing Encapsulation over IPv4 Networks*
- RFC 2003, *IP Encapsulation within IP*
- RFC 2890, *Key and Sequence Number Extensions to GRE.* The JUNOS software supports the key field, but not the sequence number field.

Integrated Local Management Interface (ILMI)

- ILMI Management Information Base MIB (only the atmfMYIPNmAddress and atmfPortMyIfname objects). For more information about the ILMI MIB, see the *JUNOS Network Management Configuration Guide* and the ATM Forum at http://www.atmforum.com/.

IP Multicast

- RFC 1112, *Host Extensions for IP Multicasting* (defines IGMP Version 1)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2327, *SDP: Session Description Protocol*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- RFC 2365, *Administratively Scoped IP Multicast*
- RFC 2547, *BGP/MPLS VPNs*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 2974, *Session Announcement Protocol*
- RFC 3208, *PGM Reliable Transport Protocol Specification*
- RFC 3376, *Internet Group Management Protocol, Version 3* (source-specific multicast [SSM] include mode only)
- RFC 3446, *Anycast Rendezvous Point (RP) Mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
- RFC 3569, *An Overview of Source-Specific Multicast (SSM)*
- RFC 3590, *Source Address Selection for Multicast Listener Discovery Protocol* (SSM include mode only)
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- Internet draft draft-ietf-pim-sm-bsr-03.txt, *Bootstrap Router (BSR) Mechanism for PIM Sparse Mode* (expired August 2003)

- Internet draft draft-ietf-idmr-dvmrp-v3-11.txt, *Distance Vector Multicast Routing Protocol* (expired April 2004)
- Internet draft draft-rosen-vpn-mcast-06.txt, *Multicast in MPLS/BGP VPNs,* Option 2 (expired April 2004)
- Internet draft draft-ietf-pim-sm-v2-new-10.txt, *Protocol Independent Multicast—Sparse Mode (PIM-SM): Protocol Specification (Revised)* (expires January 2005)
- Internet draft draft-ietf-pim-dm-new-v2-05.txt, *Protocol Independent Multicast Version 2 Dense Mode Specification* (expires December 2004)
- Internet draft draft-ietf-ssm-arch-06.txt, *Source-Specific Multicast for IP* (expires March 2005)
- Internet draft draft-ietf-mboned-ssm232-08.txt, *Source-Specific Protocol Independent Multicast in 232/8* (expires September 2004)
- Internet draft draft-holbrook-idmr-igmpv3-ssm-07.txt, *Using IGMPv3 and MLDv2 for Source-Specific Multicast* (expires December 2004)
- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs* (expires December 2004)

IPSec and IKE

- RFC 2085, *HMAC-MD5 IP Authentication with Replay Prevention*
- RFC 2401, *Security Architecture for the Internet Protocol*
- RFC 2402, *IP Authentication Header*
- RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
- RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
- RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
- RFC 2406, *IP Encapsulation Security Payload*
- RFC 2407, *The Internet IP Security Domain of Interpretation for ISAKMP*
- RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
- RFC 2409, *Internet Key Exchange*
- RFC 2410, *The NULL Encryption Algorithm and Its Use With IPSec*
- RFC 2412, *The OAKLEY Key Determination Protocol*
- RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*

IPv6

- RFC 1157, *A Simple Network Management Protocol (SNMP)*
- RFC 1213, *Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II.* The JUNOS software supports the following areas:
  - MIB II and its SNMP version 2 derivatives, including:
    - Statistics counters
    - IP (except for ipRouteTable, which has been replaced by ipCidrRouteTable [RFC 2096, *IP Forwarding Table MIB*])
    - SNMP management
    - Interface management
  - SNMP version 1 Get and GetNext requests and version 2 GetBulk requests.
  - JUNOS-specific secured access list
  - Master configuration keywords
  - Reconfigurations upon SIGHUP
- RFC 1215, *A Convention for Defining Traps for Use with the SNMP* (only MIB II SNMP version 1 traps and version 2 notifications)
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 1771, *A Border Gateway Protocol 4 (BGP-4)*
- RFC 1772, *Application of the Border Gateway Protocol in the Internet*
- RFC 1901, *Introduction to Community-based SNMPv2*
- RFC 1902, *Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC 1981, *Path MTU Discovery for IP version 6*
- RFC 2080, *RIPng for IPv6*
- RFC 2081, *RIPng Protocol Applicability Statement*
- RFC 2283, *Multiprotocol Extensions for BGP-4*
- RFC 2373, *IP Version 6 Addressing Architecture*
- RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*

- RFC 2461, *Neighbor Discovery for IP Version 6 (IPv6)*
- RFC 2462, *IPv6 Stateless Address Autoconfiguration*
- RFC 2463, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
- RFC 2464*, Transmission of IPv6 Packets over Ethernet Networks*
- RFC 2472, *IP Version 6 over PPP*
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*
- RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
- RFC 2740, *OSPF for IPv6*
- RFC 2878, *PPP Bridging Control Protocol*
- RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers*
- Internet draft draft-kato-bgp-ipv6-link-local-00.txt, *BGP4+ Peering Using IPv6 Link-local Address* (expires April 2002)
- Internet draft draft-ietf-dhc-dhcpv6-16.txt, *Dynamic Host Configuration Protocol for Ipv6* (expires May 2001)
- Internet draft draft-ietf-ngtrans-bgp-tunnel-04.txt, *Connecting IPv6 Islands across IPv4 Clouds with BGP* (only MP-BGP over IPv4 approach)
- Internet draft draft-ietf-isis-ipv6-02.txt, *Routing IPv6 with IS-IS*

IS-IS

- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*
- RFC 2763, *Dynamic Hostname Exchange Mechanism for IS-IS*
- RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*
- RFC 2973, *IS-IS Mesh Groups*
- Internet draft draft-ietf-katz-ward-bfd-02.txt, *Bidirectional Forwarding Detection* (except the transmission of echo packets) (expires January 2005)
- Internet draft draft-ietf-isis-hmac-03.txt, *IS-IS Cryptographic Authentication* (expires January 2002)
- Internet draft draft-ietf-isis-traffic-04.txt, *IS-IS Extensions for Traffic Engineering* (expires February 2002)

- Internet draft draft-ietf-isis-wg-multi-topology-04.txt, *M-ISIS: Multi Topology (MT) Routing in IS-IS* (expires December 2004)

- Internet draft draft-ietf-isis-snp-checksum-02.txt, *Optional Checksums for IS-IS* (expires 2001)

- Internet draft draft-ietf-isis-igp-p2p-over-lan-03.txt, *Point-to-point operation over LAN in link-state routing protocols* (expires February 2004)

- Internet draft draft-ietf-isis-restart-05.txt, *Restart signaling for IS-IS* (expires February 2002)

- Internet draft draft-ietf-isis-ipv6-02.txt, *Routing IPv6 with IS-IS* (expires September 2001)

- Internet draft draft-ietf-isis-3way-03.txt, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies* (expires January 2001)

LDP

- RFC 3036, *LDP Specification*. The JUNOS software does not support the following areas:

  - Label distribution control mode: ordered (only independent)

  - Label retention mode: liberal (only conservative)

  - Label advertisement mode: downstream unsolicited (only downstream on demand)

  - Loop detection

  - Constraint-Based Routed LDP (CR-LDP)

- Internet draft draft-ietf-mpls-ldp-restart-06.txt, *Graceful Restart Mechanism for LDP*

Link Management Protocol (LMP)

- Internet draft draft-ietf-ccamp-lmp-09.txt, *Link Management Protocol (LMP)*

Layer 2 Tunneling Protocol (L2TP)

- RFC 2661, *Layer Two Tunneling Protocol "L2TP"*

- RFC 2866, *Radius Accounting*

MIBs

- IEEE, 802.3ad, *Aggregation of Multiple Link Segments*

  Only the following are supported:

  - dot3adAggPortTable, dot3adAggPortListTable, dot3adAggTable, and dot3adAggPortStatsTable
  - dot3adAggPortDebugTable (only dot3adAggPortDebugRxState, dot3adAggPortDebugMuxState, dot3adAggPortDebugActorSyncTransitionCount, dot3adAggPortDebugPartnerSyncTransitionCount, dot3adAggPortDebugActorChangeCount, and dot3adAggPortDebugPartnerChangeCount)
  - dot3adTablesLastChanged

- RFC 1155, *Structure and Identification of Management Information for TCP/IP-based Internets*
- RFC 1157, *A Simple Network Management Protocol (SNMP)*
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments* (only isisSystem, isisMANAreaAddr, isisAreaAddr, isisSysProtSupp, isisSummAddr, isisCirc, isisCircLevel, isisPacketCount, isisISAdj, isisISAdjAreaAddr, isisAdjIPAddr, isisISAdjProtSupp, isisRa, and isisIPRA)
- RFC 1212, *Concise MIB Definitions*
- RFC 1213, *Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II.* The JUNOS software supports the following areas:

  - MIB II and its SNMP version 2 derivatives, including:

    - Statistics counters
    - IP (except for ipRouteTable, which has been replaced by ipCidrRouteTable [RFC 2096, *IP Forwarding Table MIB*])
    - SNMP management
    - Interface management

  - SNMP version 1 Get and GetNext requests and version 2 GetBulk requests.
  - JUNOS-specific secured access list
  - Master configuration keywords
  - Reconfigurations upon SIGHUP

- RFC 1215, *A convention for Defining Traps for Use with the SNMP* (only MIB II SNMP version 1 traps and version 2 notifications)

- RFC 1406, *Definitions of Managed Objects for the DS1 and E1 Interface Types* (T1 MIB is supported)

- RFC 1407, *Definitions of Managed Objects for the DS3/E3 Interface Type* (T3 MIB is supported)

- RFC 1657, *Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2*

- RFC 1850, *OSPF Version 2 Management Information Base* (except for the ospfOriginateNewLsas and ospfRxNewLsas objects, the Host Table, and the traps ospfOriginateLSA, ospfLsdbOverflow, and ospfLsdbApproachingOverflow)

- RFC 1901, *Introduction to Community-based SNMPv2*

- RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*

- RFC 1907, *Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)*

- RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol using SMIv2*

- RFC 2012, *SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2*

- RFC 2013, *SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2*

- RFC 2096, *IP Forwarding Table MIB*

- RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

- RFC 2287, *Definitions of System-Level Managed Objects for Applications* (only sysApplInstallPkgTable, sysApplInstallElmtTable, sysApplElmtRunTable, and sysApplMapTable)

- RFC 2465, *Management Information Base for IP Version 6: Textual Conventions and General Group* (except IPv6 or ICMP version 6 [ICMPv6] statistics)

- RFC 2495, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types* (except for dsx1FarEndConfigTable, dsx1FarEndCurrentTable, dsx1FarEndIntervalTable, dsx1FarEndTotalTable, and dsx1FracTable)

- RFC 2496, *Definitions of Managed Objects for the DS3/E3 Interface Type* (except dsx3FarEndConfigTable, dsx3FarEndCurrentTable, dsx3FarEndIntervalTable, dsx3FarEndTotalTable, and dsx3FracTable)

- RFC 2515, *Definitions of Managed Objects for ATM Management* (except atmVpCrossConnectTable, atmVcCrossConnectTable, and aal5VccTable)

- RFC 2558, *Definitions of Managed Objects for the SONET/SDH Interface Type*

- RFC 2571, *An Architecture for Describing SNMP Management Frameworks* (read-only access)

- RFC 2576, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

- RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

- RFC 2579, *Textual Conventions for SMIv2*

- RFC 2662. *Definitions of Managed Objects for ADSL Lines* (J-series Services Routers; all MIB tables, objects, and traps applicable for the ADSL ATU-R agent)

- RFC 2665, *Definitions of Managed Objects for the Ethernet-like Interface Types*

- RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol* (except row creation, set operation, and the object vrrpStatsPacketLengthErrors)

- RFC 2790, *Host Resources MIB*

  - Only the hrStorageTable. The file systems /, /config, /var, and /tmp will always return the same index number. When SNMP restarts, the index numbers for the remaining file systems might change.

  - Only the objects of the hrSystem and hrSWInstalled groups.

- RFC 2819, *Remote Network Monitoring Management Information Base* (the etherStatsTable for Ethernet interfaces only and the objects alarmTable, eventTable, and logTable)

- RFC 2863, *The Interfaces Group MIB*

- RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations* (only the objects pingCtlTable, pingResultsTable, pingProbeHistoryTable, pingMaxConcurrentRequests, traceRouteCtlTable, traceRouteResultsTable, traceRouteProbeHistoryTable, and traceRouteHopsTable)

- RFC 2932, *IPv4 Multicast Routing MIB*

- RFC 2933, *Internet Group Management Protocol (IGMP) MIB*

- RFC 2934, *Protocol Independent Multicast MIB for IPv4*

- RFC 3413, *Simple Network Management Protocol (SNMP) Applications* (except for the proxy MIB*)*

- RFC 3414, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

- RFC 3415, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

- RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures* (implemented under the Juniper Networks enterprise branch)

- RFC 3811, *Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management*

- RFC 3813, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)* (except mplsInterfacePerfTable, mplsInSegmentPerfTable, mplsOutSegmentPerfTable, mplsInSegmentMapTable, mplsXCUp, and mplsXCDown)

- RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

- *IANA iftype Textual Convention MIB*, Internet Assigned Numbers Authority (referenced by RFC 2233, available at ftp://ftp.isi.edu/mib/ianaiftype.mib)

- ESO Consortium MIB, which can be found at http://www.snmp.com/eso/

- Internet draft draft-ietf-idr-bgp4-mibv2-04.txt, *Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4), Second Version* (only jnxBgpM2PrefixInPrefixes, jnxBgpM2PrefixInPrefixesAccepted, and jnxBgpM2PrefixInPrefixesRejected objects) (June 12, 2004)

- Internet draft draft-reeder-snmpv3-usm-3desede-00.txt, *Extension to the User-Based Security Model (USM) to Support Triple-DES EDE in 'Outside' CBC Mode*

- Internet draft draft-ietf-ppvpn-mpls-vpn-mib-05.txt, *MPLS/BGP Virtual Private Network Management Information Base Using SMIv2* (only mplsVpnScalars, mplsVpnVrfTable, mplsVpnVrfPerfTable, and mplsVpnVrfRouteTargetTable)

- Internet draft draft-ietf-isis-wg-mib-16.txt, *Management Information Base for IS-IS* (only isisISAdjTable, isisISAdjAreaAddrTable, isisISAdjIPAddrTable, and isisISAdjProtSuppTable) (expires January 2005)

- Internet draft draft-ietf-msdp-mib-08.txt, *Multicast Source Discovery protocol MIB* (except msdpEstablished, msdpBackwardTransition, and msdpRequestsTable) (expires April 2004)

MPLS

- RFC 2205, *Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification*
- RFC 2209, *Resource ReSerVation Protocol (RSVP)—Version 1 Message Processing Rules*
- RFC 2210, *The Use of RSVP with IETF Integrated Services*
- RFC 2211, *Specification of the Controlled-Load Network Element Service*
- RFC 2215, *General Characterization Parameters for Integrated Service Network Elements*
- RFC 2216, *Network Element Service Specification Template*
- RFC 2702, *Requirements for Traffic Engineering Over MPLS*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3032, *MPLS Label Stack Encoding*
- RFC 3208, *PGM Reliable Transport Protocol Specification* (only the network element)
- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services (e-lsps only)*
- RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*
- RFC 3469, *Framework for Multi-Protocol Label Switching (MPLS)-based Recovery*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 3811, *Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management*
- RFC 3813, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)* (except mplsInterfacePerfTable, mplsInSegmentPerfTable, mplsOutSegmentPerfTable, mplsInSegmentMapTable, mplsXCUp, and mplsXCDown)
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels* (except node protection in facility backup)
- RFC 4124, *Protocol Extensions for Support of Differentiated-Service-aware MPLS Traffic Engineering*

- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- Internet draft draft-ietf-l3vpn-rfc2547bis-03.txt, *BGP/MPLS IP VPNs*
- Internet draft draft-ietf-mpls-lsp-ping-version-05.txt, *Detecting MPLS Data Plane Failures* (only the LDP IPv4 prefix type, length, and value [TLV], RSVP IPv4 Session Query TLV, and VPN IPv4 prefix TLV)
- Internet draft draft-raggarwa-mpls-p2mp-te-02.txt, *Establishing Point to Multipoint MPLS TE LSPs*, (except nonadjacent signaling for branch LSPs, make-before-break and fast reroute, and LSP hierarchy using point-to-point LSPs)
- Internet draft draft-ietf-mpls-icmp-01.txt, *ICMP Extensions for Multiprotocol Label Switching*
- Internet draft draft-kompella-ppvpn-l2vpn-00.txt, *MPLS-based Layer 2 VPNs*
- Internet draft draft-ietf-mpls-label-encaps-07.txt, *MPLS Label Stack Encoding*
- Internet draft draft-ietf-mpls-soft-preemption-00.txt, *MPLS Traffic Engineering Soft preemption*
- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-09.txt, *OSPF Extensions in Support of Generalized MPLS* (interface switching only)
- Internet draft draft-marques-ppvpn-ibgp-00.txt, *RFC2547bis networks using internal BGP as PE-CE*
- Internet draft draft-ietf-ccamp-gmpls-routing-0.6txt, *Routing Extensions in Support of Generalized MPLS* (interface switching only)
- Internet draft draft-martini-l2circuit-trans-mpls-14.txt, *Transport of Layer 2 Frames Over MPLS*

Network Address Translation (NAT)

- RFC 3022, *Traditional IP Network Address Translator (Traditional NAT)*

OSPF

- RFC 1587, *The OSPF NSSA Option*
- RFC 2328, *OSPF Version 2*
- RFC 2740, *OSPF for IPv6*
- RFC 3623, *OSPF Graceful Restart*
- Internet draft draft-ietf-katz-ward-bfd-02.txt, *Bidirectional Forwarding Detection* (except the transmission of echo packets)

- Internet draft draft-ietf-isis-igp-p2p-over-lan-03.txt, *Point-to-point operation over LAN in link-state routing protocols* (expires February 2004)

- Internet draft draft-katz-yeung-ospf-traffic-01.txt, *Traffic Engineering Extensions to OSPF* (expires April 2000)

- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-12.txt*, OSPF Extensions in Support of Generalized MPLS* (except link local/remote identifiers, link protection type, shared risk link group [SRLG], and implications of graceful restart) (expires April 2004)

Point-to-Point Protocol (PPP)

- RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

- RFC 1661, *The Point-to-Point Protocol (PPP)*

- RFC 1662, *PPP in HDLC-like Framing*

- RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

- RFC 1990, *The PPP Multilink Protocol (MP*)

- RFC 2509, *IP Header Compression over PPP*

- RFC 2615, *PPP over SONET/SDH*

- RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

  - Prefix elision is not supported.

  - The JUNOS software PPP implementation does not support the negotiation of address field compression and protocol field compression PPP NCP options. The software always send a full 4-byte PPP header.

RIP

- RFC 1058, *Routing Information Protocol*

- RFC 2082, *RIP-2 MD-5 Authentication*

- RFC 2453, *RIP Version 2*

RSVP

- RFC 2205, *Resource ReSerVation Protocol (RSVP), Version 1, Functional Specification*
- RFC 2209, *Resource ReSerVation Protocol (RSVP), Version 1, Message Processing Rules*
- RFC 2210, *The Use of RSVP with IETF Integrated Services*
- RFC 2211, *Specification of the Controlled-Load Network Element Service*
- RFC 2212, *Specification of Guaranteed Quality of Service*
- RFC 2215, *General Characterization Parameters for Integrated Service Network Elements*
- RFC 2216, *Network Element Service Specification Template*
- RFC 2747, *RSVP Cryptographic Authentication*
- RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
- RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*
- Internet draft draft-ietf-mpls-generalized-rsvp-te-09.txt, *Generalized MPLS Signaling - RSVP-TE Extensions* (fault handling only*)*
- Internet draft draft-ietf-mpls-rsvp-te-p2mp-01.txt, *Extensions to RSVP-TE for Point to Multipoint TE LSPs* (expires June 2005)

Secure Sockets Layer (SSL)

- RFC 1319, *The MD2 Message-Digest Algorithm*
- RFC 1321, *The MD5 Message-Digest Algorithm*
- RFC 2246, *The TLS Protocol Version 1.0*
- RFC 2459, *Internet X.509 Public Key Infrastructure Certificate and CRL Profile*

TCP/IPv4

- RFC 768, *User Datagram Protocol*
- RFC 791, *Internet Protocol*
- RFC 792, *Internet Control Message Protocol*
- RFC 793, *Transmission Control Protocol*
- RFC 826, *Ethernet Address Resolution Protocol*
- RFC 854, *Telnet Protocol Specification*
- RFC 862, *Echo Protocol*
- RFC 863, *Discard Protocol*
- RFC 896, *Congestion Control in IP/TCP Internetworks*
- RFC 919, *Broadcasting Internet Datagrams*
- RFC 922, *Broadcasting Internet Datagrams in the Presence of Subnets*
- RFC 959, *File Transfer Protocol*
- RFC 1027, *Using ARP to Implement Transparent Subnet Gateways*
- RFC 1042, *Standard for the Transmission of IP Datagrams over IEEE 802 Networks*
- RFC 1157, *Simple Network Management Protocol (SNMP)*
- RFC 1166, *Internet Numbers*
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 1256, *ICMP Router Discovery Messages*
- RFC 1305, *Network Time Protocol (Version 3) Specification, Implementation, and Analysis*
- RFC 1519, *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*
- RFC 1812, *Requirements for IP Version 4 Routers*
- RFC 1948, *Defending Against Sequence Number Attacks*
- RFC 2338, *Virtual Router Redundancy Protocol*
- RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*

Voice Services

- RFC 2508, *Compressing IP/UDP/RTP Headers for Low-Speed Serial Links*
- RFC 2509, *IP Header Compression over PPP*

VPNs

- RFC 1918, *Address Allocation for Private Internets*
- RFC 2283, *Multiprotocol Extensions for BGP4*
- RFC 2547, *BGP/MPLS VPNs*
- RFC 3107, *Carrying Label Information in BGP-4*
- Internet draft draft-ietf-l3vpn-rfc2547bis-03.txt, *BGP/MPLS IP VPNs*
- Internet draft draft-kompella-l2ppvpn-version.txt, *MPLS based Layer 2 VPNs*
- Internet draft draft-marques-ppvpn-rt-constrain-01.txt, *Constrained VPN Route Distribution*
- Internet draft draft-rosen-rfc2547bis, *BGP/MPLS VPNs*
- Internet draft draft-martini-l2circuit-encap-mpls-07.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*

  The JUNOS software has the following exceptions:

  - A packet with a sequence number of 0 is treated as out of sequence.
  - Any packet that does not have the next incremental sequence number is considered out of sequence.
  - When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.
- Internet draft draft-ietf-ppvpn-vpls-bgp-00.txt, *Virtual Private LAN Service*
- Internet draft draft-ietf-l2vpn-vpls-bgp-05.txt, *Virtual Private LAN Service*
- Internet draft draft-marques-ppvpn-rt-constrain-01.txt, *Constrained VPN Route Distribution*
- Internet draft draft-martini-l2circuit-trans-mpls-14.txt, *Transport of Layer 2 Frames Over MPLS*

Supported ISO Standards

IS-IS

- ISO/IEC 10589, *Information technology, Telecommunications and information exchange between systems, Intermediate system to intermediate system intradomain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)*

Supported SDH and SONET Standards

- ANSI T1.105, *Synchronous Optical Network (SONET) Basic Description Including Multiplex Structures, Rates, and Formats*
- ANSI T1.105.02, *Synchronous Optical Network (SONET) Payload Mappings*
- ANSI T1.105.06, *SONET: Physical Layer Specifications*
- GR-253-CORE, *SONET Transport Systems: Common Generic Criteria*
- GR-499-CORE, *Transport System Generic Requirements (TSGR): Common Requirements*
- GR-1377-CORE, *SONET OC-192 Transport System Generic Criteria*
- ITU-T Recommendation G.691, *Optical interfaces for single channel SDH systems with optical amplifiers, and STM-64 systems*
- ITU-T Recommendation G.707 (1996), *Network node interface for the synchronous digital hierarchy (SDH)*
- ITU-T Recommendation G.783 (1994), *Characteristics of Synchronous Digital Hierarchy (SDH) equipment functional blocks*
- ITU-T Recommendation G.813 (1996), *Timing characteristics of SDH equipment slave clocks (SEC)*
- ITU-T Recommendation G.825 (1993), *The control of jitter and wander within digital networks which are based on the Synchronous Digital Hierarchy (SDH)*
- ITU-T Recommendation G.826 (1999), *Error performance parameters and objectives for international, constant bit rate digital paths at or above the primary rate*
- ITU-T Recommendation G.831 (1993), *Management capabilities of transport networks based on Synchronous Digital Hierarchy (SDH)*
- ITU-T Recommendation G.957 (1995), *Optical interfaces for equipment and systems relating to the synchronous digital hierarchy*
- ITU-T Recommendation G.958 (1994), *Digital line systems based on the Synchronous Digital Hierarchy for use on optical fibre cables*
- ITU-T Recommendation I.432 (1993), *B-ISDN User-Network Interface Physical layer specification*

Other Supported Standards

The following sections describe other standards supported by JUNOS software:

- ATM on page 41
- Ethernet on page 41
- Frame Relay on page 41
- Serial on page 42
- T3 on page 42

ATM

- ITU-T Recommendation I.363, *B-ISDN ATM adaptation layer sublayers: service-specific coordination function to provide the connection-oriented transport service* (JUNOS software conforms only to the AAL5/IP over ATM portion of this standard)
- ITU-T Recommendation I.432.3, *B-ISDN User-network Interface Physical Layer Specifications: 51,840 kbits/s operation*

Ethernet

- IEEE 802.3ad, *Link Aggregation* (*Aggregation of Multiple Link Segments and Link Aggregation Control Protocol only)*
- IEEE 802.3, *Carrier Sense Multiple Access with Collision Detection (CSMA/CD)*
- IEEE 802.1Q, *Virtual LANs*

Frame Relay

- ANSI T1.617-1991, *Annex D, Additional procedures for permanent virtual connections (PVCs) using unnumbered information frames*
- FRF.12, *Frame Relay Fragmentation Implementation Agreement*
- FRF.15, *End-to-End Multilink Frame Relay Implementation Agreement*
- FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*
- ITU Q.933a, *Annex A, Additional Procedures for Permanent Virtual Connections (PVC) status management (using Unnumbered Information frames)*
- Internet draft draft-martini-frame-encap-mpls-01.txt (except translation of the command/response bit and sequence numbers and padding), *Frame Relay Encapsulation over Pseudo-Wires* (expires June 2002)

Serial

- ITU-T Recommendation V.35, *Data Transmission at 48 kbit/s Using 60-108 kHz Group Band Circuits*

NOTE: The Juniper Networks Serial PIC supports V.35 interfaces with speeds higher than 48 kilobits per second (Kbps).

- ITU-T Recommendation X.21, *Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment for Synchronous Operation on Public Data Networks*
- TIA/EIA Standard 530, *High-Speed 25-Position Interface for Data Terminal Equipment and Data Circuit-Terminating Equipment*
- TIA/EIA Standard 232, *Interface between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange*

T3

- ITU-T Recommendation G.703, *Physical/electrical characteristics of hierarchical digital interfaces*

Chapter 2

Product Architecture

The JUNOS software provides IP routing protocol software as well as software for interface, network, and chassis management. The JUNOS software runs on all Juniper Networks J-series, M-series, and T-series routing platforms.

- J-series Services Routers (J2300, J4300, and J6300) are deployed at the remote edge of distributed networks.

- Most M-series routers are deployed in small and medium cores, in peering, route reflector, data center applications, or at the IP or Multiprotocol Label Switching (MPLS) edge to support high-performance Layer 2 and Layer 3 services. All M-series routers have redundant power and cooling and the M10i, M20, M40e, M160, and M320 routers have fully redundant hardware including Routing Engines and Switching, Forwarding Engine Boards, or Switch Interface Boards.

- T-series routing platforms (T320 router, T640 router, and TX Matrix platform) are deployed at the core of provider networks. These routing platforms have fully redundant hardware, including power and cooling, Routing Engines, and Switch Interface Boards.

A routing matrix is a multichassis architecture composed of one TX Matrix platform, and from one to four T640 routing nodes. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix platform controls all the T640 routing nodes on the routing matrix.

For more information about the architecture in your routing platform, see the hardware guide for your routing platform.

This chapter provides an overview of the router hardware for routing platforms and discusses the relationships between the hardware and software:

Hardware Overview

The JUNOS software runs on three types of Juniper Networks routing platforms: J-series, M-series, and T-series. The routing platforms consist of the following major hardware components:

- Routing Engine
- Control Board
- Chassis
- Switch interface boards
- Physical Interface Card (PIC)
- Flexible PIC Concentrators (FPCs), each populated by PICs for various interface types. On some routing platforms, the PICs are installed directly in the chassis.
- Power supplies
- Cooling system

For information about specific components in your routing platform, see the hardware guide for your routing platform.

Product Architecture

The routing platforms are made up of two components (see Figure 1 on page 45):

- Packet Forwarding Engine—The Packet Forwarding Engine uses application-specific integrated circuits (ASICs) to perform Layer 2 and Layer 3 packet switching, route lookups, and packet forwarding.
- Routing Engine—The Routing Engine controls the routing updates and system management. The Routing Engine consists of routing protocol software processes running inside a protected memory environment on a general purpose computer platform.

Because this architecture dedicates separate control operations such as routing updates and system management from packet forwarding, the router can deliver superior performance and highly reliable Internet operation.

Figure 1: Product Architecture

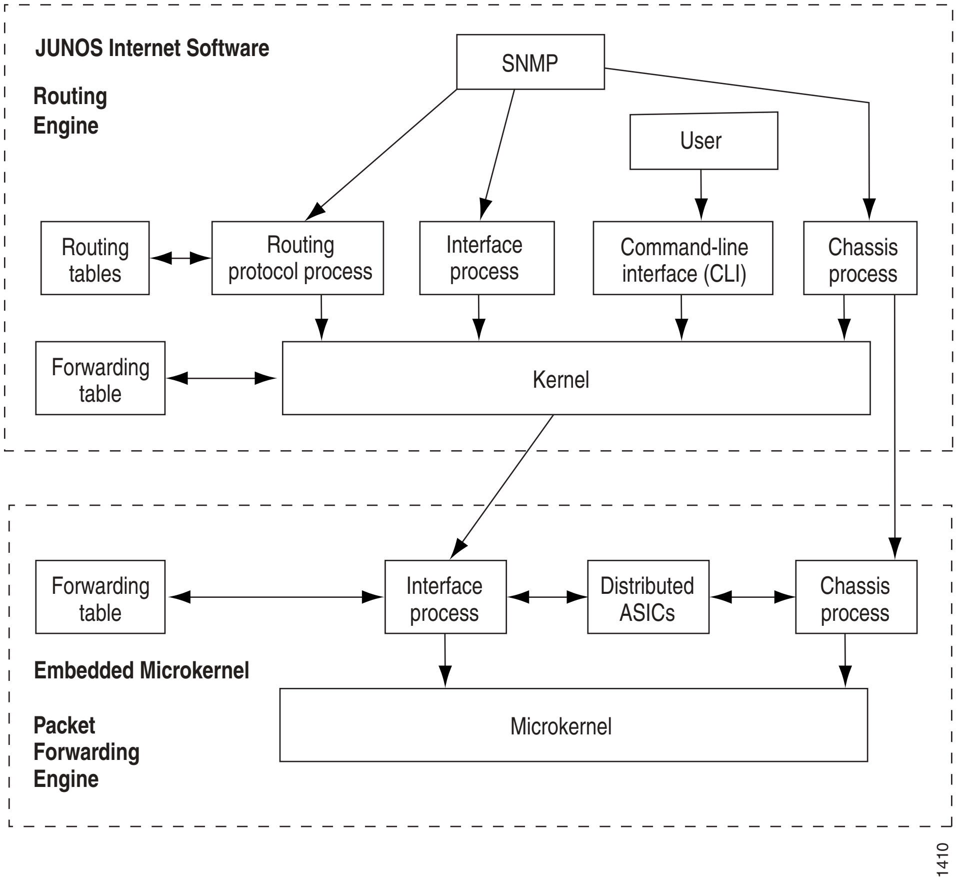

Packet Forwarding Engine

The Packet Forwarding Engine forwards packets between input and output interfaces. The M-series routers (except the M320 router) have a single Packet Forwarding Engine. The J-series Services Routers have a software-based Packet Forwarding Engine. The M320 router, and T-series routing platforms have multiple Packet Forwarding Engines. For more information about the Packet Forwarding Engine, see the hardware guide for your routing platform.

Routing Engine

The Routing Engine handles all the routing protocol processes and other software processes that control the routing platform's interfaces, some of the chassis components, system management, and user access to the routing platform. These routing platform and software processes run on top of a kernel that interacts with the Packet Forwarding Engine. The M320 router and T-series routing platforms have redundant Routing Engines. For more information about routers with redundant Routing Engines, see the hardware guide for your routing platform.

The Routing Engine has these features:

- Routing protocol packets processing—All routing protocol packets from the network are directed to the Routing Engine, and therefore do not delay the Packet Forwarding Engine unnecessarily.

- Software modularity—By dividing software functions into separate processes, a failure of one process has little or no effect on other software processes.

- In-depth IP functionality—Each routing protocol is implemented with a complete set of IP features and provides full flexibility for advertising, filtering, and modifying routes. Routing policies are set according to route parameters, such as prefix, prefix lengths, and Border Gateway Protocol (BGP) attributes.

- Scalability—The JUNOS routing tables are designed to hold all the routes in current and near-future networks. Additionally, the JUNOS software can efficiently support large numbers of interfaces and virtual circuits.

- Management interfaces—System management is possible with a command-line interface (CLI), a craft interface, and Simple Network Management Protocol (SNMP).

- Storage and change management—Configuration files, system images, and microcode can be held and maintained in one primary and two secondary storage systems, permitting local or remote upgrades.

- Monitoring efficiency and flexibility—Alarms can be generated and packets can be counted without adversely affecting packet forwarding performance.

The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the *forwarding table*, which is then copied into the Packet Forwarding Engine. The forwarding table in the Packet Forwarding Engine can be updated without interrupting the routing platform's forwarding.

In a JUNOS-FIPS environment, hardware configurations with two Routing Engines must use IPSec and a private routing instance for all communications between the Routing Engines. IPSec communication between the Routing Engines and AS II FIPS PICs is also required.

Chapter 3

Complete Configuration Mode Commands and Statements for M-series and T-series Platforms

This chapter shows the complete configuration mode commands and the complete configuration statement hierarchy. Using these commands and statements is described in other chapters.

- Complete Configuration Mode Commands on page 48
- Complete Configuration Statement Hierarchy on page 49

For information about command-line interface (CLI) operational mode commands, see the *JUNOS System Basics and Services Command Reference*. For information about IP version 6 (IPv6) configuration statements, see the *JUNOS Routing Protocols Configuration Guide*. For information about configuration mode commands and statements for J-series Services Routers, see “Complete Configuration Mode Commands and Statements for J-series Services Routers” on page 125.

Complete Configuration Mode Commands

The following is the complete list of configuration mode commands, listing all possible commands in the hierarchy.

```
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 an ASCII file
  quit         Quit from this level
  rename       Rename a statement
  rollback     Roll back database to last committed version
  run          Run an operational-mode command
  save         Save configuration to an ASCII file
  set          Set a parameter
  show         Show a parameter
  status       Display database user status
  top          Exit to top level of configuration
  up           Exit one level of configuration
```

Complete Configuration Statement Hierarchy

This section shows the complete configuration statement hierarchy, listing all possible configuration statements and showing their level in the configuration hierarchy. When you are configuring the JUNOS software, your current hierarchy level is shown in the banner on the line preceding the user@host# prompt.

This section is organized as follows:

[edit access] Hierarchy Level

```
access {
    address-pool name {
      address address-or-prefix value;
      address-range low <lower-limit> high <upper-limit >;
    }
    group-profile group-profile-name {
      l2tp {
        interface-id interface-identifier;
        lcp-renegotiation;
        local-chap;
        maximum-sessions-per-tunnel number;
        multilink {
          drop-timeout time;
          fragmentation-threshold bytes;
        }
      ppp {
        framed-pool pool-identifier;
        idle-timeout seconds;
        interface-id interface-identifier;
        keepalive seconds;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
      }
    }
    profile profile-name {
      authentication-order [ authentication-methods ];
      client client-name {
        chap-secret chap-secret;
        group-profile name;
        ike {
          allowed-proxy pair {
            remote remote-proxy-address local local-proxy-address;
          }
          pre-shared-key [ascii-text key-string] [hexadecimal key-string];
          interface-id <string-value>;
        }
        l2tp {
          interface-id interface-identifier;
          lcp-renegotiation;
          local-chap;
          maximum-sessions-per-tunnel number;
          multilink {
            drop-timeout time;
            fragmentation-threshold bytes;
          }
          ppp-authentication (chap | pap);
          ppp-profile profile-name;
          shared-secret shared-secret;
        }
        pap-password pap-password;
        ppp {
          framed-ip-address ip-address;
          framed-pool framed-pool;
          idle-timeout seconds;
          interface-id interface-identifier;
          keepalive seconds;
          primary-dns primary-dns;
          primary-wins primary-wins;
          secondary-dns secondary-dns;
          secondary-wins secondary-wins;
        }
      }
      radius-server server-address {
        accounting-port number;
        port number ;
        retry number ;
        routing-instance routing-instance-name;
        secret password ;
        timeout seconds ;
      }
    }
    radius-disconnect-port port-number {
    radius-disconnect {
      client-address {
        secret password;
      }
    }
    radius-server server-address {
      accounting-port number;
      port number ;
      retry number ;
      routing-instance routing-instance-name;
      secret password ;
      timeout seconds ;
    }
    traceoptions {
      file filename {
        files number;
        size maximum-file-size;
      }
      flag all;
      flag authentication;
      flag chap;
      flag configuration;
      flag kernel;
      flag radius;
      flag world-readable;
    }
} # End of [edit access] hierarchy level
```

[edit accounting-options] Hierarchy Level

```
accounting-options {
    class-usage-profile profile-name {
        file filename;
        interval minutes;
        source-classes {
            source-class-name;
        }
    destination-class-profile profile-name {
        destination-class {
            destination-class-name;
        }
        file filename;
        interval minutes;
    }
    file filename {
        archive-sites {
            site-name;
        }
        file file-number;
        size bytes;
        transfer-interval minutes;
    }
    filter-profile profile-name {
        counters {
            counter name;
        }
        file filename;
        interval minutes;
    }
    interface-profile profile-name {
        fields {
            field-name;
        }
        file filename;
        interval minutes;
    }
    routing-engine-profile profile-name {
        fields {
            field-name;
        }
        file filename;
        interval minutes;
    }
} # End of [edit accounting-options] hierarchy level
```

[edit applications] Hierarchy Level

```
applications {
    application application-name {
        application-protocol protocol-name;
        destination-port port-number;
        icmp-code value;
        icmp-type value;
        inactivity-timeout value;
        learn-sip-register;
        protocol type;
        rpc-program-number number;
        sip-call-hold-timeout minutes;
        snmp-command command;
        source-port port-number;
        ttl-threshold value;
        uuid hex-value;
    }
    application-set application-set-name {
        [ application application-name ];
    }
} # End of [edit applications] hierarchy level
```

[edit chassis] Hierarchy Level

```
chassis {
    aggregated-devices {
        ethernet {
            device-count number;
        }
        sonet {
            device-count number;
        }
    }
    alarm {
        interface-type {
            alarm-name (red | yellow | ignore);
        }
    }
    config-button {
        no-clear;
        no-rescue;
    }
    fpc slot-number {
        pic pic-number {
            atm-cell-relay-accumulation;
            atm-l2circuit-mode (cell | aal5 | trunk trunk);
            ce1 {
                e1 port-number {
                    channel-group group-number timeslots slot-number;
                }
            }
            ct3 {
                port port-number {
                    t1 link-number {
                        channel-group group-number timeslots slot-number;
                    }
                }
            }
            framing (sdh | sonet);
            idle-cell-format {
                itu-t;
                payload-pattern payload-pattern-byte;
            }
            max-queues-per-interface (8 | 4);
            mlfr-uni-nni-bundles number;
            no-concatenate;
            q-pic-large-buffer;
            t1;
            vtmapping (itu-t | klm);
        }
    }
    lcc number {
        fpc slot-number {
            pic pic-number {
                atm-cell-relay-accumulation;
                atm-l2-circuit-mode (cell | aal5 | trunk trunk);
                framing (sdh | sonet);
                idle-cell-format {
                    itu-t;
                    payload-pattern payload-pattern-byte;
                }
                max-queues-per-interface (8 | 4);
                no-concatenate;
            }
        }
        offline;
        online-expected;
    }
    (packet-scheduling | no-packet-scheduling);
    pem {
        minimum number;
    }
    redundancy {
        failover {
            on-disk-failure;
            on-loss-of-keepalives;
        }
        graceful-switchover (disable | enable);
        keepalive-time seconds;
        routing-engine (Redundancy) slot-number (master | backup | disabled);
        sfm slot-number (always | preferred);
        ssb slot-number (always | preferred);
    }
    routing-engine {
        on-disk-failure reboot;
    }
    sfm slot-number {
        power off;
    }
    sib {
        minimum number;
    }
    (source-route | no-source-route);
    vrf-mtu-check;
    vtmapping (km | itu-t);
} # End of [edit chassis] hierarchy level
```

[edit class-of-service] Hierarchy Level

```
class-of-service {
    adaptive-shapers {
        adaptive-shaper-name {
            trigger type shaping-rate (percent percent | rate);
        }
    }
    classifiers {
        type classifier-name {
            forwarding-class class-name {
                loss-priority (low | high) code-points [ alias | bits ];
            }
            import (classifier-name | default);
        }
    }
    code-point-aliases {
        (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) {
            alias-name bits;
        }
    }
    drop-profiles {
        profile-name {
            fill-level percentage drop-probability percentage;
            interpolate {
                drop-probability value;
                fill-level value;
            }
        }
    }
    fabric {
        scheduler-map {
            priority (low | high) scheduler scheduler-name;
        }
    }
    forwarding-classes {
        queue queue-number class-name priority (low | high);
    }
    forwarding-policy {
        class class-name {
            classification-override {
                forwarding-class class-name;
            }
        }
        next-hop-map map-name {
            forwarding-class class-name {
                next-hop [ next-hop-name ];
                lsp-next-hop [ lsp-regular-expression ];
            }
        }
    }
```

```
fragmentation-maps {
map-name {
    forwarding-class class-name {
        fragment-threshold bytes;
        multilink-class number;
        no-fragmentation;
    }
}
interfaces {
    interface-name {
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        unit logical-unit-number {
            adaptive-shaper adaptive-shaper-name;
            classifiers {
                (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence)
                        (classifier-name | default);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            loss-priority-maps {
                (map-name | default);
            }
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                frame-relay-de (rewrite-name | default);
                ieee-802.1 default;
                inet-precedence (rewrite-name | default);
            }
            scheduler-map map-name;
            shaping-rate rate;
            virtual-channel-group virtual-channel-group-name;
        }
    }
}
loss-priority-maps {
    frame-relay-de (map-name | default) {
        loss-priority level code-points [ values ];
    }
}
restricted-queues {
    forwarding-class class-name queue queue-number;
}
rewrite-rules {
    (dscp | exp | inet-precedence) rewrite-name {
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
    }
}
```

```
        routing-instances routing-instance-name {
            classifier {
                exp (classifier-name | default);
            }
        }
        scheduler-maps {
            map-name {
                forwarding-class class-name scheduler scheduler-name;
            }
        }
        schedulers
            scheduler-name {
                buffer-size (percent percentage | remainder | temporal microseconds);
                drop-profile-map loss-priority (low | high) protocol (non-tcp | tcp | any)
                    drop-profile profile-name;
                priority (low | high | strict-high);
                transmit-rate (rate | percent percentage | remainder | exact);
            }
        }
        tri-color;
        virtual-channels {
            virtual-channel-name;
        }
        virtual-channel-groups {
            virtual-channel-group-name {
                virtual-channel-name {
                    scheduler-map map-name;
                    shaping-rate (percent percent | rate);
                    default;
                }
            }
        }
} # End of [edit class-of-service] hierarchy level
```

[edit firewall] Hierarchy Level

```
firewall {
    family family-name {
      filter filter-name term term-name then {
        accounting-profile name;
        interface-specific;
        virtual-channel virtual-channel-name;
      }
      prefix-action name {
        count;
        destination-prefix-length prefix-length;
        policer policer-name;
        source-prefix-length prefix-length;
        subnet-prefix-length prefix-length;
      }
      prefix-policer {
        policer policer-name;
      }
    }
    interface-set interface-set-name {
      [ interface-names ];
    }
    load-balance-group group-name {
      next-hop-group [ group-names ];
    }
    policer policer-name {
      filter-specific;
      if-exceeding {
        bandwidth-limit bps;
        bandwidth-percent number;
        burst-size-limit bytes;
      }
      logical-interface-policer;
      then {
        policer-action;
      }
    }
    three-color-policer name {
      two-rate {
        (color aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        excess-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
      }
    }
} # End of [edit firewall] hierarchy level
```

[edit forwarding-options] Hierarchy Level

```
forwarding-options {
    accounting group-name {
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      interface [ interface-names ] {
        engine-id number;
        engine-type number;
        source-address address;
      }
      output {
        autonomous-system-type (origin | peer);
        cflowd [ host-names ] {
          aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
              caida-compliant;
            }
            source-prefix;
          }
        }
        port port-number;
        version format;
      }
    }
    family family-name {
      filter {
        input filter-name;
      }
      flood {
        input filter-name;
      }
    }
    hash-key {
      family inet {
        layer-3;
        layer-4;
      }
      family mpls {
        label-1;
        label-2;
        ip;
        payload;
      }
    }
```

```
helpers {
  bootp {
    description description-of-service;
    interface interface-group {
      description description-of-interface;
      maximum-hop-count number;
      minimum-wait-time seconds;
      no-listen;
      server [ addresses ];
    }
    server address < [ routing-instance routing-instance-name ] >;
  }
  domain {
    description description-of-service;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server address < [ routing-instance routing-instance-name ] >;
    }
    server address;
  }
  tftp {
    description description-of-service;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server address;
    }
    server address < [ routing-instance routing-instance-name ] >;
  }
  traceoptions {
    file filename;
      files number;
      size bytes;
    }
      flag flag;
    level level;
  }
}
monitoring group-name {
  family inet {
    export-format cflowd-version-5;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-names ] {
      engine-id number;
      engine-type number;
      input-interface-index number;
      output-interface-index number;
      source-address address;
    }
    output {
      cflowd [ host-names ] {
        port port-number;
    }
  }
```

```
    next-hop-group [ group-names ] {
        interface interface-name {
            next-hop [ addresses ];
        }
    }
}
pcap {
    enable;
}
port-mirroring {
    family (inet | inet6) {
        input {
            family inet {
                rate number;
                run-length number;
            }
        }
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
    }
    input {
        family inet {
            rate num;
            run-length num;
        }
    }
    output {
        interface interface-name {
            next-hop address;
        }
        no-filter-check;
    }
}
sampling {
    disable;
    input {
        family inet {
        max-packets-per-second number;
            rate number;
            run-length number;
        }
    }
    output {
        cflowd [ host-names ] {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
```

```
          autonomous-system-type (origin | peer);
          (local-dump | no-local-dump);
          port port-number;
          source-address address;
          version format;
        }
        file {
          disable;
          filename filename;
          files number;
          size bytes;
          (stamp | no-stamp);
          (world-readable | no-world-readable);
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface [ interface-names ] {
          engine-id number;
          engine-type number;
          source-address address;
        }
      }
      traceoptions {
        file filename {
          files number;
          size bytes;
          (world-readable | no-world-readable);
        }
      }
    }
} # End of [edit forwarding-options] hierarchy level
```

[edit groups] Hierarchy Level

```
groups {
    group-name {
      configuration-data;
    }
} # End of [edit groups] hierarchy level
```

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-routers *logical-router-name*] hierarchy level.

```
interfaces {
    interface-name {
        disable;
        accounting-profile name;
        description text;
        aggregated-ether-options {
            (flow-control | no-flow-control);
            lacp mode {
                periodic interval;
            }
            link-speed speed;
            (loopback | no-loopback);
            minimum-links number;
            source-address-filter {
                mac-address;
            }
            (source-filtering | no-source-filtering);
        }
        aggregated-sonet options {
            link-speed speed;
            minimum-links number;
        }
        atm-options {
            cell-bundle-size cells;
            ilmi;
            linear-red-profiles profile-name {
                high-plp-max-threshold percent;
                low-plp-max-threshold percent;
                queue-depth cells high-plp-threshold percent low-plp-threshold percent;
            }
            mpls {
                pop-all-labels {
                    required-depth number;
                }
            }
            pic-type (atm1 | atm2);
            plp-to-clp;
            promiscuous-mode {
                vpi vpi-identifier;
            }
            scheduler-maps map-name {
                forwarding-class class-name {
                    epd-threshold cells plp1 cells;
                    linear-red-profile profile-name;
                    priority (low | high);
                    transmit-weight (cells number | percent number);
                }
                vc-cos-mode (alternate | strict);
            }
    vpi vpi-identifier {
        maximum-vcs maximum-vcs;
        oam-liveness {
            down-count cells;
            up-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length |
                    vbr peak rate sustained rate burst length);
            queue-length number;
        }
    }
}
clocking clock-source;
dce;
description text;
dialer-options {
    pool pool-name {
        priority priority;
    }
}
disable;
ds0-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx64 | nx56);
    fcs (32 | 16);
    idle-cycle-flag (flags | ones);
    invert data;
    loopback (payload | remote);
    start-end-flag (shared | filler);
}
e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (32 | 16);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert data;
    loopback (local | remote);
    start-end-flag (shared | filler);
    timeslots time-slot-range;
}
e3-options {
    atm-encapsulation (direct | PLCP);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
    fcs (32 | 16);
    framing (g.751 | g.832);
    idle-cycle-flag value;
```

```
   loopback (local | remote);
   (payload-scrambler | no-payload-scrambler);
   start-end-flag value;
   (unframed | no-unframed);
}
encapsulation type;
ether-vpls-over-atm-llc;
es-options {
   backup-interface es-fpc/pic/port;
}
failure-options {
   [trigger-link-failure interface-name ];
}
fastether-options {
   802.3ad aex;
   (flow-control | no-flow-control);
   ingress-rate-limit rate;
   (loopback | no-loopback);
   source-address-filter {
      mac-address;
   }
   (source-filtering | no-source-filtering);
}
gigether-options {
   802.3ad aex;
   (flow-control | no-flow-control);
   (loopback | no-loopback);
   source-address-filter {
      mac-address;
   }
   (source-filtering | no-source-filtering);
   ethernet-switch-profile {
      (mac-learn-enable | no-mac-learn-enable);
      tag-protocol-id [ tpids ];
      ethernet-policer-profile {
         input-priority-map {
            ieee802.1p premium [ values ];
         }
         output-priority-map {
            classifier {
               premium {
                  forward-class class-name {
                     loss-priority (high | low);
                  }
               }
            }
         }
         policer cos-policer-name {
            aggregate {
               bandwidth-limit bps;
               burst-size-limit bytes;
            }
```

```
                premium {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
            }
        }
    }
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
isdn-options {
    bandwidth (64 | 56);
    interface-type (user | network);
    switch-type (NI2 | ETSI | ATT5E);
    spid1 spid-string;
    spid2 spid-string;
    tei-option (first-call | power-up);
    t306 seconds;
    t310 seconds;
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
    lmi-type (ansi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
mac mac-address;
mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    hello-timer milliseconds;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
}
mtu bytes;
multiservice-options {
    boot-command filename;
    (core-dump | no-core-dump);
    (syslog | no-syslog);
}
```

```
no-gratuitous-arp-request;
no-keepalives;
no-partition {
   interface-type type;
}
partition partition-number oc-slice oc-slice-range interface-type type {
   timeslots time-slot-range;
}
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
   chap {
      access-profile name;
      local-name name;
      passive;
   }
   compression {
      acfc;
      pfc;
   }
}
receive-bucket {
   overflow (discard | tag);
   rate percentage;
   threshold bytes;
}
redundancy-options {
   primary sp-fpc/pic/port;
   secondary sp-fpc/pic/port;
}
serial-options {
   clock-rate rate;
   clocking-mode (dce | dte | loop);
   control-leads {
      control-signal (assert | de-assert | normal);
      cts (ignore | normal | require);
      dcd (ignore | normal | require);
      dsr (ignore | normal | require);
      dtr signal-handling-option;
      ignore-all;
      indication (ignore | normal | require);
      rts (assert | de-assert | normal);
      tm (ignore | normal | require);
   }
   control-polarity (positive | negative);
   cts-polarity (positive | negative);
   dcd-polarity (positive | negative);
   dsr-polarity (positive | negative);
   dtr-circuit (balanced | unbalanced);
   dtr-polarity (positive | negative);
   encoding (nrz | nrzi);
   indication-polarity (positive | negative);
   line-protocol protocol;
   loopback mode;
```

```
    rts-polarity (positive | negative);
    tm-polarity (positive | negative);
    transmit-clock invert;
}
service-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    syslog {
        host host-name {
            facility-override facility-name;
            log-prefix prefix-number;
            [ services priority-level ];
        }
    }
}
sonet-options {
    aggregate asx;
    aps {
        advertise-interval milliseconds;
        authentication-key key;
        force;
        hold-time milliseconds;
        lockout;
        neighbor address;
        paired-group group-name;
        protect-circuit group-name;
        request;
        revert-time seconds;
        switching-mode (bidirectional | unidirectional);
        working-circuit group-name;
    }
    bytes {
        c2 value;
        e1-quiet value;
        f1 value;
        f2 value;
        s1 value;
        z3 value;
        z4 value;
    }
    fcs (32 | 16);
    loopback (local | remote);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    path-trace trace-string;
    (payload-scrambler | no-payload-scrambler);
    rfc-2615;
    trigger {
        defect ignore;
            hold-time up milliseconds down milliseconds;
        }
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
}
speed (10m | 100m);
stacked-vlan-tagging;
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx64 | nx56);
    fcs (32 | 16);
    framing (sf | esf);
    idle-cycle-flags (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (shared | filler);
    timeslots slot-number ;
}
t3-options {
    atm-encapsulation (direct | PLCP);
    bert-algorithm algorithm ;
    bert-error-rate rate ;
    bert-period seconds ;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
        value>;
    fcs (32 | 16);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value ;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local |payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag value ;
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}
transmit-bucket {
    overflow (tag | discard);
    rate percentage ;
    threshold bytes ;
}
(traps | no-traps);
vlan-tagging;
unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address;
    policer {
        input policer-name;
        output policer-name;
        }
    }
}
accounting-profile name;
allow_any_vci;
atm-scheduler-map (map-name | default);
backup-options {
    interface interface-name;
}
bandwidth rate;
cell-bundle-size cells;
clear-dont-fragment-bit;
compression {
    rtp {
        f-max-period number;
        queues [queue-numbers];
        port {
            minimum port-number;
            maximum port-number;
        }
    }
}
compression-device interface-name;
description text;
dial-options {
    activation-delay seconds;
    bearer-cap (udi | rdi | speech);
    deactivation-delay seconds;
    (dedicated | shared);
    pool pool;
    dial-string [dial-string-numbers];
    idle-timeout seconds;
    initial-route-check seconds;
    ipsec-interface-id name;
    l2tp-interface-id name;
    load-threshold number;
    re-enable seconds;
    watch-list {
        routes;
    }
}
disable;
dlci dlci-identifier;
drop-timeout milliseconds;
encapsulation type;
epd-threshold cells plp1 cells;
family family {
    address address {
        destination address;
    }
}
```

```
  fragment-threshold bytes;
  input-vlan-map {
    pop;
    push;
    swap;
    vlan-id number;
    tag-protocol-id tpid;
  }
  interleave-fragments;
  inverse-arp;
  minimum-links number;
  mrru bytes;
  multicast-dlci dlci-identifier ;
  multicast-vci vpi-identifier.vci-identifier ;
  multipoint;
  oam-liveness {
    up-count cells;
    down-count cells;
  }
  oam-period (disable | seconds);
  output-vlan-map {
    pop;
    push;
    swap;
    vlan-id number;
    tag-protocol-id tpid;
  }
  passive-monitor-mode;
  peer-unit unit-number;
  plp-to-clp;
  point-to-point;
ppp-options {
  chap {
    access-profile name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  proxy-arp;
  service-domain (inside | outside);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length |
        vbr peak rate sustained rate burst length);
    queue-length number;
  }
  short-sequence;
  transmit-weight number;
  (traps | no-traps);
  trunk-bandwidth rate;
  trunk-id number;
```

```
tunnel {
    backup-destination address;
    destination destination-address;
    key number;
    routing-instance {
            destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier ;
vpi vpi-identifier;
vlan-id number;
vlan-tag inner tpid.vlan-id outer tpid.vlan-id;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            [ input output ];
        }
    }
    bundle interface-name;
    filter {
        input filter-name;
        output filter-name;
        group filter-group-number ;
    }
    ipsec-sa sa-name;
    keep-address-and-control;
    mtu bytes;
    multicast-only;
    negotiate-address;
    no-redirects;
    policer {
        arp policer-template-name;
        output policer-template-name;
        group policer-template-name;
    }
    primary;
    proxy inet-address address;
    receive-options-packets;
    receive-ttl-exceeded;
    remote (inet-address address | mac-address address);
    rpf-check <fail-filter filter-name> {
        <mode loose>;
    }
    sampling {
        (input | output | input output);
    }
    service {
        input {
            [ service-set service-set-name <service-filter filter-name> ];
            post-service-filter filter-name; {
        }
        output {
          [ service-set service-set-name <service-filter filter-name> ];
        }
      }
      (translate-discard-eligible | no-translate-discard-eligible);
      (translate-fecn-and-becn | no-translate-fecn-and-becn);
      unnumbered-address interface-name destination address
        destination-profile profile-name;
      address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination destination-address;
        destination-profile name;
        eui-64;
        multipoint-destination destination-address (dlci dlci-identifier |
          vci vci-identifier);
        multipoint-destination destination-address {
          epd-threshold cells plp1 cells;
          inverse-arp;
          oam-liveness {
            up-count cells;
            down-count cells;
          }
          oam-period seconds;
          shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length |
                vbr peak rate sustained rate burst length);
            queue-length number;
          }
          vci vpi-identifier.vci-identifier;
        }
        preferred;
        primary;
        vrrp-group group-number {
          (accept-data | no-accept-data):
          advertise-interval seconds;
          authentication-type authentication;
          authentication-key key;
          fast-interval milliseconds;
          (preempt | no-preempt) {
            hold-time seconds;
          }
          priority number;
          }
          track {
            interface interface-name priority-cost cost;
          }
          virtual-address [ addresses ];
        }
      }
    }
  }
}
```

```
    traceoptions {
        file filename <files number> <size size> <(world-readable | no-world-readable)>;
        flag flag <disable>;
    }
}# End of [edit interfaces] hierarchy level
```

[edit logical-routers] Hierarchy Level

```
logical-routers {
    logical-router-name {
        interfaces {
            interfaces-configuration;
        }
        policy-options {
            policy-options-configuration;
        }
        protocols {
            protocols-configurations;
        }
    routing-instances {
        routing-instances-configuration;
    }
    routing-options {
        routing-options-configuration;
    }
} # End of [edit logical-routers] hierarchy level
```

[edit policy-options] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-routers *logical-router-name*] hierarchy level.

```
policy-options {
    as-path name regular-expression;
    as-path-group group-name;
    community name {
        invert-match;
        members [ community-ids ];
    }
    damping name {
        disable;
        half-life minutes;
        max-suppress minutes;
        reuse number;
        suppress number;
    }
    policy-statement policy-name {
        term term-name {
            default-action (accept | reject);
            from {
                family family-name;
                match-conditions;
                policy subroutine-policy-name;
                prefix-list name;
                route-filter destination-prefix match-type <actions>;
                source-address-filter destination-prefix match-type <actions>;
            }
            to {
                match-conditions;
                policy subroutine-policy-name;
            }
            then actions;
        }
    }
    prefix-list name {
        ip-addresses;
    }
} # End of [edit policy-options] hierarchy level
```

[edit protocols] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-routers *logical-router-name*] hierarchy level.

```
protocols {
```

Bidirectional Forwarding Detection (BFD)

```
    bfd {
        traceoptions {
            file {
                filename;
                files number;
                no-world-readable;
                size bytes;
                world-readable;
            }
        }
        flag {
            adjacency;
            all;
            error;
            event;
            packet;
            pipe;
            state;
            timer;
        }
    } # End of [edit protocols bfd] hierarchy level
```

Border Gateway Protocol (BGP)

```
bgp {
    advertise-inactive;
    advertise-peer-as;
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    disable;
    export [ policy-names ];
    family {
        (iso-vpn | inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
            (any | multicast | unicast) {
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
            }
            labeled-unicast {
                aggregate-label {
                    community community-name;
                }
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                resolve-vpn;
                rib inet.3;
                rib-group group-name;
            }
        }
        route-target {
            advertise-default;
            external-paths number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
        }
        signaling {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
        }
    }
    graceful-restart {
        disable;
        restart-time seconds;
        stale-routes-time seconds;
    }
    hold-time seconds;
    import [ policy-names ];
    include-mp-next-hop;
```

```
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
multihop {
    <ttl-value>;
    no-nexthop-change;
}
no-advertise-peer-as;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
path-selection (cisco-non-deterministic | always-compare-med |
    external-router-id);
peer-as autonomous-system;
preference preference;
remove-private;
traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
group group-name {
    advertise-inactive;
    advertise-peer-as;
    [network/mask-length];
    as-override;
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (iso-vpn | inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
            (any | multicast | unicast) {
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
            }
            flow {
                no-validate policy-name;
            }
            labeled-unicast {
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
            resolve-vpn;
            rib inet.3;
            rib-group group-name;
        }
    }
    route-target {
        advertise-default;
        external-paths number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
    }
    signaling {
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
    }
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath;
no-advertise-peer-as;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
protocol protocol;
remove-private;
traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
}
type type;
vpn-apply-export;
neighbor address {
    advertise-inactive;
    advertise-peer-as;
    as-override;
```

```
authentication-key key;
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family {
    (iso-vpn | inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
        (any | multicast | unicast) {
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group group-name;
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
            rib-group group-name;
        }
    }
    route-target {
        advertise-default;
        external-paths number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
    }
    signaling {
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
    }
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
```

```
            local-preference local-preference;
            log-updown;
            metric-out (metric | minimum-igp <offset> | igp <offset>);
            mtu-discovery;
            multihop <ttl-value>;
            multipath;
            no-advertise-peer-as;
            no-aggregator-id;
            no-client-reflect;
            out-delay seconds;
            passive;
            peer-as autonomous-system;
            preference preference;
            remove-private;
            traceoptions {
                file name <replace> <size size> <files number> <no-stamp>
                    <(world-readable | no-world-readable)>;
                flag flag <flag-modifier> <disable>;
            }
            vpn-apply-export;
        }
    }
} # End of [edit protocols bgp] hierarchy level
```

Connections

```
connections {
    interface-switch connection-name {
        interface interface-name.unit-number;
        interface interface-name.unit-number;
    }
    lsp-switch connection-name {
        transmit-lsp label-switched-path;
        receive-lsp label-switched-path;
    }
    p2mp-transmit-switch point-to-multipoint-transmit-switch-name {
        input-interface input-interface-name.unit-number;
        transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
    }
    remote-interface-switch connection-name {
        interface interface-name.unit-number;
        transmit-lsp label-switched-path;
        receive-lsp label-switched-path;
    }
} # End of [edit protocols connections] hierarchy level
```

Data-Link Switching (DLSw)

```
dlsw {
    connection-idle-timeout time;
    local-peer peer-address;
    multicast-address address;
    promiscuous;
    receive-initial-pacing count;
    remote-peer peer-address;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
} #End of [edit protocols dlsw] hierarchy level
```

Distance Vector Multicast Routing Protocol (DVMRP)

```
dvmrp {
   disable;
   export [ policy-names ];
   import [ policy-names ];
   interface interface-name {
      disable;
      hello-interval seconds;
      hold-time seconds;
      metric metric;
      mode (forwarding | unicast-routing);
   }
   rib-group group-name;
      inet;
   }
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
   }
} # End of [edit protocols dvmrp] hierarchy level
```

Internet Group Management Protocol (IGMP)

```
igmp {
   interface interface-name {
      disable;
      ssm-map ssm-map-name:
      static {
         group group {
            source source;
         }
      }
      version version;
   }
   query-interval seconds;
   query-last-member-interval seconds;
   query-response-interval seconds;
   robust-count number;
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
      }
   }
} # End of [edit protocols igmp] hierarchy level
```

End System-to-Intermediate System (ES-IS)

```
esis {
   interface [(interface-name | all)] {
      hello-interval seconds;
      hold-time seconds;
}
} # End of [edit protocols esis] hierarchy level
```

Intermediate System-to-Intermediate System (IS-IS)

```
isis {
    clns-routing;
    disable;
    export [ policy-names ];
    ignore-attached-bit;
    graceful-restart {
        disable;
        helper-disable;
        restart-duration seconds;
    }
    label-switched-path name level level metric metric;
    level level-number {
        authentication-key key;
        authentication-type authentication;
        external-preference preference;
        ipv6-multicast-metric number;
        no-csnp-authentication;
        no-hello-authentication;
        no-psnp-authentication;
        preference preference;
        prefix-export-limit num;
        wide-metrics-only;
    }
    lsp-lifetime seconds;
    no-authentication-check;
    no-ipv4-routing;
    no-ipv6-routing;
    overload {
        advertise-high-metrics;
        <timeout seconds>;
    {
    reference-bandwidth reference-bandwidth;
    rib-group group name;
    spf-delay milliseconds;
    topologies {
        ipv4-multicast;
        ipv6-multicast;
        ipv6-unicast;
    }
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>;
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
    traffic-engineering {
        disable;
        ipv4-multicast-rpf-routes;
        shortcuts;
    }
    interface interface-name {
        disable;
        bfd-liveness-detection {
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier number;
            version (0 | 1);
        }
        checksum;
        csnp-interval (seconds | disable);
        lsp-interval milliseconds;
        mesh-group (value | blocked);
        no-clns-unicast;
        no–ipv4-multicast;
        no–ipv6-multicast;
        no-ipv6-unicast;
        no-unicast-topology;
        passive;
        point-to-point;
        level level-number {
            clns-unicast-metric metric;
            disable;
            hello-authentication-key key;
            hello-authentication-type authentication;
            hello-interval seconds;
            hold-time seconds;
            ipv4-multicast-metric number;
            ipv6-unicast-metric number;
            metric metric;
            passive;
            priority number;
            te-metric metric;
        }
    }
} # End of [edit protocols isis] hierarchy level
```

Layer 2 Circuits

```
l2circuit {
    neighbor address {
        interface interface-name {
            mtu mtu-number;
            protect-interface interface-name;
            virtual-circuit-id identifier;
        }
    }
    traceoptions {
        file file-name [replace] [size number] [files file-names] [nostamp];
        flag (connections | error | FEC | topology) [detail];
    }
} # End of [edit protocols l2circuit] hierarchy level
```

Label Distribution Protocol (LDP)

```
ldp {
    deaggregate | no-deaggregate;
    egress-policy policy-name;
    export policy-name;
    graceful-restart {
        disable;
        helper-disable;
        maximum-recovery-time value;
        recovery-time value;
    }
    import policy-name;
    keepalive-interval seconds;
    keepalive-timeout seconds;
    preference preference;
    transport-address (interface | loopback);
    interface interface-name {
        disable;
        hello-interval seconds;
        hold-time seconds;
        deaggregate | no-deaggregate;
        transport-address (interface | loopback);
    }
    keepalive-interval seconds;
    keepalive-timeout seconds;
    no-forwarding;
    preference preference;
    session address {
        authentication-key authentication-key;
    }
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
    track-igp-metric;
    traffic-statistics {
        file filename <replace> <size size> <files number> <(world-readable |
                no-world-readable)>;
        interval interval;
        transport-address (interface | router-id);
    }
} # End of [edit protocols ldp] hierarchy level
```

Link Management

```
link-management {
    peer peer-name {
        address address;
        control-channel [ control-channel-interfaces ];
        te-link [ te-link-names ];
    }
    te-link te-link-name {
        disable;
        interface interface-name {
            disable;
            local-address address;
            remote-address address;
            remote-id id-number;
        }
        local-address address;
        remote-address address;
        remote-id id-number;
    }
    traceoptions {
        file filename <files number> <no-stamp> <replace> <size size>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
}# End of [edit protocols link-management] hierarchy level
```

Multicast Listener Discovery (MLD)

```
mld {
    interface interface-name {
        disable;
        ssm-map ssm-map-name:
        static {
            group group {
                source source;
            }
        }
        version version;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
} # End of [edit protocols mld] hierarchy level
```

Multiprotocol Label Switching (MPLS)

```
mpls {
    disable;
    admin-groups {
        group-name group-value;
    }
    advertise-hold-time seconds;
    auto-policing {
        class all policer-action;
        class ctnumber (drop | loss-priority-high | loss-priority-low);
    }
    bandwidth bandwidth;
    class-of-service cos-value;
    diffserv-te {
        bandwidth-model {
            extended-mam;
            mam;
        }
        te-class-matrix {
            tenumber {
                priority priority;
                traffic-class {
                    ctnumber priority priority;
                }
            }
        }
    }
    explicit-null;
    hop-limit number;
    icmp-tunneling;
    interface (interface-name | all) {
        disable;
        admin-group {
            group-name;
        }
        label-map {
            in-label {
                class-of-service value;
                (next-hop (address | interface-name | address/interface-name)) | (reject |
                    discard);
                (pop | (swap <out-label>);
                preference preference;
                type type;
            }
            default-route {
                class-of-service value;
                (next-hop (address | interface-name | address/interface-name)) | (reject |
                    discard);
                (pop | (swap <out-label>);
                preference preference;
                type type;
            }
        }
    }
```

```
ipv6-tunneling;
label-switched-path lsp-path-name {
    disable;
    adaptive;
    admin-group {
        exclude group-names;
        include group-names;
    }
    auto-bandwidth {
        adjust-interval seconds;
        adjust-threshold percent;
        maximum-bandwidth bps;
        minimum-bandwidth bps;
        monitor-bandwidth;
    }
    bandwidth bps;
    class-of-service cos-value;
    description;
    fast-reroute {
        bandwidth bps;
        bandwidth-percent percent;
        (exclude group-names | no-exclude);
        hop-limit number;
        (include group-names | no-include);
    }
    from address;
    hop-limit number;
    install {
        destination-prefix/prefix-length <active>;
    }
    ldp-tunneling;
    link-protection;
    lsp-attributes {
        encoding-type (ethernet | packet | pdh | sonet-sdh);
        gpid (ethernet | hdlc | ipv4 | ppp);
        signal-bandwidth type;
        switching-type type;
    }
    metric number;
    no-cspf;
    no-decrement-ttl;
    no-install-to-address;
    node-link-protection;
    optimize-timer seconds;
    policing {
        filter filter-name;
        no-automatic-policing;
    }
    preference preference;
    priority setup-priority hold-priority;
```

```
primary path-name {
   adaptive;
   admin-group {
      exclude group-names;
      include group-names;
   }
   bandwidth bps;
   class-of-service cos-value;
   hop-limit number;
   no-cspf;
   no-decrement-ttl;
   optimize-timer seconds;
   preference preference;
   priority setup-priority hold-priority;
   (record | no-record);
   retry-limit number;
   retry-timer seconds;
   select {
      manual;
      unconditional;
   }
   standby;
}
(random | least-fill | most-fill);
(record | no-record);
revert-timer seconds;
retry-limit number;
retry-timer seconds;
secondary path-name {
   adaptive;
   admin-group {
      exclude group-names;
      include group-names;
   }
   bandwidth bps;
   class-of-service cos-value;
   hop-limit number;
   no-cspf;
   no-decrement-ttl;
   optimize-timer seconds;
   preference preference;
   priority setup-priority hold-priority;
   (record | no-record);
   select {
      manual;
      unconditional;
   }
   standby;
}
soft-preemption {
   cleanup-timer seconds;
}
standby;
to address;
```

```
      traceoptions {
         file filename <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
         flag flag <flag-modifier> <disable>;
      }
   }
   log-updown {
      (syslog | no-syslog);
      (trap | no-trap);
      trap-path-down;
      trap-path-up;
   }
   mtu-signaling;
   no-cspf;
   no-decrement-ttl;
   no-propagate-ttl;
   no-record;
   optimize-aggressive;
   optimize-timer;
   path path-name {
      address <strict | loose>;
   }
   path-mtu {
      allow-fragmentation;
      rsvp {
         mtu-signaling;
      }
   }
   policing filter filter-name;
   preference preference;
   priority setup-priority hold-priority;
   record;
   rsvp-error-hold-time seconds;
   soft-preemption {
      cleanup-timer seconds;
   }
   standby;
   static-path inet {
      prefix {
         class-of-service value;
         next-hop (address | interface-name | address/interface-name);
         push out-label;
         preference preference;
      }
   }
   statistics {
      auto-bandwidth;
      file filename size size files number <no-stamp>;
      interval seconds;
   }
   traceoptions {
      file filename <replace> <size size> <files number> <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
   }
   traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
} # End of [edit protocols mpls] hierarchy level
```

Multicast Source Discovery Protocol (MSDP)

```
msdp {
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation <(disable | enable)>;
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local address address;
  rib-group group-name;
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
  peer address {
    authentication-key peer-key;
    default-peer;
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
  }
  group group-name {
    authentication-key peer-key;
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    mode <(mesh-group | standard)>;
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
    peer address; {
      default-peer;
      disable;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
          <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
      }
    }
  }
} # End of [edit protocols msdp] hierarchy level
```

Neighbor Discovery

```
router-advertisement {
    interface interface-name {
        current-hop-limit number;
        default-lifetime seconds;
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefix prefix {
            (autonomous | no-autonomous);
            (on-link | no-on-link);
            preferred-lifetime seconds;
            valid-lifetime seconds;
        }
        reachable-time milliseconds;
        retransmit-timer milliseconds;
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                <(world-readable | no-world-readable)>;
            flag flag <detail> <disable>;
        }
    }
} # End of [edit protocols router-advertisement] hierarchy level
```

Open Shortest Path First (OSPF)

```
ospf {
    disable;
    export [ policy-names ];
    external-preference preference;
    graceful-restart {
        disable;
        helper-disable;
        notify-duration seconds;
        rest-duration seconds;
    }
    import [ policy-names ];
    overload {
        <timeout seconds>;
    }
    preference preference;
    reference-bandwidth reference-bandwidth;
    rib-group group-name;
    route-type-community (vendor | iana);
    spf-delay;
    traffic-engineering {
        no-topology;
        shortcuts {
            lsp-metric-into-summary;
        }
    }
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
```

```
area area-id {
    area-range network/masklen <restrict>;
    authentication-type authentication;
    interface interface-name {
        demand-circuit;
        disable;
        bfd-liveness-detection {
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier number;
            version (0 | 1);
        }
        authentication {
            md5 key-id; {
                key [ key-values ];
            }
            simple-password key-id;
        }
        dead-interval seconds;
        hello-interval seconds;
        interface-type type;
        metric metric;
        neighbor address <eligible>;
        passive;
        poll-interval seconds;
        priority number;
        retransmit-interval seconds;
        te-metric metric;
        transit-delay seconds;
    }
    label-switched-path name metric metric;
    nssa {
        area-range network/masklen <restrict>;
        default-lsa {
            default-metric metric;
            metric-type type;
            type-7;
        }
        (no-summaries | summaries);
    }
    peer-interface interface-name {
        disable;
        dead-interval seconds;
        hello-interval seconds;
        retransmit-interval seconds;
        transit-delay seconds;
    }
```

```
      stub <default-metric metric> < (no-summaries | summaries)>;
      virtual-link neighbor-id router-id transit-area area-id {
         authentication {
            md5 key-id;
            simple-password key-id;
         }
         dead-interval seconds;
         disable;
         hello-interval seconds;
         retransmit-interval seconds;
         transit-delay seconds;
      }
   }
} # End of [edit protocols ospf] hierarchy level
```

OSPF Version 3 (OSPFv3)

```
ospfv3 {
   disable;
   export [ policy-names ];
   external-preference preference;
   import [ policy-names ];
   overload {
      <timeout seconds>;
   }
   preference preference;
   reference-bandwidth reference-bandwidth;
   rib-group group-name;
   spf-delay;
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
   }
   area area-id {
      area-range network/mask-length <restrict>;
      interface interface-name {
         disable;
         dead-interval seconds;
         hello-interval seconds;
         ipsec-sa name;
         metric metric;
         neighbor address <eligible>;
         passive;
         priority number;
         retransmit-interval seconds;
         transit-delay seconds;
         }
      nssa {
         area-range network/mask-length <restrict>;
         default-lsa {
            default-metric metric;
            metric-type type;
            type-7;
         }
         (no-summaries | summaries)
      }
      stub <default-metric metric> <(no-summaries | summaries)>;
      virtual-link neighbor-id router-id transit-area area-id {
         disable;
         dead-interval seconds;
         hello-interval seconds;
         ipsec-sa name;
         retransmit-interval seconds;
         transit-delay seconds;
      }
   }
} # End of [edit protocols ospfv3] hierarchy level
```

Pragmatic General Multicast (PGM)

```
pgm {
   traceoptions {
      file name <replace> <size size> <files number > <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier >;
   }
} # End of [edit protocols pgm] hierarchy level
```

Protocol Independent Multicast (PIM)

```
pim {
   disable;
   assert-timeout seconds;
   dense-groups {
      addresses;
   }
   graceful-restart {
      disable;
      restart-duration seconds;
   }
   import [ policy-names ];
   interface interface-name {
      disable;
      mode (dense | sparse | sparse-dense);
      priority number;
      version version;
   }
   rib-group group-name;
   rp {
      auto-rp (announce | discovery | mapping);
      bootstrap-export [ policy-names ];
      bootstrap-import [ policy-names ];
      bootstrap-priority number;
      embedded-rp {
         maximum-rps limit;
         group-ranges {
            destination-mask;
         }
      }
      local {
         family (inet | inet6) {
            disable;
            address address;
            anycast-pim {
               rp-set {
                  address address [forward-msdp-sa];
               }
               local-address address;
            }
            group-ranges {
               destination-mask;
            }
            hold-time seconds;
            priority number;
         }
   }
   static {
      address address {
         version version;
         group-ranges {
            destination-mask;
         }
         traceoptions {
         file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
         flag flag <flag-modifier> <disable>;
      }
   }
} # End of [edit protocols pim] hierarchy level
```

Routing Information Protocol (RIP)

```
rip {
   authentication-key password;
   authentication-type type;
   (check-zero | no-check-zero);
   graceful-restart {
      disable;
      restart-time seconds;
   }
   hold-down seconds;
   import [ policy-names ];
   message-size number;
   metric-in metric ;
   receive receive-options;
   rib-group group-name;
   send send-options;
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
   }
```

```
    group group-name {
        export [ policy-names ];
        metric-out metric;
        preference preference;
        neighbor neighbor-name {
            authentication-key password;
            authentication-type type;
            (check-zero | no-check-zero);
            import [ policy-names ];
            message-size number;
            metric-in metric;
            receive receive-options;
            send send-options;
        }
    }
} # End of [edit protocols rip] hierarchy level
```

RIP Next Generation (RIPng)

```
ripng {
    graceful-restart {
        disable;
        restart-time seconds;
    }
    holddown seconds;
    import [ policy-names ];
    metric-in metric;
    receive <none>;
    send <none>;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
    group group-name {
        export [ policy-names ];
        metric-out metric;
        preference number;
        neighbor interface-name {
            import [ policy-names ];
            metric-in metric;
            receive <none>;
            send <none>;
        }
    }
} # End of [edit protocols ripng] hierarchy level
```

Router Advertisement

```
router-advertisement {
   interface interface-name {
      current-hop-limit number;
      default-lifetime seconds;
      (managed-configuration | no-managed-configuration);
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
      (other-stateful-configuration | no-other-stateful-configuration);
      prefix prefix {
         (autonomous | no-autonomous);
         (on-link | no-on-link);
         preferred-lifetime seconds;
         valid-lifetime seconds;
      }
      reachable-time milliseconds;
      retransmit-timer milliseconds;
      traceoptions {
         file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
         flag flag <detail> <disable>;
      }
   }
} # End of [edit protocols router-advertisement] hierarchy level
```

Router Discovery

```
router-discovery {
   disable;
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
         <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
   }
   interface interface-name {
      min-advertisement-interval seconds;
      max-advertisement-interval seconds;
      lifetime seconds;
   }
   address address {
      (advertise | ignore);
      (broadcast | multicast);
      (priority number | ineligible);
   }
} # End of [edit protocols router-discovery] hierarchy level
```

Resource Reservation Protocol (RSVP)

```
rsvp {
   disable;
   graceful-deletion-timeout seconds;
   graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time seconds;
      maximum-helper-restart-time seconds;
   }
```

```
    interface interface-name {
        disable;
        (aggregate | no-aggregate);
        authentication-key key;
        bandwidth bps;
        hello-interval seconds;
        link-protection {
            disable;
            bandwidth bandwidth;
            bypass bypass-name {
                bandwidth bps {
                hop-limit number;
                path address <strict | loose>;
                priority setup-priority reservation-priority;
                to address;
            }
            class-of-service cos-value;
            hop-limit number;
            max-bypasses number;
            no-node-protection;
            optimize-timer seconds;
            path address <strict | loose>;
            subscription percentage;
        }
        (reliable | no-reliable);
        subscription percentage;
        update-threshold percentage;
    }
    keep-multiplier number;
    load-balance {
        bandwidth;
    }
    peer-interface peer-interface-name {
        (aggregate | no-aggregate);
        authentication-key key;
        disable;
        hello-interval seconds;
        (reliable | no-reliable);
    }
    preemption {
        (aggressive | disabled | normal);
        soft-preemption {
            cleanup-timer seconds;
        }
    }
    refresh-time seconds;
    traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>
          <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
} # End of [edit protocols rsvp] hierarchy level
```

Session Announcement Protocol/Session Description Protocol (SAP/SDP)

```
    sap {
      disable;
      listen [ address> <port port> ];
    } # End of [edit protocols sap] hierarchy level
```

Virtual Router Redundancy Protocol (VRRP)

```
    vrrp {
      traceoptions {
        file {
          filename filename;
          files number;
          size size;
          (world-readable | no-world-readable);
        }
        flag flag;
      }
    } # End of [edit protocols vrrp] hierarchy level
} # End of [edit protocols] hierarchy level
```

[edit routing-instances] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-routers *logical-router-name*] hierarchy level.

```
routing-instances {
    routing-instance-name {
      description text;
      forwarding-options;
      instance-type (forwarding | l2vpn | no-forwarding | virtual-router | vpls | vrf);
      interface interface-name;
      no-vrf-advertise;
      route-distinguisher (as-number:number | ip-address:number);
      vrf-import [ policy-names ];
      vrf-export [ policy-names ];
      vrf-table-label;
      vrf-target {
        export community name;
        import community name;
      }
      protocols {
        bgp {
          bgp-configuration;
        }
        isis {
          isis-configuration;
        }
        l2vpn {
          l2vpn-configuration;
        }
        ldp {
          ldp-configuration;
        }
        msdp {
          msdp-configuration;
        }
    ospf {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (vendor | iana);
        ospf-configuration;
    }
    ospf3 {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (vendor | iana);
        ospf3-configuration;
    }
    pim {
        pim-configuration;
    }
    rip {
        rip-configuration;
    }
    router-discovery {
        router-discovery-configuration;
    }
    vpls {
        vpls-configuration;
    }
}
routing-options {
    aggregate {
        defaults {
            aggregate-options;
        }
        route destination-prefix {
            policy policy-name;
            aggregate-options;
        }
    }
    auto-export {
        (disable | enable);
            family {
                inet {
                    flow {
                        (disable | enable);
                        rib-group rib-group;
                    }
                    multicast {
                        (disable | enable);
                        rib-group rib-group;
                    }
                    unicast {
                        (disable | enable);
                        rib-group rib-group;
                    }
                }
            }
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
  }
}
autonomous-system autonomous-system <loops number> {
  independent-domain;
}
confederation confederation-autonomous-systems
  members autonomous-system;
dynamic-tunnels tunnel-name {
  destination-prefix prefix;
  source-address address;
  tunnel-type type-of-tunnel;
}
fate-sharing {
  group group-name;
  cost value;
  from address {
    to address;
  }
}
flow {
  route name {
    match {
      match-conditions;
    }
    then {
      actions;
    }
  }
  validation {
    traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
}
generate {
  defaults {
    generate-options;
  }
  route destination-prefix {
    policy policy-name;
    generate-options;
  }
}
```

```
instance-export [ policy-names ];
instance-import [ policy-names ];
interface-routes {
   family (inet | inet6) {
      import [ import-policies ];
      export {
         lan;
         point-to-point;
      }
   }
   rib-group {
      inet group-name;
      inet6 group-name;
}
martians {
   destination-prefix match-type <allow>;
}
maximum-routes route-limit <log-only | threshold value>;
multicast {
   forwarding-cache {
      threshold (suppress | reuse) value value;
   }
   scope scope-name {
      interface interface-name;
      prefix destination-prefix;
   }
   ssm-groups {
      addresses;
   }
}
options {
   syslog (level level | upto level);
}
resolution {
   rib routing-table-name {
      import [ policy-names ];
      resolution-ribs [ routing-table-names ];
   }
}
rib routing-table-name {
   aggregate {
      defaults {
         aggregate-options;
      }
      route destination-prefix {
         policy policy-name;
         aggregate-options;
      }
   }
   filter {
      input filter-name;
   }
   generate {
      defaults {
         generate-options;
      }
```

```
        route destination-prefix {
            policy policy-name;
            generate-options;
        }
    }
    martians {
        destination-prefix match-type <allow>;
    }
    static {
        defaults {
            static-options;
        }
        rib-group group-name;
        route destination-prefix {
            lsp-next-hop {
                metric metric;
                preference preference;
            }
            next-hop;
            qualified-next-hop address {
                metric metric;
                preference preference;
            }
            static-options;
        }
    }
}
rib-groups {
    group-name {
        import-policy [ policy-names ];
        import-rib [ group-names ];
        export-rib [ group-names ];
    }
}
route-distinguished-id address;
route-record;
router-id address;
static {
    defaults {
        static-options;
    }
    rib-group group-name;
    route destination-prefix {
        lsp-next-hop {
            metric metric;
            preference preference;
        }
        next-hop;
        qualified-next-hop address {
            metric metric;
            preference preference;
        }
        static-options;
    }
}
```

```
        traceoptions {
          file name <replace> <size size> <files number> <no-stamp>
              <(world-readable | no-world-readable)>;
          flag flag <flag-modifier> <disable>;
        }
      }
    }
} # End of [edit routing-instances] hierarchy level
```

[edit routing-options] **Hierarchy Level**

The following statement hierarchy can also be included at the [edit logical-routers *logical-router-name*] hierarchy level.

```
routing-options {
    aggregate {
      defaults {
        aggregate-options;
      }
      route destination-prefix {
        policy policy-name;
        aggregate-options;
      }
    }
    auto-export {
      (disable | enable);
      family {
        inet {
          flow {
            (disable | enable);
            rib-group rib-group;
          }
          multicast {
            (disable | enable);
            rib-group rib-group;
          }
          unicast {
            (disable | enable);
            rib-group rib-group;
          }
        }
      }
      traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
          <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
      }
    }
    autonomous-system autonomous-system <loops number>;
    confederation confederation-autonomous-system  members autonomous-system;
    dynamic-tunnels tunnel-name {
        destination-networks prefix;
        source-address address;
        tunnel-type tunnel-type;
    }
```

```
fate-sharing {
    group group-name;
    cost value;
    from address {
        to address;
    }
}
flow {
    route name {
        match {
            match-conditions;
        }
        then {
            actions;
        }
    }
    validation {
        traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
}
forwarding-table {
    export [ policy-names ];
    unicast-reverse-paths (active-paths | feasible-paths);
}
generate {
    defaults {
        generate-options;
    }
    route destination-prefix {
        policy policy-name;
        generate-options;
    }
}
graceful-restart {
    disable;
    path-selection-defer-time-limit time-limit;
}
instance-export [ policy-names ];
instance-import [ policy-names ];
interface-routes {
    family (inet | inet6) {
        export {
            lan;
            point-to-point;
        }
    }
    rib-group group-name;
}
martians {
    destination-prefix match-type <allow>;
}
```

```
maximum-routes route-limit <log-only | threshold value>;
multicast {
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
  }
  scope-policy policy-name;
  ssm-groups {
    address;
  }
  ssm-map ssm-map-name {
    policy policy-name;
    source addresses;
  }
}
options {
  syslog (level level | upto level);
}
resolution {
  rib routing-table-name {
    import [ policy-names ];
    resolution-ribs [ routing-table-names ];
  }
}
rib routing-table-name {
  aggregate {
    defaults {
      aggregate-options;
    }
    rib-group group-name;
    route destination-prefix {
      policy policy-name;
      aggregate-options;
    }
  }
  filter {
    input filter-name;
  }
  generate {
    defaults {
      generate-options;
    }
    route destination-prefix {
      policy policy-name;
      generate-options;
    }
  }
  martians {
    destination-prefix match-type <allow>;
  }
```

```
      static {
        defaults {
          static-options;
        }
        rib-group group-name;
        route destination-prefix {
            lsp-next-hop {
              metric metric;
              preference preference;
            }
            next-hop;
            qualified-next-hop address {
              metric metric;
              preference preference;
          }
          static-options;
        }
      }
    }
    rib-groups {
      group-name {
        import-policy [ policy-names ];
        import-rib [ group-names ];
        export-rib [ group-names ];
      }
    }
    route-distinguisher-id address;
    route record;
    router-id address;
    static {
      defaults {
        static-options;
      }
      rib-group group-name;
      route destination-prefix {
        lsp-next-hop {
          metric metric;
          preference preference;
        }
        next-hop;
        p2mp-lsp-next-hop {
          metric metric;
          preference preference;
        }
        qualified-next-hop-address {
          metric metric;
          preference preference;
        }
        static-options;
      }
    }
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
} # End of [edit routing-options] hierarchy level
```

[edit security] Hierarchy Level

```
security {
    certificates {
        cache-size bytes;
        cache-timeout-negative seconds;
        certification-authority ca-profile-name {
            ca-name certificate-authority-name;
            crl file-name;
            encoding (binary | pem);
            enrollment-url url-name;
            file certificate-filename;
            ldap-url url-name;
        }
        enrollment-retry number;
        local certificate-filename;
        maximum-certificates number;
        path-length bytes;
    }
    ike {
        policy ike-peer-address {
            description policy-description;
            encoding (binary | pem);
            identity identity-name;
            local certificate-filename;
            local-key-pair private-public-key-file;
            mode (aggressive | main);
            pre-shared-key (ascii-text key | hexadecimal key);
            proposals [ proposal-names ];
        }
        proposal ike-proposal-name {
            authentication-algorithm (md5 | sha1);
            authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
            dh-group (group1 | group2);
            encryption-algorithm (3des-cbc | des-cbc);
            lifetime-seconds seconds;
        }
    }
    ipsec {
        internal {
            security-association {
                manual {
                    direction (bidirectional | inbound | outbound) {
                        protocol esp;
                        spi spi-value;
                        encryption {
                            algorithm 3des-cbc;
                            key ascii-text ascii-text-string;
                        }
                    }
                }
            }
        }
```

```
            policy ipsec-policy-name {
                perfect-forward-secrecy {
                    keys (group1 | group2);
                }
                proposals [ proposal-names ];
            }
            proposal ipsec-proposal-name {
                authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
                encryption-algorithm (3des-cbc | des-cbc);
                lifetime-seconds seconds;
                protocol (ah | esp | bundle);
            }
            security-association name {
                dynamic {
                    <security-association (32 | 64)>;
                    ipsec-policy policy-name;
                }
                manual {
                    direction (JUNOS) (inbound | outbound | bi-directional) {
                        authentication {
                            algorithm (hmac-md5-96 | hmac-sha1-96);
                            key (ascii-text key | hexadecimal key);
                        }
                        auxiliary-spi auxiliary-spi-value;
                        encryption {
                            algorithm (des-cbc | 3des-cbc);
                            key (ascii-text key | hexadecimal key);
                        }
                        protocol (ah | esp | bundle);
                        spi spi-value;
                    }
                }
                mode (tunnel | transport);
                traceoptions {
                    file <files number> <size size>;
                    flag all;
                    flag database;
                    flag general;
                    flag ike;
                    flag parse;
                    flag policy-manager;
                    flag routing-socket;
                    flag timer;
                }
            }
        }
} # End of [edit security] hierarchy level
```

[edit services] hierarchy level

```
services {
    adaptive-services-pics {
        traceoptions {
            flag flag;
        }
    }
    dynamic-flow-capture {
        capture-group client-name {
            content-destination identifier {
                address address;
                ttl hops;
            }
            control-source identifier {
                allowed-destinations [ destinations ];
                no-syslog;
                notification-targets [ address address port port-number ];
                service-port port-number;
                shared-key value;
                source-addresses [ addresses ];
            }
            input-packet-rate-threshold rate;
            interfaces interface-name;
            pic-memory-threshold percentage percentage;
        }
    }
    flow-collector {
        analyzer-address address;
        analyzer-id name;
        destinations {
            ftp:url {
                password "password";
            }
        }
        file-specification {
            variant variant-number {
                data-format format;
                name-format format;
                transfer {
                    record-level number;
                    timeout seconds;
                }
            }
        }
        interface-map {
            collector interface-name;
            file-specification variant-number;
            interface-name {
                file-specification variant-number;
                collector interface-name;
            }
        }
        retry number;
        retry-delay seconds;
        transfer-log-archive {
            archive-sites {
```

```
            ftp:url {
                password “password”;
                username username;
            }
        }
        filename-prefix prefix;
        maximum-age minutes;
    }
}
ids {
    rule rule-name {
        match-direction (input | output | input-output);
        term term-name {
            from {
                applications [ application-names ];
                application-sets [ set-names ];
                destination-address address;
                source-address address;
            }
            then {
                aggregation {
                    destination-prefix prefix-value;
                    source-prefix prefix-value;
                }
            (force-entry | ignore entry);
            logging {
                syslog;
                threshold rate;
            }
            session-limit {
                by-destination {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
                by-pair {
                    maximum number;
                    packets number;
                    rate number;
                }
                by-source {
                    hold-time seconds;
                    maximum number;
                    packets number;
                    rate number;
                }
            }
            syn-cookie {
                mss value;
                threshold rate;
            }
        }
    }
    rule-set rule-set-name {
        [ rule rule-names ];
    }
}
ipsec-vpn {
    ike {
        proposal proposal-name {
            authentication-algorithm (md5 | sha1);
            authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
            description description;
            dh-group (group1 | group2);
            encryption-algorithm (3des-cbc | des-cbc);
            lifetime-seconds seconds;
        }
        policy policy-name {
            description description;
            local-id {
                fqdn [ values ];
                ipv4_addr [ values ];
                key_id [ values ];
            }
            mode (aggressive | main);
            pre-shared-key (ascii-text key | hexadecimal key);
            proposals [ proposal-names ];
            remote-id {
                fqdn [ values ];
                ipv4_addr [ values ];
                key_id [ values ];
            }
        }
    }
    ipsec {
        proposal proposal-name {
            authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
            description description;
            encryption-algorithm (3des-cbc | des-cbc);
            lifetime-seconds seconds;
            protocol (ah | esp | bundle);
        }
        policy policy-name {
            description description;
            perfect-forward-secrecy {
                keys (group1 | group2);
            }
            proposals [ proposal-names ];
        }
    }
    rule rule-name {
        match-direction (input | output);
        term term-name {
            from {
                destination-address address;
                ipsec-inside-interface interface-name;
                source-address address;
            }
```

```
      then {
        backup-remote-gateway address;
        dynamic {
          ike-policy policy-name;
          ipsec-policy policy-name;
        }
        manual (
          direction (inbound | outbound | bidirectional) {
            authentication {
              algorithm (hmac-md5-96 | hmac-sha1-96);
              key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi spi-value;
            encryption {
              algorithm (des-cbc | 3des-cbc);
              key (ascii-text key | hexadecimal key);
            }
            protocol (ah | bundle | esp);
            spi spi-value;
          }
        }
        no-anti-replay;
        remote-gateway address;
        syslog;
      }
    }
  }
  rule-set rule-set-name {
    [ rule rule-names ];
  }
}
l2tp {
  tunnel-group group-name {
    hello-interval seconds;
    hide-avps;
    l2tp-access-profile profile-name;
    local-gateway address address;
    maximum-send-window packets;
    ppp-access-profile profile-name;
    receive-window packets;
    retransmit-interval seconds;
    service-interface interface-name;
    syslog {
      host hostname {
        facility-override facility-name;
        log-prefix prefix-number;
        services severity-level;
      }
    }
    tunnel-timeout seconds;
  }
  traceoptions {
    debug-level level;
    filter {
      protocol name;
    }
```

```
      flag flag;
      interfaces interface-name {
         debug-level level;
         flag flag;
      }
   }
}
nat {
   pool nat-pool-name {
      address (address | address-range low value high value | prefix);
      port (automatic | range low minimum-value high maximum-value);
   }
   rule rule-name {
      match-direction (input | output);
      term term-name {
         from {
            applications [ application-names ];
            application-sets [ set-names ];
            destination-address address;
            source-address (address | prefix);
         }
         then {
            translated {
               destination-pool nat-pool-name;
               source-pool nat-pool-name;
               translation-type (destination type | source type);
            }
            syslog;
         }
      }
   }
   rule-set rule-set-name {
      [ rule rule-names ];
   }
}
rpm {
   probe owner {
      test test-name {
            data-fill data;
            data-size size;
            destination-port port;
            dscp-code-point DSCP bits;
            history-size size;
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance routing-instance-name;
            source-address address;
            target-url (url | address);
            test-interval interval;
            thresholds thresholds;
            traps traps;
      }
   }
```

```
         probe-server {
            tcp port;
            udp port;
         }
         probe-limit limit;
         }
      }
      service-set service-set-name {
         ([ ids-rules rule-names ] | ids-rule-sets rule-set-name);
         ([ipsec-vpn-rules rule-names ] | ipsec-vpn-rule-sets rule-set-name);
         ([ nat-rules rule-names ] | nat-rule-sets rule-set-name);
         ([ stateful-firewall-rules rule-names ] | stateful-firewall-rule-sets rule-set-name);
         interface-service {
            service-interface interface-name;
         }
         ipsec-vpn-options {
            ike-access-profile profile-name;
            local-gateway address;
         }
         max-flows number;
         next-hop-service {
            inside-service-interface name.number;
            outside-service-interface name.number;
         }
         syslog {
            host hostname {
               facility-override facility-name;
               log-prefix prefix-number;
               services priority-level;
            }
         }
      }
      stateful-firewall {
         rule rule-name {
            match-direction (input | output | input-output);
            term term-name {
                  from {
                  applications [ application-names ];
                  application-sets [ set-names ];
                  destination-address address;
                  source-address address;
               }
               then {
                  (accept | discard | reject);
                  allow-ip-option { values ]
                  syslog;
               }
            }
         }
         rule-set rule-set-name {
            [ rule rule-names ];
         }
      }
} # End of [edit services] hierarchy level
```

[edit snmp] Hierarchy Level

```
snmp {
    community community-name {
        authorization authorization;
        clients {
            address restrict;
        }
        view view-name;
    }
    contact contact;
    description description;
    engine-id {
        (local engine-id | use-mac-address | use-default-ip-address);
    }
    interface [ interface-name ];
    filter-duplicates;
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            interval seconds;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            variable oid-variable;
        }
        event index {
            community community-name;
            description description;
            type type;
        }
    }
    traceoptions {
        file size size files number;
        flag flag;
    }
    trap-group group-name {
        categories [ categories ];
        destination-port <port-number>;
        targets {
            address;
        }
        version (all | v1 | v2);
    }
    trap-options {
        agent-address outgoing-interface;
        source-address address;
    }
```

```
v3 {
    notify name {
        tag tag-name;
        type (trap | inform);
    }
    notify-filter name {
        oid oid (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        inform-retry-count number;
        inform-timeout seconds;
        port port-number;
        tag-list tag-list;
        target-parameters target-parameters-name;
    }
    target-parameters target-parameters-name {
        notify-filter name;
        parameters {
            message-processing-model (v1 | v2c | v3);
            security-model ( usm | v1 | v2c);
            security level (authentication | none | privacy);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password password;
                }
                authentication-sha {
                    authentication-password password;
                }
                authentication-none;
                privacy-3des {
                    privacy-password password;
                }
                privacy-aes128 {
                    privacy-password password;
                }
                privacy-des {
                    privacy-password password;
                }
                privacy-none;
            }
        }
    }
```

```
      vacm {
        access {
          group group-name {
            default-context-prefix {
              security-model (any | usm | v1 | v2c) {
                security-level (authentication | none | privacy) {
                  notify-view notify-view;
                  read-view read-view;
                  write-view write-view;
                }
              }
            }
          }
        }
        security-to-group {
          security-model (usm | v1 | v2c) {
            security-name security-name {
              group group-name;
            }
          }
        }
      }
      view view-name; {
          oid object-identifier (include | exclude);
      }
    }
} # End of [edit snmp] hierarchy level
```

[edit system] Hierarchy Level

```
system {
    accounting {
      destination {
        radius {
          server {
            server-address {
              accounting-port port-number;
              retry number;
              secret password;
              source-address address;
              timeout seconds;
            }
          }
        }
        tacplus {
          server {
            server-address {
              port port-number;
              secret password;
              single-connection;
              timeout seconds;
            }
          }
        }
      }
      events [ login change-log interactive-commands ];
    }
```

```
archival {
   configuration {
      archive-sites {
         ftp://<username>:<password>@<host>:<port>/<url-path>;
      }
      transfer-interval interval;
      transfer-on-commit;
   }
}
arp {
   passive-learning;
   aging-timer minutes;
}
authentication-order [ authentication-methods ];
backup-router address <destination destination-address>;
building name;
commit synchronize;
(compress-configuration-files | no-compression-configuration-files);
default-address-selection;
diag-port-authentication (encrypted-password "password" | plain-text-password);
domain-name domain-name;
domain-search [domain-list];
host-name host-name;
internet-options address <destination destination-address>;
internet-options {
   path-mtu-discovery;
   source-quench;
   source-port upper-limit <upper-limit>;
}
location {
   altitude feet;
   building name;
   country-code code;
   floor number;
   hcoord horizontal-coordinate;
   lata service-area;
   latitude degrees;
   longitude degrees;
   npa-nxx number;
   postal-code postal-code;
   rack number;
   vcoord vertical-coordinate;
}
login {
   announcement text;
   class class-name {
      allow-commands "regular-expression" ;
      allow-configuration "regular-expression" ;
      deny-commands "regular-expression";
      deny-configuration "regular-expression";
      idle-timeout minutes;
      login-tip;
      permissions [ permissions ];
   }
```

```
    message text;
    passwords {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    user username {
        full-name complete-name;
        uid uid-value;
        class class-name;
        authentication {
            (encrypted-password "password" | plain-text-password);
            ssh-rsa "public-key";
            ssh-dsa "public-key";
        }
    }
}
login-tip number;
mirror-flash-on-disk;
name-server {
    address;
}
no-redirects;
ntp {
    authentication-key key-number type type value password;
    boot-server (NTP) address;
    broadcast <address> <key key-number> <version value> <ttl value>;
    broadcast-client;
    multicast-client <address>;
    peer address <key key-number> <version value> <prefer>;
    server address <key key-number> <version value> <prefer>;
    trusted-key [ key-numbers ];
}
pic-console-authentication {
    encrypted-password encrypted-password;
    plain-text-password;
}
ports {
    auxiliary {
        type terminal-type;
    }
    console {
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
```

```
processes {
    adaptive-services (enable | disable) failover failover-option;
    alarm-control (enable | disable) failover failover-option;
    chassis-control (enable | disable) failover failover-option;
    class-of-service (enable | disable) failover failover-option;
    craft-control (enable | disable) failover failover-option;
    disk-monitoring (enable | disable) failover failover-option;
    ecc-error-logging (enable | disable) failover failover-option;
    firewall (enable | disable) failover failover-option;
    inet-process (enable | disable) failover failover-option;
    interface-control (enable | disable) failover failover-option;
    kernel-replication (enable | disable) failover failover-option;
    l2tp-service (enable | disable) failover failover-option;
    link-management (enable | disable) failover failover-option;
    mib-process (enable | disable) failover failover-option;
    network-access-service (enable | disable) failover failover-option;
    ntp (enable | disable) failover failover-option;
    pgm (enable | disable) failover failover-option;
    pic-services-logging (enable | disable) failover failover-option;
    pppoe (enable | disable) failover failover-option;
    redundancy-device (enable | disable) failover failover-option;
    remote-operations (enable | disable) failover failover-option;
    routing (enable | disable) failover failover-option;
    sampling (enable | disable) failover failover-option;
    service-deployment (enable | disable) failover failover-option;
    snmp (enable | disable)  failover failover-option;
    timeout seconds;
    watchdog (enable | disable) failover failover-option;
    web-management (enable | disable) failover failover-option;
}
radius-server server-address {
    port number ;
    retry number ;
    routing-instance routing-instance-name;
    secret password ;
    source-address source-address;
    timeout seconds ;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key" ;
    ssh-dsa "public-key" ;
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
scripts {
    commit {
        allow-transients;
        file filename.xsl {
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        traceoptions {
            file filename <files number> <size size>;
            flag flag;
        }
    }
}
services {
    finger {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    ftp {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    service-deployment {
        servers server-address {
            port port-number;
        }
        source-address source-address;
    }
    ssh {
        root-login (allow | deny | deny-password);
        protocol-version [v1 v2];
        <connection-limit limit>;
        <rate-limit limit>;
    }
    telnet {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    web-management {
        http {
            interfaces [ interface-names ];
            port port;
        }
        https {
            interfaces [ interface-names ];
            local-certificate name;
            port port;
        }
    }
    xnm-clear-text {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    xnm-ssl {
        <connection-limit limit>;
        local-certificate name;
        <rate-limit limit>;
    }
}
```

```
    static-host-mapping {
      host-name {
        alias [ alias ];
        inet [ address ];
        sysid system-identifier;
      }
    }
    syslog {
      archive {
        files number;
        size size;
        (world-readable | no-world-readable);
      }
      console {
        facility severity;
      }
      file filename {
        facility severity;
        explicit-priority;
        match "regular-expression";
        archive {
          files number;
          size size;
          (world-readable | no-world-readable);
        }
      }
      host (hostname | other-routing-engine | scc-master) {
        facility severity;
        explicit-priority;
        facility-override facility;
        log-prefix string;
        match "regular-expression";
      }
      source-address source-address;
      time-format (year | millisecond | year millisecond);
      user (username | *) {
        facility level;
        match "regular-expression";
      }
    }
    tacplus-options service-name service-name;
    tacplus-server server-address {
      secret password;
      single-connection;
      source-address source-address;
      timeout seconds;
    }
    time-zone (GMThour-offset | time-zone);
} # End of [edit system] hierarchy level
```

Chapter 4

Complete Configuration Mode Commands and Statements for J-series Services Routers

This chapter shows the complete configuration mode commands and the complete configuration statement hierarchy for J-series Services Routers. The J-series Services Router configuration statement hierarchy includes some statements from the full JUNOS configuration (for M-series and T-series platforms) and also has some unique configuration statements. Using these commands and statements is described in the *J-series Services Router Administration Guide*.

- Complete Configuration Mode Commands on page 126
- Complete Configuration Statement Hierarchy on page 127

For information about command-line interface (CLI) operational mode commands, see the *JUNOS System Basics and Services Command Reference*. For information about the complete configuration mode commands and statements for M-series and T-series platforms, see "Complete Configuration Mode Commands and Statements for M-series and T-series Platforms" on page 47.

Complete Configuration Mode Commands

The following is the complete list of configuration mode commands, listing all possible commands in the hierarchy.

```
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 an ASCII file
  quit          Quit from this level
  rename        Rename a statement
  rollback      Roll back database to last committed version
  run           Run an operational-mode command
  save          Save configuration to an ASCII file
  set           Set a parameter
  show          Show a parameter
  status        Display database user status
  top           Exit to top level of configuration
  up            Exit one level of configuration
  wildcard      Wildcard operations
```

Complete Configuration Statement Hierarchy

This section shows the complete configuration statement hierarchy for J-series Services Routers, listing all supported configuration statements and showing their level in the configuration hierarchy. When you are configuring the JUNOS software, your current hierarchy level is shown in the banner on the line preceding the user@host# prompt.

This section is organized as follows:

[edit access] Hierarchy Level

```
access {
    profile profile-name {
      authentication-order [ authentication-methods ];
      client name {
        chap-secret chap-secret;
      }
    }
    traceoptions {
      file filename {
        files number;
        size size;
      }
      flag all;
      flag authentication;
      flag chap;
      flag configuration;
      flag kernel;
      flag radius;
      flag world-readable;
    }
} # End of [edit access] hierarchy level
```

[edit accounting-options] Hierarchy Level

```
accounting-options {
    class-usage-profile profile-name {
      destination-classes {
        class-name;
      }
      source-classes {
        class-name;
      }
    }
    file filename {
      size bytes;
      files number;
      transfer-interval minutes;
      archive-sites {
        url;
      }
    }
    filter-profile profile-name {
      file filename;
      interval minutes;
      counters {
        counter-name;
      }
    }
    interface-profile profile-name {
      fields {
        field-name;
      }
      file filename;
      interval minutes;
    }
    routing-engine-profile profile-name {
        file name;
        fields {
            field-name;
        }
        interval seconds;
    }
} # End of [edit accounting-options] hierarchy level
```

[edit applications] Hierarchy Level

```
applications {
    application application-name {
        application-protocol protocol-name;
        destination-port port-number;
        icmp-code value;
        icmp-type value;
        inactivity-timeout value;
        learn-sip-register;
        protocol type;
        rpc-program-number number;
        sip-call-hold-timeout minutes;
        snmp-command command;
        source-port port-number;
        ttl-threshold value;
        uuid hex-value;
    }
    application-set application-set-name {
        [ application application-name ];
    }
} # End of [edit applications] hierarchy level
```

[edit chassis] Hierarchy Level

```
chassis {
    alarm {
        interface-type {
            alarm-name (ignore | red | yellow);
        }
    }
    config-button {
        no-clear;
        no-rescue;
    }
    fpc slot-number {
        pic pic-number {
            mlfr-uni-nni-bundles number;
            q-pic-large-buffer;
        }
    }
    (source-route | no-source-route);
} # End of [edit chassis] hierarchy level
```

[edit class-of-service] Hierarchy Level

```
class-of-service {
    adaptive-shapers {
      adaptive-shaper-name {
        trigger type shaping-rate (percent percent | rate);
      }
    }
    classifiers {
      type classifier-name {
        forwarding-class class-name {
          loss-priority (low | high) code-points [ alias | bits ];
        }
        import (classifier-name | default);
      }
    }
    code-point-aliases {
      (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) {
        alias-name bits;
      }
    }
    drop-profiles {
      profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
          drop-probability value;
          fill-level value;
        }
      }
    }
    fabric {
      scheduler-map {
        priority (low | high) scheduler scheduler-name;
      }
    }
    forwarding-classes {
      queue queue-number class-name priority (low | high);
    }
    forwarding-policy {
      class class-name {
        classification-override {
          forwarding-class class-name;
        }
      }
      next-hop-map map-name {
        forwarding-class class-name {
          next-hop [ next-hop-name ];
          lsp-next-hop [ lsp-regular-expression ];
        }
      }
    }
```

```
fragmentation-maps {
map-name {
    forwarding-class class-name {
        fragment-threshold bytes;
        multilink-class number;
        no-fragmentation;
    }
}
interfaces {
    interface-name {
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        unit logical-unit-number {
            adaptive-shaper adaptive-shaper-name;
            classifiers {
                (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence)
                        (classifier-name | default);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            loss-priority-maps {
                (map-name | default);
            }
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                frame-relay-de (rewrite-name | default);
                ieee-802.1 default;
                inet-precedence (rewrite-name | default);
            }
            scheduler-map map-name;
            shaping-rate rate;
            virtual-channel-group virtual-channel-group-name;
        }
    }
}
loss-priority-maps {
    frame-relay-de (map-name | default) {
        loss-priority level code-points [ values ];
    }
}
restricted-queues {
    forwarding-class class-name queue queue-number;
}
rewrite-rules {
    (dscp | exp | inet-precedence) rewrite-name {
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
    }
}
```

```
    routing-instances routing-instance-name;
      classifier {
        exp (classifier-name | default);
      }
    }
    scheduler-maps {
      map-name {
        forwarding-class class-name scheduler scheduler-name;
      }
    }
    schedulers
      scheduler-name {
        buffer-size (percent percentage | remainder | temporal microseconds);
        drop-profile-map loss-priority (low | high) protocol (non-tcp | tcp | any)
          drop-profile profile-name;
        priority (low | high | strict-high);
        transmit-rate (rate | percent percentage | remainder | exact);
      }
    }
    tri-color;
    virtual-channels {
      virtual-channel-name;
    }
    virtual-channel-groups {
      virtual-channel-group-name {
        virtual-channel-name {
          scheduler-map map-name;
          shaping-rate (percent percent | rate);
          default;
        }
      }
    }
} # End of [edit class-of-service] hierarchy level
```

[edit firewall] Hierarchy Level

```
firewall {
    family family-name {
      filter filter-name {
        term term-name {
          from {
            match-conditions;
          }
          then {
            action;
            action-modifiers;
          }
        }
        virtual-channel virtual-channel-name;
      }
      service-filter filter-name {
        term term-name {
          from {
            match-conditions;
          }
```

```
               then {
                  action;
                  action-modifiers;
               }
            }
         }
      }
      filter filter-name {
         term term-name {
            from {
               match-conditions;
            }
            then {
               action;
               action-modifiers;
            }
         }
      }
      interface-set interface-set-name {
         [ interface-names ];
      }
      load-balance-group group-name {
         next-hop-group [ group-names ];
      }
      policer policer-name;
         filter-specific;
         if-exceeding {
            bandwidth-limit bps;
            bandwidth-percent number;
            burst-size-limit bytes;
         }
         logical-interface-policier;
         then {
            policer-action;
         }
      }
      three-color-policer name {
         two-rate {
            (color aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            excess-burst-size bytes;
            peak-information-rate bps;
            peak-burst-size bytes;
         }
      }
} # End of [edit firewall] hierarchy level
```

[edit forwarding-options] Hierarchy Level

```
forwarding-options {
    accounting group-name {
        output {
            cflowd [ host-names ] {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                port port-number;
                version format;
            }
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            interface [ interface-names ] {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
    family family-name {
        filter {
            input filter-name;
        }
    }
    hash-key {
        family inet {
            layer-3;
            layer-4;
        }
    }
    helpers {
        bootp {
            description description-of-service;
            interface interface-group {
                description description-of-interface;
                maximum-hop-count number;
                minimum-wait-time seconds;
                no-listen;
                server [ addresses ];
            }
            maximum-hop-count number;
            minimum-wait-time seconds;
            server address;
        }
    domain {
        description description-of-service;
        server address;
        interface interface-name {
        description description-of-interface;
        no-listen;
        server address;
        }
    }
    tftp {
        description description-of-service;
        server address < [ routing-instance routing-instance-name ] >;
        interface interface-name {
            description description-of-interface;
            no-listen;
            server address;
        }
    }
    traceoptions {
        file {
            files number;
            size kilobytes;
        }
        flag flag;
        level level;
    }
}
sampling {
    disable;
    input {
        family inet {
            max-packets-per-second number;
            rate number;
            run-length number;
        }
    }
    output {
    aggregate-export-interval seconds;
    cflowd [ host-names ] {
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        source-address address;
        port port-number;
        version format;
    }
```

```
        file {
            disable;
            filename filename;
            files number;
            size bytes;
            (stamp | no-stamp);
            (world-readable | no-world-readable);
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface [ interface-names ] {
            engine-id num;
            engine-type num;
            source-address address;
        }
    }
} # End of [edit forwarding-options] hierarchy level
```

[edit groups] Hierarchy Level

```
groups {
    group-name {
        configuration-data;
    }
} # End of [edit groups] hierarchy level
```

[edit interfaces] Hierarchy Level

```
interfaces {
    interface-name {
        disable;
        accounting-profile name;
        description text;
        clocking clock-source;
        dce;
        description text;
        dialer-options {
            pool pool-name {
                priority priority;
            }
        }
        disable;
        e1-options {
            bert-error-rate rate;
            bert-period seconds;
            fcs (32 | 16);
            framing (g704 | g704-no-crc4 | unframed);
            idle-cycle-flag (flags | ones);
            invert data;
            loopback (local | remote);
            start-end-flag (shared | filler);
            timeslots time-slot-range;
        }
        encapsulation type;
        failure-options {
            [ trigger-link-failure interface-name ];
        }
```

```
fastether-options {
   (flow-control | no-flow-control);
   ingress-rate-limit rate;
   (loopback | no-loopback);
   source-address-filter {
      mac-address;
   }
   (source-filtering | no-source-filtering);
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
isdn-options {
   bandwidth (64 | 56);
   interface-type (user | network);
   switch-type (NI2 | ETSI | ATT5E);
   spid1 spid-string;
   spid2 spid-string;
   tei-option (first-call | power-up);
   t306 seconds;
   t310 seconds;
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
   lmi-type (ansi | itu);
   n391dte number;
   n392dce number;
   n392dte number;
   n393dce number;
   n393dte number;
   t391dte seconds;
   t392dce seconds;
}
mac mac-address;
mtu bytes;
no-keepalives;
per-unit-scheduler;
ppp-options {
   chap {
      access-profile name;
      local-name name;
      passive;
   }
   compression {
      acfc;
      pfc;
   }
}
receive-bucket {
   overflow (discard | tag);
   rate percentage;
   threshold bytes;
}
redundancy-options {
   primary sp-fpc/pic/port;
   secondary sp-fpc/pic/port;
}
```

```
serial-options {
    clock-rate rate;
    clocking-mode (dce | dte | loop);
    control-leads {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    control-polarity (positive | negative);
    cts-polarity (positive | negative);
    dcd-polarity (positive | negative);
    dsr-polarity (positive | negative);
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (positive | negative);
    encoding (nrz | nrzi);
    indication-polarity (positive | negative);
    line-protocol protocol;
    loopback mode;
    rts-polarity (positive | negative);
    tm-polarity (positive | negative);
    transmit-clock invert;
}
speed (10m | 100m);
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx64 | nx56);
    fcs (32 | 16);
    framing (sf | esf);
    idle-cycle-flags (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (shared | filler);
    timeslots slot-number;
}
t3-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink)
            <subrate value>;
    fcs (32 | 16);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local |payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag value;
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}
transmit-bucket {
    overflow (tag | discard);
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
vlan-tagging;
unit logical-unit-number {
    accounting-profile name;
    backup-options {
        interface interface-name;
    }
    bandwidth rate;
    description text;
    dial-options {
        activation-delay seconds;
        bearer-cap (udi | rdi | speech);
        deactivation-delay seconds;
        (dedicated | shared);
        pool pool;
        dial-string [dial-string-numbers];
        idle-timeout seconds;
        initial-route-check seconds;
        ipsec-interface-id name;
        l2tp-interface-id name;
        load-threshold number;
        re-enable seconds;
        watch-list {
            routes;
        }
    }
    disable;
    dlci dlci-identifier;
    encapsulation type;
    inverse-arp;
    multicast-dlci dlci-identifier ;
    multipoint;
    passive-monitor-mode;
    point-to-point;
    (traps | no-traps);
    vlan-id number;
    family family {
        address address {
            destination address;
        }
```

```
            filter {
                input filter-name;
                output filter-name;
                group filter-group-number ;
            }
            mtu bytes;
            no-redirects;
            policer {
                arp policer-template-name;
                output policer-template-name;
                group policer-template-name;
            }
            primary;
            rpf-check <fail-filter filter-name> {
                <mode loose>;
            }
            sampling {
                [ input output ];
            }
            service {
                input {
                    [ service-set service-set-name <service-filter filter-name> ];
                    post-service-filter filter-name; {
                }
                output {
                    [ service-set service-set-name <service-filter filter-name> ];
                }
            }
            address address {
                arp ip-address (mac | multicast-mac) mac-address <publish>;
                destination destination-address;
                broadcast address;
                preferred;
                primary;
                vrrp-group group-number {
                    (accept-data | no-accept-data):
                    advertise-interval seconds;
                    authentication-type authentication;
                    authentication-key key;
                    fast-interval milliseconds;
                    (preempt | no-preempt) {
                        hold-time seconds;
                    }
                    priority number;
                    }
                    track {
                        interface interface-name priority-cost cost;
                    }
                    virtual-address [ addresses ];
                }
            }
        }
    }
}
```

```
    traceoptions {
      file filename <files number> <size size>
                  <(world-readable | no-world-readable)>;
      flag flag <disable>;
    }
} # End of [edit interfaces] hierarchy level
```

[edit policy-options] Hierarchy Level

```
policy-options {
    as-path name regular-expression;
    as-path-group group-name;
    community name {
      invert-match;
      members [ community-ids ];
    }
    damping name {
      disable;
      half-life minutes;
      max-suppress minutes;
      reuse number;
      suppress number;
    }
    policy-statement policy-name {
      term term-name {
        from {
          family family-name;
          match-conditions;
          policy subroutine-policy-name;
          prefix-list name;
          route-filter destination-prefix match-type <actions>;
          source-address-filter destination-prefix match-type <actions>;
        }
        to {
          match-conditions;
          policy subroutine-policy-name;
        }
        then actions;
        default-action (accept | reject);
      }
    }
    prefix-list name {
      ip-addresses;
    }
} # End of [edit policy-options] hierarchy level
```

[edit protocols] Hierarchy Level

```
protocols {
```

Bidirectional Forwarding Detection (BFD)

```
    bfd {
      traceoptions {
        file {
          filename;
          files number;
          no-world-readable;
          size bytes;
          world-readable;
        }
      }
      flag {
        adjacency;
        all;
        error;
        event;
        packet;
        pipe;
        state;
        timer;
      }
    } # End of [edit protocols bfd] hierarchy level
```

Border Gateway Protocol (BGP)

```
    bgp {
      advertise-inactive;
      advertise-peer-as;
      authentication-key key;
      cluster cluster-identifier;
      damping;
      description text-description;
      disable;
      export [ policy-names ];
      family {
          (any | multicast | unicast) {
            prefix-limit {
              maximum number;
              teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group group-name;
          }
          labeled-unicast {
            aggregate-label {
              community community-name;
            }
            explicit-null {
              connected-only;
            }
            prefix-limit {
              maximum number;
              teardown <percentage> <idle-timeout (forever | minutes)>;
            }
      resolve-vpn;
      rib inet.3;
      rib-group group-name;
    }
  }
  route-target {
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  signaling {
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
}
graceful-restart {
  disable;
  restart-time seconds;
  stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
include-mp-next-hop;
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
multihop {
  <ttl-value>;
  no-nexthop-change;
}
no-advertise-peer-as;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
path-selection (cisco-non-deterministic | always-compare-med |
  external-router-id);
peer-as autonomous-system;
preference preference;
remove-private;
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

```
vpn-apply-export;
group group-name {
   advertise-inactive;
   advertise-peer-as;
   [network/ mask-length];
   as-override;
   authentication-key key;
   cluster cluster-identifier;
   damping;
   description text-description;
   export [ policy-names ];
   family {
         (any | multicast | unicast) {
            explicit-null {
               connected-only;
            }
            prefix-limit {
               maximum number;
               teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group group-name;
         }
         flow {
            no-validate policy-name;
         }
         labeled-unicast {
            prefix-limit {
               maximum number;
               teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
            rib-group group-name;
         }
      }
      route-target {
         advertise-default;
         external-paths number;
         prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
         }
      }
      signaling {
         prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
         }
      }
   }
   graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
   }
```

```
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath;
no-advertise-peer-as;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
protocol protocol;
remove-private;
traceoptions {
   file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
   flag flag <flag-modifier> <disable>;
}
type type;
vpn-apply-export;
neighbor address {
   advertise-inactive;
   advertise-peer-as;
   as-override;
   authentication-key key;
   cluster cluster-identifier;
   damping;
   description text-description;
   export [ policy-names ];
   family {
         (any | multicast | unicast) {
            explicit-null {
               connected-only;
            }
            prefix-limit {
               maximum number;
               teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group group-name;
         }
         flow {
            no-validate policy-name;
         }
         labeled-unicast {
            prefix-limit {
               maximum number;
               teardown <percentage> <idle-timeout (forever | minutes)>;
            }
```

```
                        resolve-vpn;
                        rib inet.3;
                        rib-group group-name;
                    }
                }
                route-target {
                    advertise-default;
                    external-paths number;
                    prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                }
                signaling {
                    prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                }
            }
            graceful-restart {
                disable;
                restart-time seconds;
                stale-routes-time seconds;
            }
            hold-time seconds;
            import [ policy-names ];
            ipsec-sa ipsec-sa;
            keep (all | none);
            local-address address;
            local-as autonomous-system <private>;
            local-interface interface-name;
            local-preference local-preference;
            log-updown;
            metric-out (metric | minimum-igp <offset> | igp <offset>);
            mtu-discovery;
            multihop <ttl-value>;
            multipath;
            no-advertise-peer-as;
            no-aggregator-id;
            no-client-reflect;
            out-delay seconds;
            passive;
            peer-as autonomous-system;
            preference preference;
            remove-private;
            traceoptions {
                file name <replace> <size size> <files number> <no-stamp>
                    <(world-readable | no-world-readable)>;
                flag flag <flag-modifier> <disable>;
            }
            vpn-apply-export;
        }
    }
} # End of [edit protocols bgp] hierarchy level
```

Distance Vector Multicast Routing Protocol (DVMRP)

```
dvmrp {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    interface interface-name {
        disable;
        hold-time seconds;
        metric metric;
        mode (forwarding | unicast-routing);
    }
    rib-group group-name;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
                           <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
} # End of [edit protocols dvmrp] hierarchy level
```

Data-Link Switching (DLSw)

```
dlsw {
    connection-idle-timeout time;
    local-peer peer-address;
    multicast-address address;
    promiscuous;
    receive-initial-pacing count;
    remote-peer peer-address;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
          <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
} #End of [edit protocols dlsw] hierarchy level
```

Internet Group Management Protocol (IGMP)

```
igmp {
    interface interface-name {
        disable;
        ssm-map ssm-map-name;
        static {
            group group {
                source source;
            }
        }
        version version;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
                           <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
        }
    }
} # End of [edit protocols igmp] hierarchy level
```

End System-to-Intermediate System (ES-IS)

```
esis {
   interface [(interface-name | all)] {
      hello-interval seconds;
      hold-time seconds;
}
} # End of [edit protocols esis] hierarchy level
```

Intermediate System-to-Intermediate System (IS-IS)

```
isis {
   clns-routing;
   disable;
   export [ policy-names ];
   graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
   }
   ignore-attached-bit;
   level level-number {
      authentication-key key;
      authentication-type authentication;
      external-preference preference;
      no-csnp-authentication;
      no-hello-authentication;
      no-psnp-authentication;
      preference preference;
      prefix-export-limit num;
      wide-metrics-only;
   }
   loose-authentication-check;
   lsp-lifetime seconds;
   no-authentication-check;
   no-ipv4-routing;
   overload {
      advertise-high-metrics;
      <timeout seconds>;
   }
   reference-bandwidth reference-bandwidth;
   rib-group group name;
   spf-delay milliseconds;
   topologies {
      ipv4-multicast;
   }
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>;
                  <(world-readable | no-world-readable)>;
   flag flag <flag-modifier> <disable>;
   }
   interface interface-name {
      disable;
      bfd-liveness-detection {
         minimum-interval milliseconds;
         minimum-receive-interval milliseconds;
         minimum-transmit-interval milliseconds;
         multiplier number;
         version (0 | 1);
      }
```

```
        checksum;
        csnp-interval (seconds | disable);
        lsp-interval milliseconds;
        mesh-group (value | blocked);
        no-clns-unicast;
        no–ipv4-multicast;
        passive;
        point-to-point;
        level level-number {
            clns-unicast-metric metric;
            disable;
            hello-authentication-key key;
            hello-authentication-type authentication;
            hello-interval seconds;
            hold-time seconds;
            ipv4-multicast-metric number;
            metric metric;
            passive;
            priority number;
        }
    }
} # End of [edit protocols isis] hierarchy level
```

Multicast Source Discovery Protocol (MSDP)

```
msdp {
    active-source-limit {
        maximum number;
        threshold number;
    }
    data-encapsulation <(disable | enable)>;
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local address address;
    rib-group group-name;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
                   <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
    peer address {
        authentication-key peer-key;
        default-peer;
        disable;
        export [ policy-name ];
        import [ policy-name ];
        local-address address;
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                       <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
        }
    }
```

```
   group group-name {
      authentication-key peer-key;
      disable;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      mode <(mesh-group | standard)>;
      traceoptions {
         file name <replace> <size size> <files number> <no-stamp>
                        <(world-readable | no-world-readable)>;
         flag flag <flag-modifier> <disable>;
      }
      peer address; {
         default-peer;
         disable;
         export [ policy-name ];
         import [ policy-name ];
         local-address address;
         traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                           <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
         }
      }
   }
} # End of [edit protocols msdp] hierarchy level
```

Open Shortest Path First (OSPF)

```
ospf {
   disable;
   export [ policy-names ];
   external-preference preference;
   graceful-restart {
      disable;
      helper-disable;
      notify-duration seconds;
      rest-duration seconds;
   }
   import [policy-names];
   overload {
      <timeout seconds>;
   }
   preference preference;
   reference-bandwidth reference-bandwidth;
   rib-group group-name;
   route-type-community (vendor | iana);
   spf-delay;
   traffic-engineering {
      no-topology;
      shortcuts {
         lsp-metric-into-summary;
      }
   }
   traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
                 <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
   }
```

```
area area-id {
    area-range network/masklen <restrict>;
    authentication-type authentication;
    interface interface-name {
        demand-circuit;
        disable;
        bfd-liveness-detection {
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier number;
            version (0 | 1);
        }
        authentication {
            md5 key-id {
                key [ key-values ];
            }
            simple-password key-id;
        }
        dead-interval seconds;
        hello-interval seconds;
        interface-type type;
        metric metric;
        neighbor address <eligible>;
        passive;
        poll-interval seconds;
        priority number;
        retransmit-interval seconds;
        te-metric metric;
        transit-delay seconds;
    }
    label-switched-path name metric metric;
    nssa {
        area-range network/masklen <restrict>;
        default-lsa {
            default-metric metric;
            metric-type type;
            type-7;
        }
        (no-summaries | summaries);
    }
    peer-interface interface-name {
        disable;
        dead-interval seconds;
        hello-interval seconds;
        retransmit-interval seconds;
        transit-delay seconds;
    }
```

```
        stub <default-metric metric> < (no-summaries | summaries)>;
        virtual-link neighbor-id router-id transit-area area-id {
            authentication {
                md5 key-id;
                simple-password key-id;
            }
            dead-interval seconds;
            disable;
            hello-interval seconds;
            retransmit-interval seconds;
            transit-delay seconds;
        }
    }
} # End of [edit protocols ospf] hierarchy level
```

Pragmatic General Multicast (PGM)

```
pgm {
    traceoptions {
        file name <replace> <size size> <files number > <no-stamp>
                          <(world-readable | no-world-readable)>;
        flag flag <flag-modifier >;
    }
} # End of [edit protocols pgm] hierarchy level
```

Protocol Independent Multicast (PIM)

```
pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        disable;
        mode (dense | sparse | sparse-dense);
        priority number;
        version version;
    }
    rib-group group-name;
    rp {
        auto-rp (announce | discovery | mapping);
        bootstrap-export [ policy-names ];
        bootstrap-import [ policy-names ];
        bootstrap-priority number;
        embedded-rp {
            maximum-rps limit;
            group-ranges {
                destination-mask;
            }
        }
        local {
            family (inet | inet6) {
                disable;
                address address;
                anycast-pim {
                    rp-set {
                        address address [forward-msdp-sa];
                    }
                    local-address address;
                }
                group-ranges {
                    destination-mask;
                }
                hold-time seconds;
                priority number;
            }
        }
    static {
        address address {
            version version;
            group-ranges {
                destination-mask;
            }
            traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                        <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
        }
    }
} # End of [edit protocols pim] hierarchy level
```

Routing Information Protocol (RIP)

```
rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    graceful-restart {
        disable;
        restart-time seconds;
    }
    hold-down seconds;
    import [ policy-names ];
    message-size number;
    metric-in metric;
    receive receive-options;
    rib-group group-name;
    send send-options;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
                    <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
```

```
        group group-name {
            export [ policy-names ];
            metric-out metric;
            preference preference;
            neighbor neighbor-name {
                authentication-key password;
                authentication-type type;
                (check-zero | no-check-zero);
                import [ policy-names ];
                message-size number;
                metric-in metric;
                receive receive-options;
                send send-options;
            }
        }
    } # End of [edit protocols rip] hierarchy level
```

Router Discovery

```
    router-discovery {
        disable;
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                              <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
        }
        interface interface-name {
            min-advertisement-interval seconds;
            max-advertisement-interval seconds;
            lifetime seconds;
        }
        address address {
            (advertise | ignore);
            (broadcast | multicast);
            (priority number | ineligible);
        }
    } # End of [edit protocols router-discovery] hierarchy level
```

Session Announcement Protocol/Session Description Protocol (SAP/SDP)

```
    sap {
        disable;
        listen <address> <port port>;
    } # End of [edit protocols sap] hierarchy level
```

Virtual Router Redundancy Protocol (VRRP)

```
    vrrp {
        traceoptions {
            file {
                filename filename;
                files number;
                size size;
                (world-readable | no-world-readable);
            }
            flag flag
        }
    } # End of [edit protocols vrrp] hierarchy level
} # End of [edit protocols] hierarchy level
```

[edit routing-options] Hierarchy Level

```
routing-options {
    aggregate {
      defaults {
        aggregate-options;
      }
      route destination-prefix {
        policy policy-name;
        aggregate-options;
      }
    }
    autonomous-system autonomous-system <loops number>;
    confederation confederation-autonomous-system members autonomous-system;
    dynamic-tunnels tunnel-name {
      destination-networks prefix;
      source-address address;
      tunnel-type tunnel-type;
    }
    forwarding-table {
      export [ policy-names ];
      unicast-reverse-paths (active-paths | feasible-paths);
    }
    generate {
      defaults {
        generate-options;
      }
      route destination-prefix {
        policy policy-name;
        generate-options;
      }
    }
    graceful-restart {
      disable;
      path-selection-defer-time-limit time-limit;
    }
    interface-routes {
      family (inet | inet6) {
        export {
          lan;
          point-to-point;
        }
      }
      rib-group group-name;
    }
    martians {
      destination-prefix match-type <allow>;
    }
    maximum-routes route-limit <log-only | threshold value>;
    multicast {
      forwarding-cache {
        threshold suppress value <reuse value>;
      }
      scope scope-name {
        interface interface-name;
        prefix destination-prefix;
      }
    ssm-groups {
        address;
    }
    ssm-map ssm-map-name {
        policy policy-name;
        source addresses;‘
    }
}
options {
    syslog (level level | upto level);
}
resolution {
    rib routing-table-name {
        import [ policy-names ];
        resolution-ribs [ routing-table-names ];
    }
}
rib routing-table-name {
    aggregate {
        defaults {
            aggregate-options;
        }
        rib-group group-name;
        route destination-prefix {
            policy policy-name;
            aggregate-options;
        }
    }
    filter {
        input filter-name;
    }
    generate {
        defaults {
            generate-options;
        }
        route destination-prefix {
            policy policy-name;
            generate-options;
        }
    }
    martians {
        destination-prefix match-type <allow>;
    }
    static {
        defaults {
            static-options;
        }
        rib-group group-name;
        route destination-prefix {
            lsp-next-hop {
                metric metric;
                preference preference;
            }
```

```
                    next-hop;
                    p2mp-lsp-next-hop {
                        metric metric;
                        preference preference;
                    }
                    qualified-next-hop address {
                        metric metric;
                        preference preference;
                    }
                    static-options;
                }
            }
        }
        rib-groups {
            group-name {
                import-policy [ policy-names ];
                import-rib [ group-names ];
                export-rib [ group-names ];
            }
        }
        route-record;
        router-id address;
        static {
            defaults {
                static-options;
            }
            rib-group group-name;
            route destination-prefix {
                lsp-next-hop {
                    metric metric;
                    preference preference;
                }
                next-hop;
                qualified-next-hop address {
                    metric metric;
                    preference preference;
                }
                static-options;
            }
        }
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                <(world-readable | no-world-readable)>;
                flag flag <flag-modifier> <disable>;
        }
} # End of [edit routing-options] hierarchy level
```

[edit security] Hierarchy Level

```
security {
    certificate {
        cache-size bytes;
        cache-timeout-negative seconds;
        certification-authority ca-profile-name {
            ca-name certificate-authority-name;
            crl file-name;
            encoding (binary | pem);
            file certificate-filename;
            enrollment-url url-name;
            ldap-url url-name;
        }
        enrollment-retry number;
        local certificate-name;
        maximum-certificates maximum-number;
        path-length bytes;
    }
    ike {
        policy ike-peer-address {
            description policy-description;
            encoding (binary | pem);
            identity identity-name;
            local certificate-name;
            local-key-pair private-public-key-file;
            mode (aggressive | main);
            pre-shared-key (ascii-text key | hexadecimal key);
            proposals [ proposal-names ];
        }
        proposal ike-proposal-name {
            authentication-algorithm (md5 | sha1);
            authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
            dh-group (group1 | group2);
            encryption-algorithm (3des-cbc | des-cbc);
            lifetime-seconds seconds;
        }
    }
    ipsec {
        internal {
            security-association {
                manual {
                    direction (bidirectional | inbound | outbound) {
                        protocol esp;
                        spi spi-value;
                        encryption {
                            algorithm 3des-cbc;
                            key ascii-text ascii-text-string;
                        }
                    }
                }
            }
        }
```

```
      policy ipsec-policy-name {
        perfect-forward-secrecy {
          keys (group1 | group2);
        }
        proposals [ proposal-names];
      }
      proposal ipsec-proposal-name {
        authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
        encryption-algorithm (3des-cbc | des-cbc);
        lifetime-seconds seconds;
        protocol (ah | esp | bundle);
      }
      security-association name {
        dynamic {
          <security-association (32 | 64)>;
          ipsec-policy policy-name;
        }
        manual {
          direction (inbound | outbound | bi-directional) {
            authentication {
              algorithm (hmac-md5-96 | hmac-sha1-96);
              key (ascii-text key | hexadecimal key);
            }
            auxiliary-spi auxiliary-spi-value;
            encryption {
              algorithm (des-cbc | 3des-cbc);
              key (ascii-text key | hexadecimal key);
            }
            protocol (ah | esp | bundle);
            spi spi-value;
          }
        }
        mode (tunnel | transport);
        traceoptions {
          file <files number> <size size>;
          flag all;
          flag database;
          flag general;
          flag ike;
          flag parse;
          flag policy-manager;
          flag routing-socket;
          flag timer;
        }
      }
    }
} # End of [edit security] hierarchy level
```

[edit services] hierarchy level

```
services {
    adaptive-services-pics {
      traceoptions {
        flag flag;
      }
    }
    dynamic-flow-capture {
      capture-group client-name {
        content-destination identifier {
          address address;
          ttl hops;
        }
        control-source identifier {
          allowed-destinations [ destinations ];
          no-syslog;
          notification-targets [ address address port port-number ];
          service-port port-number;
          shared-key value;
          source-addresses [ addresses ];
        }
        input-packet-rate-threshold rate;
        interfaces interface-name;
        pic-memory-threshold percentage percentage;
      }
    }
    ids {
      rule rule-name {
        match-direction (input | output | input-output);
        term term-name {
          from {
            applications [ application-names ];
            application-sets [ set-names ];
            destination-address address;
            source-address address;
          }
          then {
            aggregation {
              destination-prefix prefix-value;
              source-prefix prefix-value;
            }
            (force-entry | ignore-entry);
            logging {
              syslog;
              threshold rate;
            }
            session-limit {
              by-destination {
                hold-time seconds;
                maximum number;
                packets number;
                rate number;
              }
```

```
                    by-pair {
                        maximum number;
                        packets number;
                        rate number;
                    }
                    by-source {
                        hold-time seconds;
                        maximum number;
                        packets number;
                        rate number;
                    }
                }
                syn-cookie {
                    mss value;
                    threshold rate;
                }
            }
        }
    }
    rule-set rule-set-name {
        [ rule rule-names ];
    }
}
ipsec-vpn {
    ike {
         proposal-name {
            authentication-algorithm (md5 | sha1);
            authentication method (dsa-signatures | pre-shared-keys | rsa-signatures);
            description description;
            dh-group (group1 | group2);
            encryption-algorithm (3des-cbc | des-cbc);
            lifetime-seconds seconds;
        }
        policy policy-name {
            description description;
            local-id {
                fqdn [ values ];
                ipv4_addr [ values ];
                key_id [ values ];
            }
            mode (aggressive | main);
            pre-shared-key (ascii-text key | hexadecimal key);
            proposals [ proposal-names ];
            remote-id {
                fqdn [ values ];
                ipv4_addr [ values ];
                key_id [ values ];
            }
        }
    }
}
```

```
nat {
    pool nat-pool-name {
        address (address | address-range low value high value | prefix);
        port (automatic | range low minimum-value high maximum-value);
    }
    rule rule-name {
        match-direction (input | output);
        term term-name {
            from {
                applications [ application-names ];
                application-sets [ set-names ];
                destination-sets address;
                ipsec-inside-interface interface-name;
                source-address (address | prefix);
            }
            then {
                translated {
                    destination-pool nat-pool-name;
                    source-pool nat-pool-name;
                    translation-type (destination type | source type);
                }
                syslog;
            }
        }
    }
    rule-set rule-set-name {
        [ rule rule-names ];
    }
}
rpm {
    probe owner {
        test test-name {
                data-fill data;
                data-size size;
                destination-port port;
                dscp-code-point DSCP bits;
                history-size size;
                probe-count count;
                probe-interval seconds;
                probe-type type;
                routing-instance routing-instance-name;
                source-address address;
                target-url (url | address);
                test-interval interval;
                thresholds thresholds;
                traps traps;
        }
    }
    probe-server {
        tcp port;
        udp port;
    }
    probe-limit limit;
    }
}
```

```
    service-set service-set-name {
        ([ ids-rules rule-names ] | ids-rule-sets rule-set-name);
        ([ ipsec-vpn-rules rule-names ] | ipsec-vpn-rule-sets rule-set-name);
        ([ nat-rules rule-names ] | nat-rule-sets rule-set-name);
        ([ stateful-firewall-rules rule-names ] | stateful-firewall-rule-sets rule-set-name);
        interface-service {
            service-interface interface-name;
        }
        ipsec-vpn-options {
            ike-access-profile profile-name;
            local-gateway address;
        }
        max-flows number;
        next-hop-service {
            inside-service-interface name.number;
            outside-service-interface name.number;
        }
        syslog {
            host hostname {
                services priority-level;
                facility-override facility-name;
                log-prefix prefix-number;
            }
        }
    }
    stateful-firewall {
        rule rule-name {
            match-direction (input | output | input-output);
            term term-name {
                from {
                    applications [ application-names ];
                    application-sets [ set-names ];
                    destination-address address;
                    source-address address;
                }
                then {
                    (accept | discard | reject);
                    allow-ip-option [ values ];
                    syslog;
                }
            }
        }
        rule-set rule-set-name {
            [ rule rule-names ];
        }
    }
} # End of [edit services] hierarchy level
```

[edit snmp] Hierarchy Level

```
snmp {
    community community-name {
        authorization authorization;
        clients {
            address restrict;
        }
        view view-name;
    }
    contact contact;
    description description;
    engine-id {
        (local engine-id | use-mac-address | use-default-ip-address);
    }
    interface [ interface-name ];
    filter-duplicates;
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            interval seconds;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            variable oid-variable;
        }
        event index {
            community community-name;
            description description;
            type type;
        }
    }
    traceoptions {
        file size size files number;
        flag flag;
    }
    trap-group group-name {
        categories [ categories ];
        destination-port <port-number>;
        targets {
            address;
        }
        version (all | v1 | v2);
    }
    trap-options {
        agent-address outgoing-interface;
        source-address address;
    }
```

```
v3 {
    notify name {
        tag tag-name;
        type (trap | inform);
    }
    notify-filter name {
        oid oid (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        inform-retry-count number;
        inform-timeout seconds:
        port port-number;
        tag-list tag-list;
        target-parameters target-parameters-name;
    }
    target-parameters target-parameters-name {
        notify-filter name;
        parameters {
            message-processing-model (v1 | v2c | V3);
            security-model ( usm | v1 | v2c);
            security level (authentication | none | privacy);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password password;
                }
                authentication-sha {
                    authentication-password password;
                }
                authentication-none;
                    privacy-password password;
                }
                privacy-aes128 {
                    privacy-password password;
                }
                privacy-des {
                    privacy-password password;
                }
                privacy-none;
            }
        }
    }
```

```
      vacm {
        access {
          group group-name {
            default-context-prefix {
              security-model (any | usm | v1 | v2c) {
                security-level (authentication | none | privacy) {
                  notify-view notify-view;
                  read-view read-view;
                  write-view write-view;
                }
              }
            }
          }
        }
        security-to-group {
          security-model (usm | v1 | v2c) {
            security-name security-name {
              group group-name;
            }
          }
        }
      }
    }
    view view-name; {
      oid object-identifier (include | exclude);
    }
} # End of [edit snmp] hierarchy level
```

[edit system] Hierarchy Level

```
system {
    accounting {
      events [ login change-log interactive-commands ];
      destination {
        radius {
          server {
            server-address {
              accounting-port port-number;
              retry number;
              secret password;
              source-address address;
              timeout seconds;
              }
            }
          }
          tacplus {
            server {
              server-address {
                port port-number;
                server password;
                single-connection;
                source-address source-address;
                timeout seconds;
              }
            }
          }
        }
    }
}
archival {
    configuration {
        archive-sites {
            ftp://<username>:<password>@<host>:<port>/<url-path>;
        }
        transfer-interval interval;
        transfer-on-commit;
    }
}
arp {
    passive-learning;
    aging-timer minutes;
}
authentication-order [ authentication-methods ];
autoinstallation {
    interfaces {
        interface-name {
            bootp;
            rarp;
            slarp;
        }
    }
    configuration-servers {
        url;
    }
}
backup-router address <destination destination-address>;
commit synchronize;
(compress-configuration-files | no-compression-configuration-files);
default-address-selection;
diag-port-authentication (encrypted-password "password" | plain-text-password);
domain-name domain-name;
domain-search [ domain-list ];
dump-device (compact-flash | remove-compact | usb);
host-name hostname;
internet-options address <destination destination-address>;
internet-options {
    path-mtu-discovery;
    source-quench;
    source-port upper-limit <upper-limit>;
}
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
        longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
```

```
login {
   announcement text;
   class class-name {
      allow-commands "regular-expression" ;
      allow-configuration "regular-expression" ;
      deny-commands "regular-expression";
      deny-configuration "regular-expression";
      idle-timeout minutes;
      login-alarms;
      login-tips:
      permissions [ permissions ];
   }
   message text;
   passwords {
      change-type (set-transitions | character-set);
      format (md5 | sha1 | des);
      maximum-length length;
      minimum-changes number;
      minimum-length length;
   }
   user username {
      full-name complete-name;
      uid uid-value;
      class class-name;
      authentication {
         (encrypted-password "password" | plain-text-password);
         ssh-rsa "public-key";
         ssh-dsa "public-key";
      }
   }
}
max-configurations-on-flash number;
mirror-flash-on-disk;
name-server {
   address;
}
no-redirects;
ntp {
   authentication-key key-number type type value password;
   boot-server address;
   broadcast <address> <key key-number> <version value> <ttl value>;
   broadcast-client;
   multicast-client <address>;
   peer address <key key-number> <version value> <prefer>;
   server address <key key-number> <version value> <prefer>;
   trusted-key [ key-numbers ];
}
pic-console-authentication {
   encrypted-password encrypted-password;
   plain-text-password;
}
ports {
   auxiliary {
      type terminal-type;
   }
```

```
    console {
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
processes {
    adaptive-services (enable | disable) failover failover-option;
    alarm-control (enable | disable) failover failover-option;
    chassis-control (enable | disable) failover failover-option;
    class-of-service (enable | disable) failover failover-option;
    craft-control (enable | disable) failover failover-option;
    dhcp (enable | disable) failover failover-option;
    disk-monitoring (enable | disable) failover failover-option;
    ecc-error-logging (enable | disable) failover failover-option;
    firewall (enable | disable) failover failover-option;
    inet-process (enable | disable) failover failover-option;
    interface-control (enable | disable) failover failover-option;
    kernel-replication (enable | disable) failover failover-option;
    l2tp-service (enable | disable) failover failover-option;
    link-management (enable | disable) failover failover-option;
    mib-process (enable | disable) failover failover-option;
    network-access-service (enable | disable) failover failover-option;
    ntp (enable | disable) failover failover-option;
    pgm (enable | disable) failover failover-option;
    pic-services-logging (enable | disable) failover failover-option;
    pppoe (enable | disable) failover failover-option;
    redundancy-device (enable | disable) failover failover-option;
    remote-operations (enable | disable) failover failover-option;
    routing (enable | disable) failover failover-option;
    sampling (enable | disable) failover failover-option;
    service-deployment (enable | disable) failover failover-option;
    snmp (enable | disable) failover failover-option;
    timeout seconds;
    usb-control (enable | disable) failover failover-option;
    watchdog (enable | disable) failover failover-option;
    web-management (enable | disable) failover failover-option;
}
radius-server server-address {
    port number;
    retry number;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files files;
```

```
scripts {
  commit {
    (allow-transients | no-allow-transients);
    file filename.xsl {
      optional;
      refresh;
      refresh-from url;
      source url;
    }
    traceoptions {
      file filename <files number> <size size>;
      flag flag;
    }
  }
}
services {
  dhcp {
    boot-file filename;
    boot-server (address | hostname);
    domain-name domain-name;
    domain-search [domain-list];
    default-lease-time seconds;
    maximum-lease-time seconds;
    name-server {
      address;
    }
    option {
      [ (id-number option-type option-value) | (id-number array option-type
    option-values) ];
    }
    pool (address /prefix-length) {
      address-range {
        low address;
        high address;
      }
      exclude-address {
        address;
      }
    }
    router {
      address;
    }
    static-binding MAC-address {
      fixed-address {
        address;
      }
      host hostname;
      client-identifier (ascii client-id | hexadecimal client-id);
    }
    server-identifier address;
    wins-server {
      address;
    }
  }
```

```
    finger {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    ftp {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    service-deployment {
        servers server-address {
            port port-number;
        }
        source-address source-address;
    }
    ssh {
        <connection-limit limit>;
        protocol-version [ versions ];
        <rate-limit limit>;
        root-login (allow | deny | deny-password);
    }
    telnet {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    web-management {
        http {
            interfaces [ interface-names ];
            port port;
        }
        https {
            interfaces [ interface-names ];
            local-certificate name;
            port port;
        }
    }
    xnm-clear-text {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    xnm-ssl {
        <connection-limit limit>;
        local-certificate name;
        <rate-limit limit>;
    }
}
static-host-mapping {
    host-name {
        alias [ alias ];
        inet [ address ];
        sysid system-identifier ;
    }
}
```

```
        syslog {
            archive {
                files number;
                size size;
                (world-readable | no-world-readable);
            }
            console {
                facility severity;
            }
            file filename {
                facility severity;
                explicit-priority;
                archive {
                    files number;
                    size size;
                    (world-readable | no-world-readable);
                }
            }
            host (hostname | other-routing-engine) {
                facility severity;
                explicit-priority;
                facility-override facility;
                log-prefix string;
            }
            source-address source-address;
            time-format (year | millisecond | year millisecond);
            user (username | *) {
                facility level;
            }
        }
        tacplus-options service-name service-name;
        tacplus-server server-address {
            secret password;
            single-connection;
            source-address source-address;
            timeout seconds;
        }
        time-zone (GMThour-offset | time-zone);
} # End of [edit system] hierarchy level
```

Part 2

Command-Line Interface

Chapter 5

CLI Overview

The command-line interface (CLI) is the interface to the software that you use whenever you access the router—whether from the console or through a remote network connection. The CLI, which automatically starts after the router finishes booting, provides commands that you use to perform various tasks, including configuring the JUNOS software and monitoring and troubleshooting the software, network connectivity, and the router hardware.

The CLI is a straightforward command interface. You type commands on a single line, and the commands are executed when you press the Enter key. The CLI provides command help and command completion, and it also provides Emacs-style keyboard sequences that allow you to move around on a command line and scroll through a buffer that contains recently executed commands.

The CLI is indicated by the presence of the > or # prompt, which is preceded by a string that defaults to the name of the user and the name of the router. For example:

```
user@host>

user@host#
```

For information about customizing your CLI session, see "Configuring the Router with the CLI" on page 217. For information on CLI differences on a TX Matrix platform and single routing platform, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

This chapter discusses the following topics:

- CLI Modes on page 176
- CLI Command Hierarchy on page 176
- Using the J-Web Graphical User Interface as an Alternative to the CLI on page 177
- Commands and Configuration Statements for JUNOS-FIPS on page 178
- Reserved Names and Values on page 179

CLI Modes

The CLI has two modes: operational and configuration. In operational mode, you monitor and troubleshoot the software, network connectivity, and the router by entering commands. For more information about operational mode, see “CLI Operational Mode” on page 181.

When in configuration mode, you configure the JUNOS software by creating a hierarchy of configuration statements. You can do this by using the CLI or by creating a text (ASCII) file that contains the statement hierarchy. (The statement hierarchy is identical in both the CLI and text configuration file.) You can configure all properties of the JUNOS software, including interfaces, general routing information, routing protocols, and user access, as well as several system hardware properties. When you have finished entering the configuration statements, you commit them, which activates the configuration on the router. For more information about configuration mode, see “Configuring the Router with the CLI” on page 217.

CLI Command Hierarchy

The CLI commands are organized in a hierarchical fashion, with commands that perform a similar function grouped together under the same level. For example, all commands that display information about the system and the 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 2 illustrates a portion of the show command hierarchy.

Figure 2: CLI Command Hierarchy Example

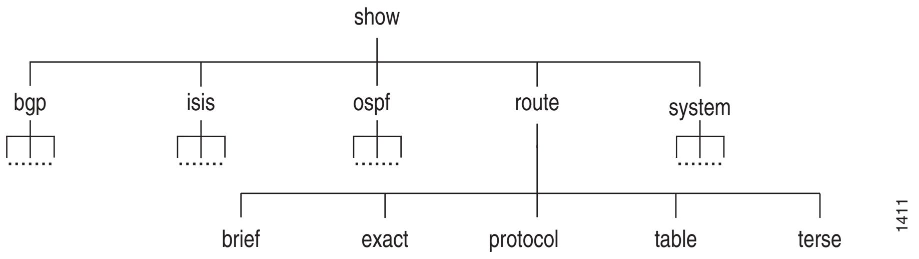

To execute a command, you enter the full command name, starting at the top level of the hierarchy. For example, to display a brief view of the routes in the router table, use the command show route brief.

The hierarchical organization results in commands that have a regular syntax and provides several 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. As examples, 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. This means that if you already are familiar with the JUNOS software or with other routing software, you can use many of the CLI commands without referring to the documentation.

- Command completion—Command completion for command names (keywords) and for 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.

Using the J-Web Graphical User Interface as an Alternative to the CLI

As an alternative to entering CLI commands, JUNOS software also supports a J-Web graphical user interface (GUI). The J-Web user interface allows you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser with Hypertext Transfer Protocol (HTTP) or HTTP over Secure Sockets Layer (HTTPS) enabled.

The J-Web user interface comes standard on J-series Services Routers. It is provided as an optional, licensed software package on M-series and T-series routers. For more information on installing the optional jweb package, see “Upgrading Software Packages” on page 341.

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 a certificate on the router and enabled HTTPS.

NOTE: If the router is running the worldwide version of the JUNOS software and you are using an Internet Explorer Web browser, you must disable the Use SSL 3.0 option in the Web browser to access the router.

2. After http:// or https:// in your Web browser, type the hostname or IP address of the router.

   The J-Web login page appears.

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

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

NOTE: J-Web menus and functions supported on J-series Services Routers differ slightly from J-Web menus and functions supported on M-series and T-series routers. For details on how to use J-Web menus and functions, see the *J-Web Interface User Guide*.

J-Web Sessions

You establish a J-Web session with the router through an HTTP- or HTTPS-enabled Web browser. To use HTTPS, you must have installed a certificate on the router and enabled HTTPS.

When you attempt to log in through the J-Web interface, the router authenticates your username with the same methods used for telnet and SSH. If the router does not detect any activity through the J-Web interface for 24 hours, the session times out and is terminated. You must log in again to begin a new session.

Commands and Configuration Statements for JUNOS-FIPS

The JUNOS-FIPS software environment requires the installation of FIPS software by a Crypto Officer. In JUNOS-FIPS, some JUNOS commands and statements have restrictions and some additional configuration statements are available.

- JUNOS-FIPS defines a restricted set of user roles. Unlike JUNOS, which allows a wide range of capabilities to users, FIPS 140-2 defines three especially important types of users: Crypto Officer, User, and Maintenance.
- JUNOS-FIPS disables many of the usual JUNOS protocols and services. Services that cannot be configured include telnet, rlogin, rsh, DHCP, FTP, finger, JUNOScript clear text, and TFTP.
- Only secure sockets layer (SSL) version 2 or TLS can be used for remote access. TLS is secure sockets layer (SSL) version 3 with FIPS-restricted encryption. The remote access service uses only the following FIPS-approved encryption:
  - AES128-SHA
  - AES256-SHA
  - DES-CBC3-SHA (not recommended)
- JUNOS-FIPS has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If JUNOS-FIPS is installed, you cannot configure passwords unless they meet this standard.
- In configurations that include multiple Routing Engines, IPSec is required for communications between Routing Engines.

For more information, see the *JUNOS-FIPS Configuration Guide*

Reserved Names and Values

The following list of names and values are reserved for use by Juniper Networks:

- Identifiers starting with junos-
- Identifiers starting with __juniper-
- Logical interface unit numbers greater than 16383

Chapter 6

CLI Operational Mode

When you log in to the router and the command-line interface (CLI) starts, you are at the top level of operational mode. Operational mode is indicated by the presence of the > prompt, which is preceded by a string that defaults to the name of the user and the name of the router. For example:

```
user@host>
```

At this level, there are a number of broad groups of CLI commands:

- Commands for controlling the CLI environment—The commands in the set hierarchy configure the CLI display screen. For information about these commands, see "Controlling the CLI Environment" on page 213.
- Commands for monitoring and troubleshooting—The following commands let you display information and statistics about the software and test network connectivity. Using these commands is discussed in the *JUNOS Interfaces Command Reference*.
  - clear—Clear statistics and protocol database information.
  - mtrace—Trace mtrace packets from source to receiver.
  - monitor—Perform real-time debugging of various software components, including the routing protocols and interfaces.
  - ping—Determine the reachability of a remote network host.
  - show—Display the current configuration and information about interfaces, routing protocols, routing tables, routing policy filters, system alarms, and the chassis.
  - test—Test the configuration and application of policy filters and autonomous system (AS) path regular expressions.
  - traceroute—Trace the route to a remote network host.
- Commands for connecting to other network systems—The ssh command opens secure shell connections, and the telnet command opens telnet sessions to other hosts on the network. For information about these commands, see the *JUNOS System Basics and Services Command Reference*.

- Commands for copying files—The copy command copies files from one location on the router to another, from the router to a remote system, or from a remote system to the router. For information about these commands, see the *JUNOS System Basics and Services Command Reference*.
- Commands for restarting software processes—The commands in the restart hierarchy restart the various JUNOS software processes, including the routing protocol, interface, and Simple Network Management Protocol (SNMP). For information about these commands, see the *JUNOS System Basics and Services Command Reference*.
- A command—request—for performing system-level operations, including stopping and rebooting the router and loading JUNOS software images. For information about this command, see the *JUNOS System Basics and Services Command Reference*.
- A command—start—to exit the CLI and start a UNIX shell. For information about this command, see the *JUNOS System Basics and Services Command Reference*.
- A command—configure—for entering configuration mode, which provides a series of commands that configure the JUNOS software, including the routing protocols, interfaces, network management, and user access. For information about the CLI configuration commands, see "Configuring the Router with the CLI" on page 217.
- A command—quit—to exit the CLI. For information about this command, see the *JUNOS System Basics and Services Command Reference*.
- For more information about the CLI operational mode commands, see the *JUNOS Interfaces Command Reference* and the *JUNOS System Basics and Services Command Reference*.

This chapter discusses the following topics:

Using the CLI

This section describes how to use the JUNOS software CLI. It discusses the following topics:

- Getting Help About Commands on page 183
- Getting Help Based on a String in a Statement Name on page 185
- Displaying Tips About CLI Commands on page 185
- Using CLI Complete Commands on page 185
- CLI Messages on page 186
- Moving Around and Editing the Command Line on page 187
- How Output Appears on the Screen on page 188

Getting Help About Commands

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 **?**. You do not need to press Enter after typing the question mark.

- If you type the question mark at the command-line prompt, the CLI lists the available commands and options. For example, 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
  mtrace         Trace mtrace packets from source to receiver.
  monitor        Real-time debugging
  ping           Ping a remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart a software process
  set            Set CLI properties, date, time, craft display text
  show           Show information about the system
  ssh            Open a secure shell to another host
  start          Start a software process
  telnet         Telnet to another host
  test           Diagnostic debugging commands
  traceroute     Trace the route to a remote host
user@host>
```

- If you type the 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> clear ?
Possible completions:
  arp           Clear address-resolution information
  bgp           Clear BGP information
  chassis       Clear chassis information
  firewall      Clear firewall counters
  igmp          Clear IGMP information
  interfaces    Clear interface information
  ilmi          Clear ILMI statistics information
  isis          Clear IS-IS information
  ldp           Clear LDP information
  log           Clear contents of a log file
  mpls          Clear MPLS information
  msdp          Clear MSDP information
  multicast     Clear Multicast information
  ospf          Clear OSPF information
  pim           Clear PIM information
  rip           Clear RIP information
  route         Clear routing table information
  rsvp          Clear RSVP information
  snmp          Clear SNMP information
  system        Clear system status
  vrrp          Clear VRRP statistics information
user@host> clear
```

- If you type the 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 c, type the following:

```
user@host> c?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
user@host>c
```

- For introductory information on using the question mark or the help command, you can also type **help** and press Enter:

```
help
```

Getting Help Based on a String in a Statement Name

In operational mode, you can use the help command to display help based on a text string contained in a statement name. This command displays help for statements at the current hierarchy level and below:

```
help apropos string
```

string is a text string about which you want to get help. This string is used to match statement names as well as the help strings that are displayed for the statements. If the string contains spaces, enclose it in quotation marks (" "). You also can specify a regular expression for the string, using standard UNIX-style regular expression syntax.

You can display help based on a text string contained in a statement name using the help topic and help reference commands:

```
help topic string
help reference string
```

The help topic command displays usage guidelines for the statement, while the help reference command displays summary information about the statement.

You can display help based on a system log tag using the help syslog command:

```
help syslog syslog-tag
```

The help syslog command displays the contents of a syslog message.

Displaying Tips About CLI Commands

To get tips about CLI commands, issue the help tip cli command. For example:

```
user@host> help tip cli <number>
JUNOS tip:
Use 'request system software validate' to validate the incoming software
against the current configuration without impacting the running system.
```

< number > associates a tip with a number.

You can enable or disable the tip command automatically when a user logs in. See “Configuring Tips” on page 412.

Using CLI Complete Commands

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 the tab key or the spacebar. If the partially typed letters begin a string that uniquely identifies a command, the complete command name appears. Otherwise, a beep indicates that you have entered an ambiguous command, and the possible completions are displayed.

Command completion also applies to other strings, such as filenames and usernames. To display all possible values, type a partial string followed immediately by a question mark. However, to complete these strings, press the tab key; pressing the space bar does not work.

Examples: Using CLI Command Completion

Issue the show interfaces command:

```
user@host> sh<Space>ow i<Space>
'i' is ambiguous.
Possible completions:
    igmp      Show information about IGMP
    interfaceShow interface information
    isis      Show information about IS-IS
user@host> show in<Space>terfaces <Enter>
Physical interface: at-0/1/0, Enabled, Physical link is Up
  Interface index: 11, SNMP ifIndex: 65
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode
  Speed: OC12, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : 0x01
...
user@host>
```

Display a list of all log files whose names start with the string "messages," and then display the contents of one of the files:

```
user@myhost> show log mes?
Possible completions:
  <filename>         Log file to display
  messages           Size: 1417052, Last changed: Mar  3 00:33
  messages.0.gz      Size: 145575, Last changed: Mar  3 00:00
  messages.1.gz      Size: 134253, Last changed: Mar  2 23:00
  messages.10.gz     Size: 137022, Last changed: Mar  2 14:00
  messages.2.gr      Size: 137112, Last changed: Mar  2 22:00
  messages.3.gz      Size: 121633, Last changed: Mar  2 21:00
  messages.4.gz      Size: 135715, Last changed: Mar  2 20:00
  messages.5.gz      Size: 137504, Last changed: Mar  2 19:00
  messages.6.gz      Size: 134591, Last changed: Mar  2 18:00
  messages.7.gz      Size: 132670, Last changed: Mar  2 17:00
  messages.8.gz      Size: 136596, Last changed: Mar  2 16:00
  messages.9.gz      Size: 136210, Last changed: Mar  2 15:00
user@myhost> show log mes<Tab>sages.4<Tab>.gz<Enter>
Jan 15 21:00:00 myhost newsyslog[1381]: logfile turned over
...
```

CLI Messages

Messages appear when you enter and exit from configuration mode, when you commit a configuration, and when you type a string or value that is not valid.

When you commit a configuration, the JUNOS software checks the configuration you are committing. If there are no problems, a message indicates that the configuration was accepted. If there are problems, a message indicates where the errors are.

In the top-level CLI commands and in configuration mode, if you type an invalid string—for example, the name of a command or statement that does not exist—you see the message “syntax error” or “unknown command”. A caret (^) indicates where the error is. Examples:

```
user@host> clear route
                       ^
syntax error, expecting <command>.

[edit]
user@host# telnet
            ^
unknown command.
```

When the number of choices is limited, a message might display the commands you can enter to correct the syntax error. For example:

```
[edit]
user@host# load myconfig-file<Enter>
                ^
syntax error, expecting ‘merge’, ‘override’, or ‘replace’.
```

Moving Around and Editing the Command Line

In the CLI, you can use keyboard sequences to move around on a command line and edit the command line. You can also use keyboard sequences to scroll through a list of recently executed commands. Table 7 lists some of the CLI keyboard sequences. They are the same as those used in Emacs.

Table 7: CLI Keyboard Sequences

Category	Action	Keyboard Sequence
Move the Cursor	Move the cursor back one character.	Ctrl-b
	Move the cursor back one word.	Esc-b or Alt-b
	Move the cursor forward one character.	Ctrl-f
	Move the cursor forward one word.	Esc-f or Alt-f
	Move the cursor to the beginning of the command line.	Ctrl-a
	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, Esc-Backspace, or Alt-Backspace
	Delete the word after the cursor.	Esc-d or Alt-d
Insert Recently Deleted Text	Insert the most recently deleted text at the cursor.	Ctrl-y

Category	Action	Keyboard Sequence
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-/
Display Previous Command Words	Scroll backward through the list of recently entered words in a command line.	Esc-. or Alt-.
Repeat Keyboard Sequences	Specify the number of times to execute a keyboard sequence. *number* can be from 1 through 9.	Esc-*number sequence* or Alt-*number sequence*

How Output Appears on the Screen

When you issue commands in operational mode, or when you issue the show command in configuration mode, the output appears on the screen. You can also filter the output of commands, either to perform simple commands on the output or to place the output into a file.

This section discusses the following topics:

- Displaying Output One Screen at a Time on page 188
- Filtering Command Output on page 190

Displaying Output One Screen at a Time

If the output is longer than the screen length, it appears one screen at a time by means of a UNIX more-type interface. The prompt —More— indicates that more output is available. The output buffer for the prompt is restricted to 32 megabytes (MB). Any new data that exceeds the buffer limit replaces the oldest data in the memory buffer. When the buffer limit is exceeded, attempts to search backward or navigate to the beginning of the output generate a warning indicating that the output is truncated. Because of the buffer size restriction, use of the Scroll Up and Search functions might be limited.

Table 8 lists the keyboard sequences you can use at the —More— prompt. As soon as the CLI knows how long the output is (usually by the second screen), it displays the percentage of the command output above the prompt.

Table 8: —More— Prompt Keyboard Sequences

Category	Action	Keyboard Sequence
Get Help	Display information about the keyboard sequences you can display at the —More— prompt.	h
Scroll Down	Scroll down one line.	Enter, Return, k, Ctrl-m, Ctrl-n, or down arrow
	Scroll down one-half screen.	Tab, d, Ctrl-d, or Ctrl-x
	Scroll down one whole screen.	Space or Ctrl-f
	Scroll down to the bottom of the output.	Ctrl-e or G
	Display the output all at once instead of one screen at a time. (Same as specifying the | no-more command.)	N
Scroll Up	Display the previous line of output.	j, Ctrl-h, Ctrl-p, or up arrow
	Scroll up one-half screen.	u or Ctrl-u
	Scroll up one whole screen.	b or Ctrl-b
	Scroll up to the top of the output.	Ctrl-a or g
Search	Search forward for a string.	/ *string*
	Search backward for a string.	?*string*
	Repeat the previous search for a string.	n
	Search for a text string. You are prompted for the string to match. (Same as specifying the | match *string* command.)	m or M
	Search, ignoring a text string. You are prompted for the string to not match. (Same as specifying the | except *string* command.)	e or E
Interrupt or End Output, Redraw the Output, and Save the Output to a File	Interrupt the display of output.	Ctrl-C, q, Q, or Ctrl-k
	Do not redisplay the CLI prompt immediately after displaying the output, but remain at the —More— prompt. (Same as specifying the | hold command.)	H
	Clear any match conditions and display the complete output.	c or C
	Redraw the output on the screen.	Ctrl-l
	Save the command output to a file. You are prompted for a filename. (Same as specifying the | save *filename* command.)	s or S

Filtering Command Output

For operational and configuration commands that display output, such as the show commands, you can filter the output. 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:

```
user@host> show configuration ?
Possible completions:
    <[Enter]>Execute this command
    |           Pipe through a command
user@host> show configuration | ?
Possible completions:
    count     Count occurrences
    except    Show only text that does not match a pattern
    find      Search for the first occurrence of a pattern
    hold      Hold text without exiting the —More– prompt
    last      Show the specified number of lines from the end of the output
    match     Show only text that matches a pattern
    no-more Don't paginate output
    resolve   Resolve IP addresses
    save      Save output text to a file
    trim      Trim specified number of columns from the start line
```

In configuration mode, two additional filters appear, display and compare:

```
[edit]
user@host # show | ?
Possible completions:
    compareCompare configuration changes with a prior version
    count     Count occurrences
    display   Display additional configuration information
    except    Show only text that does not match a pattern
    find      Search for the first occurrence of a pattern
    hold      Hold text without exiting the —More– prompt
    last      Show specified number of lines from the end of the output
    match     Show only text that matches a pattern
    no-more Don't paginate output
    resolve   Resolve IP addresses
    save      Save output text to a file
    trim      Trim specified number of columns from the start line
```

The following filtering operations are available:

- Placing Command Output in a File on page 191
- Searching for a String in the Output on page 191
- Comparing Configuration Changes with a Prior Version on page 194
- Counting the Number of Lines in the Output on page 196
- Displaying All Output at Once on page 196
- Displaying the Lines You Want to View from the End of the Output on page 196

Placing Command Output in a File

When command output is very long, when you need to store or analyze the output, or when you need to send the output in e-mail or by FTP, you can place the output into a file. Doing this is useful when the output scrolls off the screen, making it difficult to cut the output from a window and paste it into another.

To save the output to a file, specify the save command after the pipe:

```
user@host> command | save filename
```

By default, the file is placed in your home directory on the router. For information about how you can specify the name of the file, see "Specifying Filenames and URLs" on page 360.

This example stores the output of the request support information command in a file:

```
user@host> request support information | save filename
Wrote 1143 lines of output to 'filename'
user@host>
```

Searching for a String in the Output

You can search for text matching a regular expression by filtering output. You can make a regular expression match everything except a regular expression, or find the first occurrence of text matching a regular expression. Searches are not case-sensitive.

To match a regular expression, specify the match command after the pipe:

```
user@host> command | match regular-expression
```

To ignore text that matches a regular expression, specify the except command after the pipe:

```
user@host> command | except regular-expression
```

If the *regular-expression* contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

You use extended regular expressions to specify what text in the output to match. Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2. Table 9 on page 192 lists common regular expression operators.

Table 9: Common Regular Expression Operators in Operational Mode Commands

Operator	Match...
|	One of the two terms on either side of the pipe.
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces$" means that the user can issue the show interfaces command but cannot issue show interfaces detail or show interfaces extensive.
[ ]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
( )	A group of commands, indicating an expression to be evaluated; the result is then evaluated as part of the overall expression.

For example, if a command produces the following output:

```
one two
two two
three two one
four
```

The match two command displays:

```
one two
two two
three two one
```

The except one command displays:

```
two two
four
```

List all the Asynchronous Transfer Mode (ATM) interfaces in the configuration:

```
user@host> show configuration | match at-
    at-2/1/0 {
    at-2/1/1 {
    at-2/2/0 {
    at-5/2/0 {
    at-5/3/0 {
```

Display a skeleton of your router configuration:

```
[edit]
user@host # show | match {
system {
    root-authentication {
    name-server {
    login {
      class super-user {
      user juniper {
        authentication {
    services {
    syslog {
      file messages {
    processes {
chassis {
    alarm {
      sonet {
    images {
      scb {
      fpc {
interfaces {
    at-2/1/1 {
      atm-options {
      unit 0 {
    at-2/2/0 {
    ...
snmp {
    community public {
      clients {
routing-options {
    static {
      route 0.0.0.0/0 {
      route 192.168.0.0/16 {
      route 208.197.169.0/24 {
protocols {
    rsvp {
      interface so-5/1/0 {
    mpls {
      interface so-5/1/0 {
    bgp {
      group internal {
    ospf {
      area 0.0.0.0 {
        interface so-5/1/0 {
```

List all users who are logged in to the router except for the user “root”:

```
user@host> show system users | except root
 8:28PM  up 1 day, 13:59, 2 users, load averages: 0.01, 0.01, 0.00
USER     TTY FROM              LOGIN@  IDLE WHAT
sheep    p0  baa.juniper.net   7:25PM     - cli
```

Save the configuration, except for encrypted passwords, to a file:

```
user@host> show configuration | except SECRET-DATA | save my.output.file
```

Display the output, starting not at the beginning but rather at the first occurrence of text matching a regular expression, using the find command after the pipe:

```
user@host> command | find regular-expression
```

If the regular expression contains spaces, operators, or wildcard characters, enclose the expression in quotation marks.

List the routes in the routing table starting at 208.197.169.0:

```
user@host> show route | find 208.197.169.0
208.197.169.0/24   *[Static/5] 1d 13:22:11
                    > to 192.168.4.254 via so-3/0/0.0
224.0.0.5/32       *[OSPF/10] 1d 13:22:12, metric 1

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.4015.00/160
                   *[Direct/0] 1d 13:22:12
                    > via lo0.0
```

Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the compare command to display the configuration. The compare command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the compare command after the pipe:

```
[edit]
user@host# show | compare [filename | rollback n]
```

filename is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements. For information about how to save a configuration to a file, see “Saving a Configuration to a File” on page 263. For information about formatting the hierarchy of statements, see “Configuration Statement Hierarchy” on page 219.

n is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 9. If you do not specify arguments, the candidate configuration is compared against the active configuration file (/config/juniper.conf).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (–).
- Statements that are unchanged are prefixed with a single blank space ( ).

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the [edit protocols bgp] hierarchy level:

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

[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 60;
    advertise-inactive;
    allow 1.1.1.1/32;
}
group fred {
    type external;
    peer-as 33333;
    allow 2.2.2.2/32;
}
group test-peers {
    type external;
    allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-   hold-time 60;
+   hold-time 90;
-   advertise-inactive;
[edit protocols bgp group fred]
+   advertise-inactive;
[edit protocols bgp]
-group test-peers {
-   type external;
-   allow 3.3.3.3/32;
-}
```

```
[edit protocols bgp]
user@host# show
group my-group {
      type internal;
      hold-time 90;
      allow 1.1.1.1/32;
}
group fred {
      type external;
      advertise-inactive;
      peer-as 3333;
      allow 2.2.2.2/32;
}
```

Counting the Number of Lines in the Output

To count the number of lines in the output, specify the count command after the pipe:

```
user@host> command | count
```

For example:

```
user@host> show configuration | count
Count: 269 lines
user@host> show route | count
Count: 67 lines
```

Displaying All Output at Once

To display the output all at once instead of one screen at a time, specify the no-more command after the pipe. This command is equivalent to the set cli screen-length 0 command, but affects the output of the one command only.

```
user@host> command | no-more
```

Displaying the Lines You Want to View from the End of the Output

By default, the last command displays the lines from the end of the output. You can also specify the number of lines to display from the end of the output. If the number of specified lines is less than a screen full, the command will display a screen full of output. This feature is most useful for viewing log files where the end of the files contain the most recent entries.

To view the most recent entries from the end of the output, use the | last command after the pipe:

```
user@host> show command | last lines
```

lines specifies the number of lines to display from the end of output.

Retaining the Output After the Last Screen

When you view output one screen at a time, you typically return to the CLI prompt after viewing the last screen.

To not return immediately to the CLI prompt, use the hold command after the pipe. This feature is useful, for example, when you want to scroll or search through the output.

```
user@host> command | hold
```

Displaying Additional Information About the Configuration

In configuration mode only, to display additional information about the configuration, use the display detail command after the pipe in conjunction with a show command. The additional information includes the help string that explains each configuration statement and the permission bits required to add and modify the configuration statement.

```
user@host# show <hierarchy-level> | display detail
```

For example:

```
[edit]
user@host# show | display detail
    ##
    ## version: Software version information
    ## require: system
    ##
    version "3.4R1 [tlim]";
    system {
    ##
    ## host-name: Host name for this router
    ## match: ^[[:alnum:]._-]+$
    ## require: system
    ##
    host-name router-name;
    ##
    ## domain-name: Domain name for this router
    ## match: ^[[:alnum:]._-]+$
    ## require: system
    ##
    domain-name isp.net;
    ##
    ## backup-router: Address of router to use while booting
    ##
    backup-router 192.168.100.1;
    root-authentication {
      ##
      ## encrypted-password: Encrypted password string
      ##
      encrypted-password "$1$BYJQE$/ocQof8pmcm7MSGK0"; # SECRET-DATA
    }
    ##
    ## name-server: DNS name servers
    ## require: system
    ##
```

```
        name-server {
            ##
            ## name-server: DNS name server address
            ##
            208.197.1.0;
        }
        login {
            ##
            ## class: User name (login)
            ## match: ^[[:alnum:]_-]+$
            ##
            class super-user {
                ##
                ## permissions: Set of permitted operation categories
                ##
                permissions all;
        }
        ...
        ##
        ## services: System services
        ## require: system
        ##
        services {
            ## services: Service name
            ##
            ftp;
            ##
            ## services: Service name
            ##
            telnet;
            ##
        }
        syslog {
            ##
            ## file-name: File to record logging data
            ##
            file messages {
                ##
                ## Facility type
                ## Level name
                ##
                any notice;
                ##
                ## Facility type
                ## Level name
                ##
            authorization info;
            }
        }
}
```

```
chassis {
    alarm {
      sonet {
        ##
        ## lol: Loss of light
        ## alias: loss-of-light
        ##
        lol red;
        }
      }
    }
}
interfaces {
    ##
    ## Interface name
    ##
    at-2/1/1 {
      atm-options {
        ##
        ## vpi: Virtual path index
        ## range: 0 .. 255
        ## maximum-vcs: Maximum number of virtual circuits on this VP
        ##
        vpi 0 maximum-vcs 512;
      }
      ##
      ## unit: Logical unit number
      ## range: 0 .. 16384
      ##
      unit 0 {
        ##
        ## vci: ATM point-to-point virtual circuit identifier ([vpi.]vci)
        ## match: ^([[:digit:]]+.){0,1}[[:digit:]]+$
        ##
        vci 0.128;
      }
    }
...
```

Filtering Command Output Multiple Times

For the output of a single command, you can filter the output one or more times. For example:

user@host> *command* **| match** *regular-expression* **| except** *regular-expression* **|**
match *other-regular-expression* **| find** *regular-expression* **| hold**

Setting the Current Date and Time

To set the current date and time on the router, use the set date command:

```
user@host> set date YYYYMMDDhhmm.ss
```

YYYY is the four-digit year, *MM* is the two-digit month, *DD* is the two-digit date, *hh* is the two-digit hour, *mm* is the two-digit minute, and ss is the two-digit second. At a minimum, you must specify the two-digit minute. All other parts of the date and time are optional.

To set the time zone, see “Setting the Time Zone” on page 417. To configure time synchronization, see “Configuring the Network Time Protocol” on page 418.

Setting the Date and Time from NTP Servers

If the Network Time Protocol (NTP) server is unable to synchronize the current date and time on the router, a system log message similar to the following appears:

```
"time error %.0f over %d seconds; set clock manually".
```

To set the date and time from all NTP servers configured at the [edit system ntp server] hierarchy level to determine the correct time, use the set date ntp command:

```
user@host> set date ntp
```

NOTE: You do not need to reboot the router when you use the set date ntp command.

To set the date and time from a NTP server configured at the [edit system ntp server] hierarchy level to determine the correct time, use the set date ntp command:

```
user@host> set date ntp  ntp-server
```

To set the date and time from multiple NTP servers configured at the [edit system ntp server] hierarchy level to determine the correct time, use the set date ntp command:

```
user@host>  set date ntp  ntp-server
```

ntp-server is the IP address of one or more NTP servers to query. When querying more than one server, enclose the IP addresses in quotation marks using the format *“ip-address ip-address”*. For example:

```
user@host> set date ntp  "200.49.40.1 129.127.28.4"
10 Feb 13:50:21 ntpdate[794]: step time server 129.127.28.4 offset 0.000163
sec
```

For more information about how to configure NTP, see “Configuring the Network Time Protocol” on page 418 and the *JUNOS System Basics and Services Command Reference*.

Setting the Source Address to Contact the NTP Server

To specify a source address that the JUNOS software uses to contact the remote NTP server configured at the [edit system ntp server], use the set date ntp *source-address* command:

```
user@host> set date ntp source-address source-address
```

source-address is a valid IP address.

Displaying CLI Command History

You can display a list of recent commands that you issued. To display the command history, use the show cli history command:

```
user@host> show cli history
  03-03 01:00:50 -- show cli history
  03-03 01:01:12 -- show interfaces terse
  03-03 01:01:22 -- show interfaces lo0
  03-03 01:01:44 -- show bgp next-hop-database
  03-03 01:01:51 -- show cli history
```

By default, this command displays the last 100 commands issued in the CLI. If you specify a number with the command, it displays that number of recent commands. For example:

```
user@host> show cli history 3
  01:01:44 – show bgp next-hop-database
  01:01:51 – show cli history
  01:02:51 – show cli history 3
```

Displaying CLI Word History

You can type Esc-. or Alt-. to insert the last word of the previous command. Repeat Esc-. or Alt-. to scroll backwards through the list of recently entered words. For example:

```
user@host> show interfaces terse fe-0/0/0
Interface            Admin   Link   Proto  Local              Remote
fe-0/0/0             up      up
fe-0/0/0.0           up      up     inet   192.168.220.1/30
user@host> fe-0/0/0
```

If you scroll completely to the beginning of the list, typing Esc-. or Alt-. again restarts scrolling from the last word entered.

Monitoring Who Uses the CLI

Depending upon how you configure the JUNOS software, multiple users can log in to the router, use the CLI, and configure or modify the software configuration.

The JUNOS software provides a general syslog-like mechanism to log system operations, such as when users log in to the router and when they issue CLI commands. To configure system logging, include the syslog statement in the configuration, as described in "Configuring System Log Messages" on page 427.

If, when you enter configuration mode, another user is also in configuration mode, a notification message is displayed that indicates who the user is and what portion of the configuration they are viewing or editing:

```
user@host> configure
Entering configuration mode
Current configuration users:
    root terminal p3 (pid 1088) on since 1999-05-13 01:03:27 EDT
      [edit interfaces so-3/0/0 unit 0 family inet]
The configuration has been changed but not committed
```

Using the Comment Character

You can copy operational mode commands that include comments from a file and paste them into the command-line interface. A # at the beginning of the command line indicates a comment line. This is useful for describing frequently used operational mode commands; for example, a user's work instructions on how to monitor the network. To add a comment to a command file, the first character of the line must be #. When you start a command with #, the rest of the line is disregarded by the JUNOS software.

To add comments in operational mode, start with a # and end with a new line (carriage return):

```
user@host> # comment-string
```

comment-string is the text of the comment. The comment text can be any length, and you must type it on a single line.

Example: Using Comments

File with Comments

```
#Command 1: Show the router version
show version
#Command 2: Show all router interfaces
show interfaces terse
```

Copy and Paste Contents of the File into the CLI

```
user@host> #Command 1: Show the router version

user@host> show version
Hostname: myhost
Model: m5
JUNOS Base OS boot [6.4-20040511.0]
JUNOS Base OS Software Suite [6.4-20040511.0]
JUNOS Kernel Software Suite [6.4-20040511.0]
JUNOS Packet Forwarding Engine Support (M5/M10) [6.4-20040511.0] JUNOS
Routing Software Suite [6.4-20040511.0] JUNOS Online Documentation
[6.4-20040511.0] JUNOS Crypto Software Suite [6.4-20040511.0]

user@host> # Command 2: Show all router interfaces
regress@fbi> show interfaces terse
Interface          Admin  Link    Proto Local              Remote
fe-0/0/0             up   up
fe-0/0/1             up   down
fe-0/0/2             up   down
mo-0/1/0             up
mo-0/1/0.16383       up   up      inet  10.0.0.1          –> 10.0.0.17
so-0/2/0             up   up
so-0/2/1             up   up
dsc                  up   up
fxp0                 up   up
fxp0.0               up   up      inet  192.168.70.62/21
fxp1                 up   up
fxp1.0               up   up      tnp   4
gre                  up   up
ipip                 up   up
lo0                  up   up
lo0.0                up   up      inet  127.0.0.1          –> 0/0
lo0.16385            up   up      inet inet6
```

JUNOS-FIPS Commands

For routers running JUNOS-FIPS, some commands have restrictions and some additional commands are available. For more information, see the *JUNOS-FIPS Configuration Guide*.

Routing Matrix CLI Enhancements

This section describes how the CLI has been enhanced to accommodate managing a routing matrix. Rather than listing all routing matrix commands and options, it gives key examples to help you in understand how managing the routing matrix can differ from managing a standalone T640 routing node.

This section discusses the following topics:

Routing Matrix Overview

A routing matrix is a multichassis architecture consists of one TX Matrix platform, and from one to four T640 routing nodes, as shown in Figure 3. Each component has two Routing Engines.

Figure 3: Routing Matrix

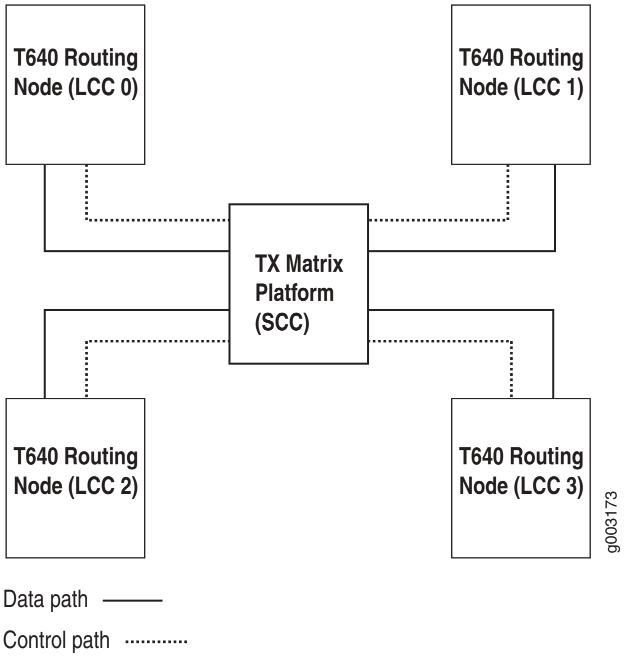
T640 Routing Node (LCC 0)
T640 Routing Node (LCC 1)
TX Matrix Platform (SCC)
T640 Routing Node (LCC 2)
T640 Routing Node (LCC 3)
g003173
Data path
Control path

How the Routing Matrix Is Identified in the CLI

The CLI uses the terms SCC and scc (which stand for *switch-card chassis*) to refer to the TX Matrix platform. Similarly, the CLI uses the terms LCC and lcc (which stand for *line-card chassis*) to refer to the T640 routing nodes in a routing matrix.

The T640 routing nodes are assigned index numbers, LCC0 through LCC3, depending on the hardware setup and how they are connected to the TX Matrix platform. For more information, see the *TX Matrix Platform Hardware Guide.*

Viewing the Routing Matrix as a Single Router

From the CLI, you can view the routing matrix as a single router with many Flexible PIC Concentrators (FPCs) and Physical Interface Cards (PICs).

For example, you can view a list of all hardware components in the routing matrix, view alarms for the entire routing matrix, view interfaces on all T640 routing nodes, and so on. To do so, issue standard operational commands on the TX Matrix platform.

CLI Options for Selecting Routing Matrix Components

When you issue operational mode commands on the TX Matrix platform, CLI command options allow you to apply the command to a component of the routing matrix rather than to the routing matrix as a whole.

These are the options shown in the CLI:

- scc—The TX Matrix platform.
- lcc *number*—A specific T640 routing node.
- all-lcc—All T640 routing nodes.

If you specify none of these options, then the command applies by default to the whole routing matrix: the TX Matrix platform and all connected T640 routing nodes.

Examples of Routing Matrix Command Options

The following output samples, using the show version command, demonstrate the four different options for viewing information about the routing matrix: none, scc, lcc *number*, and all-lcc.

```
user@host> show version ?
Possible completions:
  <[Enter]>            Execute this command
  all-lcc              Show software version on all LCC chassis
  brief                Display brief output
  detail               Display detailed output
  lcc                  Show software version on specific LCC (0..3)
  scc                  Show software version on the SCC
  |                    Pipe through a command
```

Sample Output: No Routing Matrix Options Specified

```
user@host> show version
scc-re0:
--------------------------------------------------------------------------
Hostname: scc
Model: TX Matrix
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]

lcc0-re0:
--------------------------------------------------------------------------
Hostname: lcc0
Model: t640
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]
JUNOS Support Tools Package [7.0-20040630.0]
```

```
lcc1-re0:
---------------------------------------------------------------------------
Hostname: lcc1
Model: t640
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]
JUNOS Support Tools Package [7.0-20040630.0]
```

Sample Output: TX Matrix Platform Only (scc Option)

```
user@host> show version scc
Hostname: scc
Model: TX Matrix
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]
```

Sample Output: Specific T640 Routing Node (lcc *number* Option)

```
user@host> show version lcc 0
lcc0-re0:
---------------------------------------------------------------------------
Hostname: lcc0
Model: t640
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]
JUNOS Support Tools Package [7.0-20040630.0]
```

Sample Output: All T640 routing nodes (all-lcc option)

```
user@host> show version all-lcc
lcc0-re0:
---------------------------------------------------------------------------
Hostname: lcc0
Model: t640
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]
JUNOS Support Tools Package [7.0-20040630.0]
```

```
lcc1-re0:
--------------------------------------------------------------------------
Hostname: lcc1
Model: t640
JUNOS Base OS boot [7.0-20040630.0]
JUNOS Base OS Software Suite [7.0-20040629.0]
JUNOS Kernel Software Suite [7.0-20040630.0]
JUNOS Packet Forwarding Engine Support (T-Series) [7.0-20040630.0]
JUNOS Routing Software Suite [7.0-20040630.0]
JUNOS Online Documentation [7.0-20040630.0]
JUNOS Crypto Software Suite [7.0-20040630.0]
JUNOS Support Tools Package [7.0-20040630.0]
```

Using FPC Numbers in Routing Matrix CLI Commands

A standalone T640 routing node can be configured with up to eight FPCs (numbered 0 through 7). Since a routing matrix can have up to four T640 routing nodes, and each T640 routing node has up to eight FPCs, the routing matrix as a whole can have up to 32 FPCs (0 through 31).

Table 10 shows the correspondence between the FPC hardware slot numbers in the T640 routing nodes and the FPC assignments recognized by the JUNOS software for a routing matrix.

Table 10: FPC Correspondence Between T640 Routing Nodes and the Routing Matrix

T640 Routing Node	T640 FPC Range	Routing Matrix FPC Range
LCC 0	0–7	0–7
LCC 1	0–7	8–15
LCC 2	0–7	16–23
LCC 3	0–7	24–31

The FPC slot numbers appear in two main ways in the CLI:

- In chassis (hardware-based) commands, such as the operational mode show chassis commands and the configuration mode [edit chassis] hierarchy level.
- When configuring or displaying information about interfaces, such as in the configuration mode [edit interfaces] hierarchy level and as part of the interface name in operational mode show interfaces commands.

Using FPC Numbers in Operational Mode Chassis Commands

You can specify the FPC slot number in operational mode chassis commands in two ways:

- Specify the number of the T640 routing node using the lcc *number* option and use a value from 0 through 7 for the FPC slot number.

  This is the recommended method. The show command output also lists information in this way, sorted first by LCC index number and then by FPC slot 0 through 7.

- The lcc *number* option and use a value from 0 through 31 for the FPC slot number.

For example, the following commands have the same result, but the first is the recommended usage:

```
user@host> request chassis fpc lcc 1 slot 1 offline
user@host> request chassis fpc slot 9 offline
```

Specifying FPC Numbers at the [edit chassis lcc *number*] Hierarchy Level

When you are configuring at the [edit chassis lcc *number*] hierarchy level, you must specify the LCC index number and the actual FPC hardware slot number as labeled on the T640 routing node chassis (0 through 7). For more information on configuring the chassis on the routing matrix, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

```
[edit chassis]
lcc lcc-number {
    fpc slot-number { # Use the hardware FPC slot number
        pic pic-number {
            ...
        }
    }
}
```

Specifying FPC Numbers When Configuring Interfaces

When configuring interfaces (or specifying an interface name in show commands), you use the routing matrix FPC range of 0 through 31. For example, the 11 in t1-11/2/0 refers to FPC hardware slot 3 on LCC 1. For more information on configuring interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Operational Commands Issued on Routing Engines

Operational mode commands that you issue on the TX Matrix master Routing Engine are distributed to all master Routing Engines on the T640 routing nodes in the routing matrix.

Commands that you issue on the TX Matrix backup Routing Engine are distributed to all backup Routing Engines on the T640 routing nodes in the routing matrix.

General Operational Tasks

This section describes general operational tasks you can perform on a TX Matrix platform. This section discusses the following topics:

- Upgrading Software on a Routing Matrix on page 210
- Managing Backup Routing Engines on page 210
- Halting and Rebooting Routing Matrix Components on page 210
- Bringing Routing Nodes Offline or Online on page 210
- Managing Files on Routing Engines on page 211
- Displaying Logs on Any Routing Engine on page 211

Upgrading Software on a Routing Matrix

By default, when you upgrade software on the TX Matrix platform, the new image is sent to the master Routing Engines of the connected T640 routing nodes.

Once installation on the TX Matrix platform is complete, a reboot issued on the TX Matrix platform activates the new software on the master Routing Engines of the TX Matrix platform and on all connected T640 routing nodes.

Managing Backup Routing Engines

To manage the backup Routing Engines on all components (for example, to upgrade JUNOS software on backup Routing Engines), you must log in to the TX Matrix platform backup Routing Engine.

Halting and Rebooting Routing Matrix Components

A reboot issued on the TX Matrix platform will reboot the T640 routing node master Routing Engines.

You can halt the TX Matrix platform or a specific T640 routing node within a routing matrix. However, halting both Routing Engines on a TX Matrix platform will halt both Routing Engines on the T640 routing nodes.

Bringing Routing Nodes Offline or Online

You can bring offline or online a specific T640 routing node within a routing matrix:

```
user@host> request chassis lcc ?
Possible completions:
  offline             Take LCC offline
  online              Bring LCC online
  slot                LCC Slot (0..3)
```

Managing Files on Routing Engines

You can manage files on all Routing Engines, for example, copy a file from the TX Matrix master Routing Engine to a T640 routing node Routing Engine:

```
user@host> file list

/var/home/user/:
.ssh/
fred.txt

user@host> file copy fred.txt lcc0:fred.txt

user@host> file list lcc0:
lcc0-master:
--------------------------------------------------------------------------

/var/home/user/:
.ssh/
fred.txt
```

Displaying Logs on Any Routing Engine

You can display logs on any Routing Engines from the TX Matrix platform:

```
user@host> show log lcc0:messages
lcc0-master:
--------------------------------------------------------------------------
Aug 19 17:17:23  lcc0 mgd[7099]: UI_LOAD_EVENT: User 'user' is performing a
'rollback'
Aug 19 17:17:24  lcc0 mgd[7099]: UI_LOAD_EVENT: User 'user' is performing a
'load update'
Aug 19 17:17:25  lcc0 mgd[7099]: UI_COMMIT: User 'user' performed commit: no
comment
```

Checking the Status of T640 Routing Nodes

You can check the status of the T640 routing nodes using the following command:

```
user@host> show chassis lccs
Slot  State            Uptime
0     Online           39 minutes, 16 seconds
1     Online           39 minutes, 16 seconds
2     Empty
3     Empty
```

Configuring the Routing Matrix

This section summarizes some routing matrix configuration guidelines. For more information, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852 and the *JUNOS Network Interfaces Configuration Guide*.

You configure all components of the routing matrix from the TX Matrix platform master Routing Engine:

- Only configuration changes committed on the TX Matrix platform are distributed to all connected T640 routing nodes.
- Any configuration committed on a T640 routing node is not distributed to the TX Matrix platform or other T640 routing nodes.
- A commit on the TX Matrix platform overrides any changes committed on a T640 routing node.
- A standard commit on the master Routing Engine of the TX Matrix platform automatically updates all the connected T640 routing node master Routing Engines.
- A standard commit on the backup Routing Engine of the TX Matrix platform will automatically update all the T640 routing node backup Routing Engines.
- A commit synchronize on the master Routing Engine of the TX Matrix platform updates the master and backup Routing Engines for all components in the routing matrix.

Additional Groups

You can specify two special group names: re0 and re1. These two special group names apply to the Routing Engines in slots 0 and 1 of the TX Matrix platform. In addition, the routing matrix supports special group names for the two Routing Engines in each T640 routing node: lcc *number*-re0 and lcc *number*-re1. *number* identifies a T640 routing node from 0 through 3, for example, lcc0-re0. For more information, see “Creating a Configuration Group” on page 618, and “Example: Creating and Applying Configuration Groups on a TX Matrix Platform” on page 621.

Chapter 7

Controlling the CLI Environment

In operational mode, you can control the command-line interface (CLI) environment. For example, you can specify the number lines that are displayed on the screen or your terminal type. The following output lists the options that you can use to control the CLI environment:

```
user@host> set cli ?
Possible completions:
    complete-on-space Toggle word completion on space
    directory          Set the current working directory
    idle-timeout       Set the cli maximum idle time
    prompt             Set the cli command prompt string
    restart-on-upgrade Set cli to prompt for restart after a software upgrade
    screen-length      Set number of lines on screen
    screen-width       Set number of characters on a line
    terminal           Set terminal type
    timestamp          Timestamp cli output
```

NOTE: When you use SSH to log in to the router or log in from the console when its terminal type is already configured (as described in "Configuring Console and Auxiliary Port Properties" on page 456), your terminal type, screen length, and screen width are already set.

This chapter discusses the following topics:

- Setting the Terminal Type on page 214
- Setting the Screen Length on page 214
- Setting the Screen Width on page 214
- Setting the CLI Prompt on page 214
- Setting the CLI Directory on page 214
- Setting the CLI Timestamp on page 215
- Setting the Idle Timeout on page 215
- Setting the CLI to Prompt after a Software Upgrade on page 215
- Setting Command Completion on page 215

- Displaying CLI Settings on page 216
- Example: Controlling the CLI Environment on page 216

Setting the Terminal Type

To set the terminal type, use the set cli terminal command:

```
user@host> set cli terminal terminal-type
```

The terminal type can be one of the following: ansi, vt100, small-xterm, or xterm.

Setting the Screen Length

The default CLI screen length is 24 lines. To change the length, use the set cli screen-length command:

```
user@host> set cli screen-length length
```

Setting the screen length to 0 lines disables the display of output one screen at a time. Disabling this UNIX more-type interface can be useful when you are issuing CLI commands from scripts.

Setting the Screen Width

The default CLI screen width is 80 columns. To change the width, use the set cli screen-width command:

```
user@host> set cli screen-width width
```

Setting the CLI Prompt

The default CLI prompt is user@host>. To change this prompt, use the set cli prompt command. If the prompt string contains spaces, enclose the string in quotation marks (" ").

```
user@host> set cli prompt string
```

Setting the CLI Directory

To the set the current working directory, use the set cli directory command:

```
user@host> set cli directory directory
```

directory is the pathname of working directory.

Setting the CLI Timestamp

By default, CLI output does not include a timestamp. To include a timestamp in CLI output, use the set cli timestamp command:

```
user@host> set cli timestamp [format time-date-format | disable]
```

If you do not specify a timestamp format, the default format is *Mmm dd hh:mm:ss* (for example, Feb 08 17:20:49). Enclose the format in single quotation marks (‘).

Setting the Idle Timeout

By default, an individual CLI session never times out after extended times, unless the idle-timeout statement has been included in the user’s login class configuration. To set the maximum time an individual session can be idle before the user is logged off the router, use the set cli idle-timeout command:

```
user@host> set cli idle-timeout timeout
```

timeout can be 0 through 100,000 minutes. Setting *timeout* to 0 disables the timeout.

Setting the CLI to Prompt after a Software Upgrade

By default, the CLI prompts you to restart after a software upgrade. To disable the prompt for an individual session, use the set cli restart-on-upgrade off command:

```
user@host> set cli restart-on-upgrade off
```

To re-enable the prompt, use the set cli restart-on-upgrade on command:

```
user@host> set cli restart-on-upgrade on
```

Setting Command Completion

By default, you can press the spacebar or tab key to have the CLI complete a command.

To have the CLI allow only a tab to complete a command, use the set cli complete-on-space off command:

```
user@host> set cli complete-on-space off
Disabling complete-on-space
user@host>
```

To re-enable the use of both spaces and tabs for command completion, use the set cli complete-on-space on command:

```
user@host> set cli complete-on-space on
Enabling complete-on-space
user@host>
```

Displaying CLI Settings

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

```
user@host> show cli
CLI screen length set to 24
CLI screen width set to 80
CLI complete-on-space set to on
```

Example: Controlling the CLI Environment

Change the default CLI environment:

```
user@host> set cli screen-length 66
Screen length set to 66
user@host> set cli screen-width 40
Screen width set to 40
user@host> set cli prompt “router1-san-jose > ”
router1-san-jose > 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 66
CLI screen width set to 40
CLI terminal is ‘xterm’
router1-san-jose >
```

Chapter 8
Configuring the Router with the CLI

You configure the JUNOS software by entering configuration mode and creating a hierarchy of configuration mode statements. In configuration mode, the command-line interface (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.

When you first log on to the router, you enter the CLI operational mode. Operational mode is indicated by the presence of the > prompt, which is preceded by string that defaults to the name of the user and the name of the router. For example:

```
user@host>
```

You enter configuration mode by issuing the configure command or the edit command from the CLI operation mode. When you do this, the CLI prompt changes from user@host> to user@host#. Configuration mode is indicated by the presence of the # prompt, which is preceded by string that defaults to the name of the user and the name of the router. For example:

```
user@host>
user@host> configure
entering configuration mode
[edit]
user@host#
```

To view a list of configuration mode statements, see "Entering Configuration Mode" on page 222. For information about CLI enhancements for the TX Matrix platform and its connected T640 routing nodes, see "Routing Matrix CLI Enhancements" on page 204.

This chapter discusses the following topics:

For information about the configuration statements to use to configure particular system functionality, see the chapter about that feature.

For general guidelines for using the CLI to configure a routing matrix, see “Configuring the Routing Matrix” on page 212.

Configuration Statement Hierarchy

The JUNOS software configuration consists of a hierarchy of *statements*. There are two types of statements: *container statements*, which are statements that 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 at the top level of the configuration hierarchy resides at the trunk (or root level) of a hierarchy tree. The top-level statements are container statements, containing other statements that form the tree branches. The leaf statements are the leaves of the hierarchy tree. An individual hierarchy of statements, which starts at the trunk of the hierarchy tree, is called a *statement path*. Figure 4 illustrates the hierarchy tree, showing a statement path for the portion of the protocol configuration hierarchy that configures the hello interval on an interface in an Open Shortest Path First (OSPF) area. The protocols statement is a top-level statement at the trunk of the configuration tree. The ospf, area, and interface statements are all subordinate container statements of a higher statement (they are branches of the hierarchy tree), and the hello-interval statement is a leaf on the tree, which, in this case, contains a data value, the length of the hello interval, in seconds.

Figure 4: Configuration Mode Hierarchy of Statements

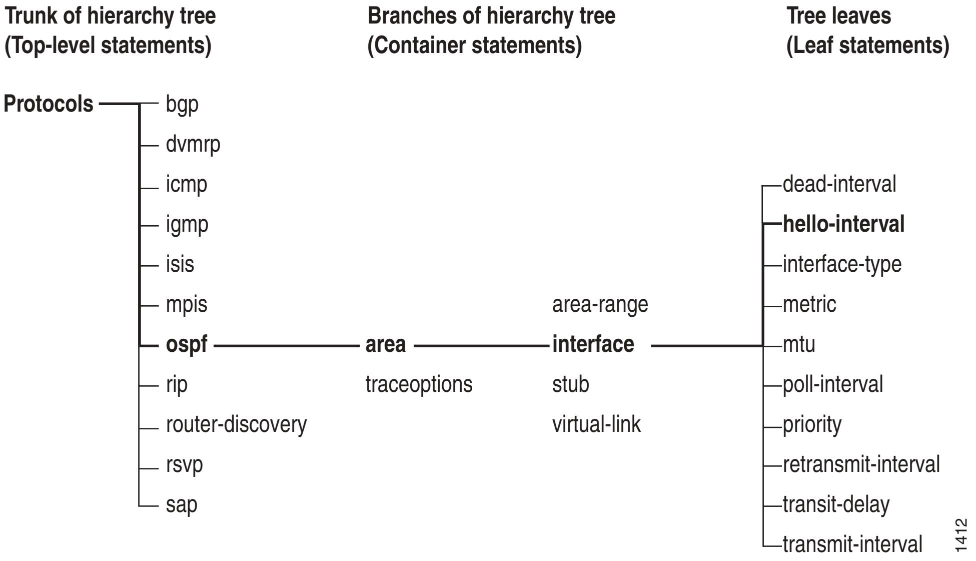

The CLI represents the statement path shown in Figure 4 as [protocols ospf area *area-number* interface *interface-name*], and displays the configuration as follows:

```
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/0 {
                hello-interval 5;
            }
            interface so-0/0/1 {
                hello-interval 5;
            }
        }
    }
}
```

The CLI indents each level in the hierarchy to indicate each statement's relative position in the hierarchy and generally sets off each level 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 printed. Each leaf statement ends with a semicolon. If the hierarchy does not extend as far as a leaf statement, the last statement in the hierarchy ends with a semicolon.

The CLI uses this indented representation when it displays the current system configuration, and you use this format when creating ASCII files that contain the software configuration. However, the format of ASCII configuration files is not as strict as the CLI output of the configuration. Although the braces and semicolons are required, the indention and use of new lines, as shown above, are not required in ASCII configuration files.

How the Configuration Is Stored

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 in the CLI immediately, so if multiple users are editing the configuration at the same time, all users can see all changes.

To have a candidate configuration take effect, you *commit* the changes. At this point, 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.

In addition to saving the current configuration, the CLI 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, which is the default configuration that the system returns to if you roll back to a previous configuration), and the oldest saved configuration is version 49. The currently operational JUNOS software configuration is stored in the file juniper.conf, and the last three committed configurations are stored in the files juniper.conf.1, juniper.conf.2, and juniper.conf.3. These four files are located in the directory /config, which is on the router's flash drive. The remaining 46 previous versions of committed configurations, the files juniper.conf.4 through juniper.conf.49, are stored in the directory /var/db/config on the hard disk.

Figure 5 illustrates the various router configuration states and the configuration mode commands you use to load, commit, copy, save, or roll back the configuration.

Figure 5: Commands for Storing and Modifying the Router Configuration

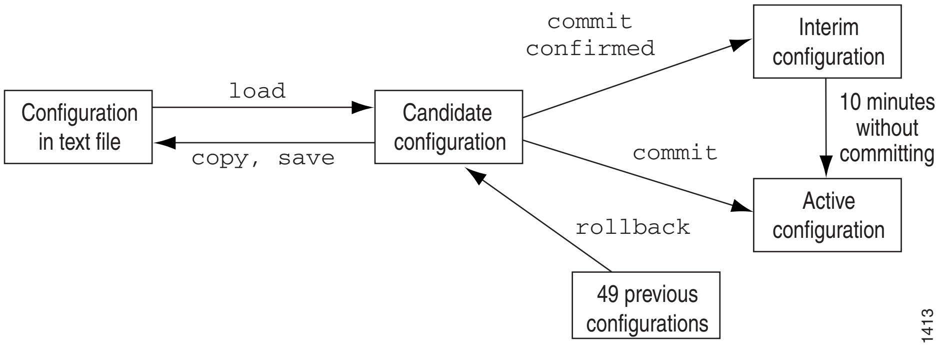

Entering Configuration Mode

You enter configuration mode by entering the configure operational mode command.

The following configuration mode commands are available:

```
user@host> configure
entering configuration mode
[edit]
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 an ASCII file
  quit         Quit from this level
  rename       Rename a statement
  rollback     Roll back database to last committed version
  run          Run an operational-mode command
  save         Save configuration to an ASCII file
  set          Set a parameter
  show         Show a parameter
  status       Display database user status
  top          Exit to top level of configuration
  up           Exit one level of configuration
```

The access privilege level required to enter configuration mode is controlled by the configure permission bit. Users for whom this permission bit is not set do not see the configure command as a possible completion when they enter a ? in operational mode, and they cannot enter configuration mode. Users for whom this bit is set do see this command and can enter configuration mode. When in configuration mode, a user can view and modify only those statements for which they have access privileges set. For more information, see “Configuring Access Privilege Levels” on page 402.

This section discusses the following topics:

- Using the Configure Command on page 223
- Using the Configure Exclusive Command on page 224
- Using the Configure Private Command on page 225

Using the Configure Command

If you and other users enter configuration mode with the configure command, everyone can make configuration changes and commit all changes made to the configuration. This means that if you and another user have made configuration changes and the other user commits, the changes you made are committed as well. That is, no one has a lockout on the configuration file.

If, when you enter configuration mode, another user is also in configuration mode, a message shows who the user is and what part of the configuration that user is viewing or editing:

```
user@host> configure
Entering configuration mode
Current configuration users:
    root terminal p3 (pid 1088) on since 1999-05-13 01:03:27 EDT
      [edit interfaces so-3/0/0 unit 0 family inet]
The configuration has been changed but not committed
[edit]
user@host>
```

If, when you enter configuration mode, the configuration contains changes that have not been committed, a message appears:

```
user@host> configure
Entering configuration mode
The configuration has been changed but not committed
[edit]
user@host>
```

If, while in configuration mode, you try to make a change while the configuration is locked by another user, a message indicates that the configuration database is locked, who the user is, and what portion of the configuration the user is viewing or editing:

```
user@host# set system host-name ipswitch
error: configuration database locked by:
    user2 terminal d0 (pid 1828) on since 19:47:58 EDT, idle 00:02:11
      exclusive [edit protocols]
```

Using the Configure Exclusive Command

If you enter configuration mode with the configure exclusive command, you lock the candidate configuration for as long as you remain in configuration mode, allowing you to make changes without interference from other users. Other users can enter and exit configuration mode, but they cannot change the configuration. If another user has locked the configuration, and you need to forcibly log him or her out, enter the operational mode command request system logout pid *pid_number*. When a user exits from configure exclusive mode when another user is in configure private mode, the JUNOS software will roll back any uncommitted changes.

If, when you enter configuration mode, another user is also in configuration mode and has locked the configuration, a message indicates who the user is and what portion of the configuration that user is viewing or editing:

```
user@host> configure
Entering configuration mode
Users currently editing the configuration:
    root terminal p3 (pid 1088) on since 2000-10-30 19:47:58 EDT, idle
00:00:44
    exclusive [edit interfaces so-3/0/0 unit 0 family inet]
```

NOTE: If you are using the configure exclusive command, you cannot exit configuration mode with uncommitted changes while another user is in a configure private session. For more information about the configure private command, see “Using the Configure Private Command” on page 225.

Users in configure exclusive mode cannot exit configuration mode with uncommitted changes. A warning message appears notifying the user in configure exclusive mode that any changes will be discarded if the user exits from the configuration:

```
user@host> configure exclusive
warning: uncommitted changes will be discarded on exit
Entering configuration mode

[edit]
user@host# set system host-name cool

[edit]
user@host# quit
The configuration has been changed but not committed
warning: Auto rollback on exiting 'configure exclusive'
Discard uncommitted changes? [yes,no] (yes)

warning: discarding uncommitted changes
load complete
Exiting configuration mode
```

When you use the yes option to exit configure exclusive mode, the JUNOS software discards your uncommitted changes and rolls backs your configuration. The no option allows you to continue editing or to commit your changes in configure exclusive mode. These options enforce the restriction that the global configuration must be unmodified for configure private users to commit changes.

Using the Configure Private Command

The configure private command allows multiple users to edit different parts of the configuration at the same time and to commit only their own changes, or to roll back without interfering with one another's changes. When you issue the configure private command, you work in a private candidate configuration, which is a copy of the most recently committed configuration.

When you commit a private candidate configuration, the JUNOS software temporarily locks the global configuration, enforces the restriction that the global configuration must be unmodified to commit private changes, and validates the private candidate configuration. If a merge conflict occurs, the commit fails and the configuration lock is released. You can then modify your private candidate configuration and commit it again. If there are no errors, the changes made in the private candidate configuration are merged into the most recently committed global configuration, are activated, and begin running on the router, and the configuration lock is released.

NOTE: You cannot commit changes in configure private mode when another user is in configure exclusive mode.

If the global configuration has changed, users in configure private mode can issue the rollback or update command to obtain the most recently committed global configuration. For more information about the update command, see "Updating the Configure Private Configuration" on page 227.

You must issue the commit command from the top of the configuration.

You cannot save a configure private session; uncommitted changes are discarded.

You cannot issue the commit confirm command when you are in configure private mode.

Users in configure exclusive mode cannot exit configuration mode with uncommitted changes while another user is in configure private mode. A warning message appears notifying the user in configure exclusive mode that any changes will be discarded if the user exits from the configuration:

```
[edit]
user@host# set system host-name fu

[edit]
user@host# quit
The configuration has been changed but not committed
warning: private edits in use. Auto rollback on exiting 'configure exclusive'
Discard uncommitted changes? [yes,no] (yes)

load complete
Exiting configuration mode
user@host
```

When you use the **yes** option to exit configure exclusive mode, the JUNOS software discards your uncommitted changes and rolls backs your configuration. The **no** option allows you to continue editing or to commit your changes in configure exclusive mode. These options enforce the restriction that the global configuration must be unmodified for users to commit configure private changes.

NOTE: You cannot enter configure private mode when the global configuration has been modified.

If a configure private edit is in session, users who issue the configure command can only view the global configuration; a message appears indicating that these users must use the configure exclusive or configure private commands to modify the configuration:

```
[edit]
user@host# set system host-name ipswitch
error: private edits in use. Try 'configure private' or 'configure
exclusive'.
[edit]
user@host#
```

If the global configuration has been modified, users cannot enter configure private mode because they cannot commit changes when the global configuration has been modified. For example:

```
user@host# configure private
error: shared configuration database modified
Users currently editing the configuration:
root terminal d0 (pid 7951) on since 2002-02-21 14:18:46 PST
[edit]
user@host#
```

NOTE: Users in configure or configure exclusive mode cannot exit the global configuration with uncommitted changes.

If another user commits a change to the same section of the configuration that the private user has modified, a merge conflict may result. The JUNOS software then updates the private user’s configuration with the most recently committed global configuration and the private user can commit the changes. For example:

```
[edit]
user@host# set system host-name foo

[edit]
user@host# show | compare
[edit system]
-  host-name host;
+  host-name foo;
```

```
[edit]
user@host# commit
[edit system host-name]
  'host-name bar'
    statement does not match patch; 'bar' != 'host'
load complete (1 errors)

[edit]
user@host# show | compare
[edit system]
-  host-name bar;
+  host-name foo;

[edit]
user@host#
```

In this example, after the JUNOS software detects the merge conflict and fixes it, the user in configure private mode issues the show | compare command. This command displays the private user’s database changes against the most recently committed global configuration.

Updating the Configure Private Configuration

When you are in configure private mode, you must work with a copy of the most recently committed global configuration. If the global configuration changes, you can issue the update command to update your private candidate configuration. When you do this, your private candidate configuration contains a copy of the most recently committed configuration with your private changes merged in. For example:

```
[edit]
user@host# update

[edit]
user@host#
```

NOTE: You can get merge conflicts when you issue the update command.

You can also issue the rollback command to discard your private candidate configuration changes and obtain the most recently committed configuration:

```
[edit]
user@host# rollback

[edit]
user@host#
```

Configuration Mode Prompt

In configuration mode, the prompt changes from a > to a #. For example:

```
user@host> configure
entering configuration mode
[edit]
user@host#
```

Configuration Mode Banner

The portion of the prompt in braces, [edit], is a *banner*. The banner indicates that you are in configuration mode and shows your location in the statement hierarchy. When you first enter configuration mode, you always are at the top level of the hierarchy, which is indicated by the [edit] banner. For example:

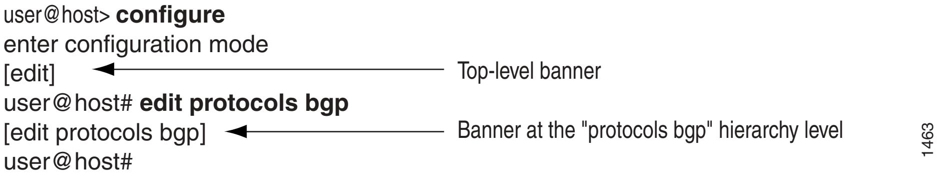

```
user@host> configure
enter configuration mode
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host#
```

NOTE: When the word ORPHANED appears to the right of the [edit] banner, it indicates that part of the configuration has been deleted by another user, and you are not working with the most recent global candidate configuration; for example:

```
[edit protocols bgp] ORPHANED
user@host#
```

To refresh your view of the global configuration, move to the top level of the hierarchy.

Configuration Statements and Identifiers

You configure all router properties by including *statements* in the configuration. A statement consists of a keyword, which is fixed text, and, optionally, an *identifier*. An identifier is an identifying name that you define, such as the name of an interface or a username, and that allows you and the CLI to discriminate among a collection of statements.

The following list shows the statements available at the top level of configuration mode (that is, the trunk of the hierarchy tree). Table 11 on page 230 describes each statement.

```
user@host# set ?
Possible completions:
> accounting-options    Accounting data configuration
+ apply-groups          Groups from which to inherit configuration data
> chassis               Chassis configuration
> class-of-service      Class-of-service configuration
> firewall              Define a firewall configuration
> forwarding-options    Configure options to control packet sampling
> groups                Configuration groups
> interfaces            Interface configuration
> policy-options        Routing policy option configuration
> protocols             Routing protocol configuration
> routing-instances     Routing instance configuration
> routing-options       Protocol-independent routing option configuration
> snmp                  Simple Network Management Protocol
> system                System parameters
```

An angle bracket (>) before the statement name indicates that it is a container statement and that you can define other statements at levels below it.

If there is no angle bracket (>) before the statement name, the statement is a leaf statement; you cannot define other statements at hierarchy levels below it.

A plus sign (+) before the statement name indicates that it can contain a set of values. To specify a set, include the values in brackets. For example:

```
[edit]
user@host# set policy-options community my-as1-transit members [65535:10
65535:11]
```

In some statements, you can include an identifier. For some identifiers, such as interface names, you must specify the identifier in a precise format. For example, the interface name so-0/0/0 refers to a SONET/SDH interface that is on the Flexible PIC Concentrator (FPC) in slot 0, in the first PIC location, and in the first port on the Physical Interface Card (PIC). For other identifiers, such as interface descriptive text and policy and firewall term names, you can specify any name, including special characters, spaces, and tabs.

You must enclose in quotation marks (double quotes) identifiers and any strings that include the following characters: space tab ( ) [ ] { } ! @ # $ % ^ & | ’ = ?

Table 11: Configuration Mode Top-Level Statements

Statement	Description
access	Configure the Challenge Handshake Authentication Protocol (CHAP). For information about the statements in this hierarchy, see "Access" on page 645.
accounting-options	Configure accounting statistics data collection for interfaces and firewall filters. For information about the statements in this hierarchy, see the *JUNOS Network Management Configuration Guide*.
chassis	Configure properties of the router chassis, including conditions that activate alarms and SONET/SDH framing and concatenation properties. For information about the statements in this hierarchy, see "Router Chassis" on page 797.
class-of-service	Configure class-of-service parameters. For information about the statements in this hierarchy, see the *JUNOS Class of Service Configuration Guide.*
firewall	Define filters that select packets based on their contents. For information about the statements in this hierarchy, see the *JUNOS Policy Framework Configuration Guide*.
forwarding-options	Define forwarding options, including traffic sampling options. For information about the statements in this hierarchy, see the *JUNOS Network Interfaces Configuration Guide.*
groups	Configure configuration groups. For information about statements in this hierarchy, see "Configuration Groups" on page 615.
interfaces	Configure interface information, such as encapsulation, interfaces, virtual channel identifiers (VCIs), and data-link connection identifiers (DLCIs). For information about the statements in this hierarchy, see the *JUNOS Network Interfaces Configuration Guide.*
policy-options	Define routing policies, which allow you to filter and set properties in incoming and outgoing routes. For information about the statements in this hierarchy, see the *JUNOS Routing Protocols Configuration Guide*.
protocols	Configure routing protocols, including Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Label Distribution Protocol (LDP), Multiprotocol Label Switching (MPLS), OSPF, Routing Information Protocol (RIP), and Resource Reservation Protocol (RSVP). For information about the statements in this hierarchy, see the chapters that discuss how to configure the individual routing protocols in the *JUNOS Routing Protocols Configuration Guide* and the *JUNOS MPLS Applications Configuration Guide*.
routing-instances	Configure multiple routing instances. For information about the statements in this hierarchy, see the *JUNOS Routing Protocols Configuration Guide*.
routing-options	Configure protocol-independent routing options, such as static routes, autonomous system numbers, confederation members, and global tracing (debugging) operations to log. For information about the statements in this hierarchy, see the *JUNOS Routing Protocols Configuration Guide*.
security	Configure IP Security (IPSec) services. For information about the statements in this hierarchy see "Security Services" on page 709.
snmp	Configure Simple Network Management Protocol (SNMP) community strings, interfaces, traps, and notifications. For information about the statements in this hierarchy, see the *JUNOS Network Management Configuration Guide*.
system	Configure systemwide properties, including the hostname, domain name, Domain Name System (DNS) server, user logins and permissions, mappings between hostnames and addresses, and software processes. For information about the statements in this hierarchy, see "System Management Configuration Statements" on page 365.

Getting Help About Configuration Mode Commands, Statements, and Identifiers

Configuration mode provides two different types of help:

- Using Command Completion in Configuration Mode on page 231
- Getting Help Based on a String in a Statement Name on page 233

Using Command Completion in Configuration Mode

The CLI command completion functions, described in "Using CLI Complete Commands" on page 185, for operational mode commands, also apply to the commands in configuration mode and to configuration statements. Specifically, to display all possible commands or statements, type the partial string followed immediately by a question mark; to complete a command or statement that you have partially typed, press the tab key or spacebar.

Command completion also applies to identifiers, with one slight difference. To display all possible identifiers, type a partial string followed immediately by a question mark. To complete an identifier, you must press the tab key. This scheme allows you to enter identifiers with similar names; then press the spacebar when you are done typing the identifier name.

Examples: Using Command Completion in Configuration Mode

List the configuration mode commands:

```
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 an ASCII file
  quit           Quit from this level
  rename         Rename a statement
  rollback       Roll back database to last committed version
  run            Run an operational-mode command
  save           Save configuration to an ASCII file
  set            Set a parameter
  show           Show a parameter
  status         Display database user status
  top            Exit to top level of configuration
  up             Exit one level of configuration
```

List all the statements available at a particular hierarchy level:

```
[edit]
user@host# edit ?
Possible completions:
> accounting-options  Accounting data configuration
> chassis             Chassis configuration
> class-of-service    Class-of-service configuration
> firewall            Define a firewall configuration
> forwarding-options  Configure options to control packet sampling
> groups              Configuration groups
> interfaces          Interface configuration
> policy-options      Routing policy option configuration
> protocols           Routing protocol configuration
> routing-instances    Routing instance configuration
> routing-options     Protocol-independent routing option configuration
> snmp                Simple Network Management Protocol
> system              System parameters

user@host# edit protocols ?
Possible completions:
<[Enter]>           Execute this command
> bgp                 BGP options
> connections         Circuit cross-connect configuration
> dvmrp               DVMRP options
> igmp                IGMP options
> isis                IS-IS options
> ldp                 LDP options
> mpls                Multiprotocol Label Switching options
> msdp                MSDP options
> ospf                OSPF configuration
> pim                 PIM options
> rip                 RIP options
> router-discovery    ICMP router discovery options
> rsvp                RSVP options
> sap                 Session Advertisement Protocol options
> vrrp                VRRP options
  |                   Pipe through a command
[edit]
user@host# edit protocols
```

List all commands that start with a particular letter or string:

```
user@host# edit routing-options a?
Possible completions:
    > aggregate          Coalesced routes
    > autonomous-systemAutonomous system number
[edit]
user@host# edit routing-options a
```

List all configured Asynchronous Transfer Mode (ATM) interfaces:

```
user@host# edit interfaces at?
Possible completions:
    <interface_name>     Interface name
    at-2/1/1
    at-2/2/0
    at-5/1/0
[edit]
user@host# edit interfaces at
```

Display a list of all configured policy statements:

```
[edit]
user@host# show policy-options policy-statement ?
Possible completions:
<policy_name>       Name to identify a policy filter
[edit]
user@host# edit policy-options policy-statement
```

Getting Help Based on a String in a Statement Name

In configuration mode, you can use the help command to display help based on a text string contained in a statement name. This command displays help for statements at the current hierarchy level and below.

```
help apropos string
```

string is a text string about which you want to get help. This string is used to match statement names as well as the help strings that are displayed for the statements. If the string contains spaces, enclose it in quotation marks (" "). You also can specify a regular expression for the string, using standard UNIX-style regular expression syntax.

You can also display help based on a text string contained in a statement name using the help topic and help reference commands:

```
help topic string
help reference string
```

The help topic command displays usage guidelines for the statement, while the help reference command displays summary information about the statement.

For introductory information on using the help command, type help and press Enter:

```
help
```

Example: Getting Help Based on a String in a Statement Name

Get help about statements that contain the string “traps”:

```
[edit]
user@host# help apropos traps
set interfaces <interface_name>
    Enable SNMP notifications on state changes
set interfaces <interface_name> unit <interface_unit_number>
    Enable SNMP notifications on state changes
set snmp trap-group
    Configure traps and notifications
set snmp trap-group <group_name> version <version> all
    Send SNMPv1 and SNMPv2 traps
set snmp trap-group <group_name> version <version> v1
    Send SNMPv1 traps
set snmp trap-group <group_name> version <version> v2
    Send SNMPv2 traps
set protocols mpls log-updown
    Send SNMP traps
set firewall filter <filter-name> term <rule-name> from source-port snmptrap
    SNMP traps
set firewall filter <filter-name> term <rule-name> from source-port-except
snmptrap
    SNMP traps
set firewall filter <filter-name> term <rule-name> from destination-port snmptrap
    SNMP traps
set firewall filter <filter-name> term <rule-name> from destination-port-except
snmptrap
    SNMP traps
set firewall filter <filter-name> term <rule-name> from port snmptrap
    SNMP traps
set firewall filter <filter-name> term <rule-name> from port-except snmptrap
    SNMP traps
[edit]
user@host# edit interfaces at-5/3/0
[edit interfaces at-5/3/0]
user@host# help apropos traps
set <interface_name>
    Enable SNMP notifications on state changes
set <interface_name> unit <interface_unit_number>
    Enable SNMP notifications on state changes
```

Creating and Modifying the Configuration

To configure the router or to modify an existing router configuration, you add statements to the configuration, in the process creating a statement hierarchy. For each statement hierarchy, you create the hierarchy starting with a statement at the top level and continuing with statements that move progressively lower in the hierarchy.

For example, to configure an interface in OSPF area 0, you must configure the following hierarchy of statements:

```
protocols
    ospf
      area 0
        interface interface-name
```

To create the hierarchy, you use two configuration mode commands:

- set—Creates a statement hierarchy and sets identifier values. After you issue a set command, you remain at the same level in the hierarchy.

  The set command has the following syntax:

  ```
  user@host# set <statement-path> statement <identifier>
  ```

  statement-path is the hierarchy to the configuration statement and the statement itself. If you have already moved to the statement's hierarchy level, you omit this.

  statement is the configuration statement itself.

  identifier is a string that identifies an instance of a statement. Not all statements require identifiers. In the example shown at the beginning of this section, the area name and the interface names are identifiers. In many cases, the identifier can contain a space. When you type these identifiers in the configuration, you must enclose them in quotation marks. When the CLI displays these identifiers in the output of a show or other command, it encloses them in quotation marks.

  The set command is analogous to an operating system command in which you specify the full pathname of the statement you are performing an action on, for example, mkdir /usr/home/boojum/files or mkdir f:\home\boojum\files.

  For statements that can have more than one identifier, when you issue a set command to set an identifier, only that identifier is set. The other identifiers that are specified in the statement remain.

- edit—Moves to a particular hierarchy level. If that hierarchy level does not exist, the edit command creates it and then moves to it. After you issue an edit command, the banner changes to indicate your current level in the hierarchy.

  The edit command has the following general syntax:

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

  The edit command is analogous to the combination of operating system commands that you would use to first change to a directory and then perform an action; for example, cd /usr/home/boojum;mkdir files.

Examples: Creating and Modifying the Configuration

To configure an interface to run OSPF, you could issue a single set command from the top level of the configuration hierarchy. The initial [edit] banner indicates that you are at the top level. Notice that after you issue the set command, you remain at the top level of the statement hierarchy, as indicated by the second [edit] banner.

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

You also can use the edit command to create and move to the [edit protocols ospf area 0.0.0.0 interface so-0/0/0] hierarchy level and then issue a set command to set the value of the hello-interval statement. After you issue the edit command, you move down in the hierarchy, as indicated by the [edit protocols ospf area 0.0.0.0 interface so-0/0/0] banner.

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

Because hello-interval is an identifier and not a statement, you cannot use the edit command to set the hello interval value. You must use the set command. You can determine that hello-interval is an identifier by listing the available commands at the [edit protocols ospf area 0.0.0.0 interface so-0/0/0] banner. All the statements *not* preceded by a > are identifiers.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0 interface so-0/0/0
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# set ?
Possible completions:
+ apply-groups        Groups from which to inherit configuration data
> authentication-key  Authentication key
  dead-interval       Dead interval (seconds)
  disable             Disable OSPF on this interface
  hello-interval      Hello interval (seconds)
  interface-type      Type of interface
  metric              Interface metric (1..65535)
> neighbor            NBMA neighbor
```

```
  passive               Do not run OSPF, but advertise it
  poll-interval         Poll interval for NBMA interfaces
  priority              Designated router priority
  retransmit-interval   Retransmission interval (seconds)
  transit-delay         Transit delay (seconds)
  transmit-interval     OSPF packet transmit interval (milliseconds)
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# set
```

In both examples above, using either just the set command or a combination of the set and edit commands, you create the same configuration hierarchy:

```
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/0 {
          hello-interval 5;
        }
      }
    }
}
```

Notice that the CLI uses indentation to visually represent the hierarchy levels, and it also places braces at the beginning and end of each hierarchy level to set them off. The CLI also places a semicolon at the end of the line that configures the hello-interval statement.

You also use the set command to modify the value of an existing identifier. The following example changes the hello interval in the configuration shown above:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/0 hello-interval 20
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/0 {
          hello-interval 20;
        }
      }
    }
}
```

When a statement can have more than one identifier, use the set command to add additional identifiers. Any identifiers that you have already set remain set.

Moving Among Levels of the 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:

```
user@host> configure
entering configuration mode
[edit]
user@host#
```

This section discusses the following topics:

- Moving Down to a Specific Level on page 238
- Moving Back Up to Your Previous Level on page 238
- Moving Up One Level on page 239
- Moving Directly to the Top of the Hierarchy on page 239
- Warning Messages When Moving Up on page 239
- Issuing Relative Configuration Commands on page 240

Moving Down to a Specific Level

To move down through an existing configuration command hierarchy, or to create a hierarchy and move down to that level, use the edit configuration mode command, specifying the hierarchy level at which you want to be. After you issue an edit command, the banner changes to indicate your current level in the hierarchy.

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

For example:

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

Moving Back Up to Your Previous Level

To move up the hierarchy, use the exit configuration mode command. This command is, in effect, the opposite of the edit command. That is, the exit command moves you back to your previous level. For example:

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

Moving Up One Level

To move up the hierarchy one level at a time, use the up configuration mode command. For example:

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

Moving Directly to the Top of the Hierarchy

To move directly to the top level, use the top configuration mode command. For example:

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

Warning Messages When Moving Up

If you have omitted a required statement at a particular level, when you issue a show command that displays that hierarchy level, a warning message indicates which statement is missing. For example:

```
[edit protocols mpls]
user@host# set statistics file
[edit protocols mpls]
user@host# show
statistics {
    file; # Warning: missing mandatory statement(s): <filename>
}
interface all;
interface so-3/0/0 {
  disable;
}
```

Issuing Relative Configuration Commands

You can issue configuration mode commands from the top of the hierarchy, or from a level above the area you are configuring. This enables you to perform configurations without having to move from your current location in the hierarchy. To do this, use the top or up commands followed by another configuration command, including edit, insert, delete, deactivate, annotate, or show.

To issue configuration mode commands from the top of the hierarchy, use the top command; then specify a configuration command. For example:

```
[edit interfaces fxp0 unit 0 family inet]
user@host# top edit system login
[edit system login]
user@host#
```

To issue configuration mode commands from a location higher in the hierarchy, use the up configuration mode command; then specify a configuration command. For example:

```
[edit protocols bgp]
user@host# up 2 activate system
```

Exiting Configuration Mode

To exit configuration mode, use the exit configuration-mode configuration mode command from any level, or use the exit command from the top level. For example:

```
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# exit configuration-mode
exiting configuration mode
user@host>

[edit]
user@host# exit
exiting configuration mode
user@host>
```

If you try to exit from configuration mode using the exit command and the configuration contains changes that have not been committed, you see a message and prompt:

```
[edit]
user@host# exit
The configuration has been changed but not committed
Exit with uncommitted changes? [yes,no] (yes) <Enter>
Exiting configuration mode
user@host>
```

To exit with uncommitted changes without having to respond to a prompt, use the exit configuration-mode command. This command is useful when you are using scripts to perform remote configuration.

```
[edit]
user@host# exit configuration-mode
The configuration has been changed but not committed
Exiting configuration mode
user@host>
```

Displaying the Current Configuration

To display the current configuration, use the show configuration mode command. This command displays the configuration at the current hierarchy level or at the specified level.

```
user@host> show <statement-path>
```

When displaying the configuration, the CLI indents each subordinate hierarchy level, inserts braces to indicate the beginning and end of each hierarchy level, and places semicolons at the end of statements that are at the lowest level of the hierarchy. You use the same format when creating an ASCII configuration file, and the CLI uses the same format when saving a configuration to an ASCII file.

The configuration statements appear in a fixed order, and interfaces appear alphabetically by type, and then in numerical order by slot number, PIC number, and port number. Note that when you configure the router, you can enter statements in any order.

You also can use the CLI operational mode show configuration command to display the last committed current configuration, which is the configuration currently running on the router:

```
user@host> show configuration
```

If you have omitted a required statement at a particular hierarchy level, when you issue the show command in configuration mode, a message indicates which statement is missing. As long as a mandatory statement is missing, the CLI continues to display this message each time you issue a show command. For example:

```
[edit]
user@host# show
protocols {
    pim {
      interface so-0/0/0 {
        priority 4;
        version 2;
        # Warning: missing mandatory statement(s): 'mode'
      }
    }
}
```

Examples: Displaying the Current Configuration

Display the entire configuration:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/0 hello-interval 5
[edit]
user@host# show
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/0 {
                hello-interval 5;
            }
        }
    }
}
```

Display a particular hierarchy in the configuration:

```
[edit]
user@host# show protocols ospf area 0.0.0.0
interface so-0/0/0 {
    hello-interval 5;
}
```

Move down to a level and display the configuration at that level:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@host# show
interface so-0/0/0 {
    hello-interval 5;
}
```

Display all of the last committed configuration:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/0 hello-interval 5
[edit]
user@host# commit
commit complete
[edit]
user@host# quit
exiting configuration mode
user@host> show configuration
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/0 {
                hello-interval 5;
            }
        }
    }
}
```

Displaying set Commands from the Configuration

In configuration mode, you can display the configuration as a series of configuration mode commands required to recreate the configuration. This is useful if you are not familiar with how to use configuration mode commands or if you want to cut, paste, and edit the displayed configuration. For information about the set command, see "Creating and Modifying the Configuration" on page 235.

To display the configuration as a series of configuration mode commands required to recreate the configuration from the top level of the hierarchy as set commands, issue the show configuration mode command with the | display set option:

```
user@host# show | display set
```

Example: Displaying set Commands from the Configuration

Display the set commands from the configuration at the [edit interfaces] hierarchy level:

```
[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
    family inet {
      address 192.107.1.230/24;
    }
    family iso;
    family mpls;
}
inactive: unit 1 {
    family inet {
      address 10.0.0.1/8;
    }
}
user@host# show | display set
set interfaces fe-0/0/0 unit 0 family inet address 192.107.1.230/24
set interfaces fe-0/0/0 unit 0 family iso
set interfaces fe-0/0/0 unit 0 family mpls
set interfaces fe-0/0/0 unit 1 family inet address 10.0.0.1/8
deactivate interfaces fe-0/0/0 unit 1
```

To display the configuration as a series of configuration mode commands required to recreate the configuration from the current hierarchy level, issue the show configuration mode command with the | display set relative option:

```
user@host# show | display set relative
```

Example: Displaying Required set Commands at the Current Hierarchy Level

Display the configuration as a series of configuration mode commands required to recreate the configuration from the current hierarchy level:

```
[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
    family inet {
      address 192.107.1.230/24;
    }
    family iso;
    family mpls;
}
inactive: unit 1 {
    family inet {
      address 10.0.0.1/8;
    }
}
user@host# show | display set relative
set unit 0 family inet address 192.107.1.230/24
set unit 0 family iso
set unit 0 family mpls
set unit 1 family inet address 10.0.0.1/8
deactivate unit 1
```

To display the configuration as set commands and search for text matching a regular expression by filtering output, specify the match option after the pipe:

```
user@host# show | display set | match regular-expression
```

For more information about how to use the match option, see "Searching for a String in the Output" on page 191.

Example: Displaying set Commands with the Match Option

Display IP addresses associated with an interface:

```
ge-2/3/0 {
    unit 0 {
        family inet {
            address 192.107.9.106/30;
        }
    }
}
so-5/1/0 {
    unit 0 {
        family inet {
            address 192.107.9.15/32 {
                destination 192.107.9.192;
            }
        }
    }
}
```

```
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
        }
    }
}
user@host# show interfaces | display set | match address
set interfaces ge-2/3/0 unit 0 family inet address 192.168.9.106/30
set interfaces so-5/1/0 unit 0 family inet address 192.168.9.15/32 destination
192.168.9.192
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

Displaying Users Currently Editing the Configuration

To display the users currently editing the configuration, use the status configuration mode command:

```
user@host# status
Current configuration users:
user terminal p0 (pid 518) on since 2000-03-12 18:24:27 PST
[edit protocols]
```

The system displays who is editing the configuration (user), from where the user is logged in (terminal p0), the date and time the user logged in (2000-03-12 18:24:27 PST), and what level of the hierarchy the user is editing ([edit protocols]).

If you issue the status configuration mode command and a user has scheduled a candidate configuration to become active for a future time, the system displays who scheduled the commit (root), where the user is logged in (terminal d0), the date and time the user logged in (2002-10-31 14:55:15 PST), and that a commit is pending (commit at).

```
[edit]
user@host# status
Users currently editing the configuration:
root terminal d0 (pid 767) on since 2002-10-31 14:55:15 PST, idle 00:03:09
commit at
```

For information about how to schedule a commit, see “Scheduling a Commit” on page 258.

If you issue the status configuration mode command and a user is editing the configuration in configure exclusive mode, the system displays who is editing the configuration (root), where the user is logged in (terminal d0), the date and time the user logged in (2002-11-01 13:05:11 PST), and that a user is editing the configuration in configure exclusive mode (exclusive [edit]).

```
[edit]
user@host# status
Users currently editing the configuration:
root terminal d0 (pid 2088) on since 2002-11-01 13:05:11 PST
exclusive [edit]
```

For more information about configure exclusive, see “Using the Configure Exclusive Command” on page 224.

Removing a Statement from the Configuration

To delete a statement or identifier, use the delete configuration mode command. Deleting a statement or an identifier effectively “unconfigures” the functionality associated with that statement or identifier, returning that functionality to its default condition.

```
user@host# delete <statement-path> <identifier>
```

When you delete a statement, the statement and all its subordinate statements and identifiers are removed from the configuration.

For statements that can have more than one identifier, when you delete one identifier, only that identifier is deleted. The other identifiers in the statement remain.

To delete the entire hierarchy starting at the current hierarchy level, do not specify a statement or an identifier in the delete command. When you omit the statement or identifier, you are prompted to confirm the deletion:

```
[edit]
user@host# delete
Delete everything under this level? [yes, no] (no) ?
    Possible completions:
      no     Don't delete everything under this level
      yes    Delete everything under this level
Delete everything under this level? [yes, no] (no)
```

Examples: Removing a Statement from the Configuration

Delete the ospf statement, effectively unconfiguring OSPF on the router:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/0 hello-interval 5
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/0 {
          hello-interval 5;
        }
      }
    }
}
[edit]
user@host# delete protocols ospf
[edit]
user@host# show
[edit]
user@host#
```

Delete all statements from the current level down:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@host# set interface so-0/0/0 hello-interval 5
[edit protocols ospf area 0.0.0.0]
user@host# delete
Delete everything under this level? [yes, no] (no) yes
[edit protocols ospf area 0.0.0.0]
user@host# show
[edit]
user@host#
```

Unconfigure a particular property:

```
[edit]
user@host# set interfaces so-3/0/0 speed 100mb
[edit]
user@host# show
interfaces {
    so-3/0/0 {
      speed 100mb;
    }
}
[edit]
user@host# delete interfaces so-3/0/0 speed
[edit]
user@host# show
interfaces {
    so-3/0/0;
}
```

For information how to use regular expressions to remove related configuration items, see “Using Regular Expressions to Remove Related Configuration Items” on page 248.

Using Regular Expressions to Remove Related Configuration Items

You can delete related configuration items simultaneously, such as channelized interfaces or static routes, by using a single command and regular expressions. Deleting a statement or an identifier effectively “unconfigures” the functionality associated with that statement or identifier, returning that functionality to its default condition.

You can only delete several parts of the configuration where you normally put multiple items; for example, interfaces. However, you cannot delete “groups” of different items; for example:

```
user@host# show system services
ftp;
rlogin;
rsh;
ssh {
    root-login allow;
}
telnet;

[edit]
user@host# wildcard delete system services *
                               ^
syntax error.
```

When you delete a statement, the statement and all its subordinate statements and identifiers are removed from the configuration.

To delete related configuration items, issue the wildcard configuration mode command with the delete option and specify the statement path, the items to be summarized with a regular expression, and the regular expression.

```
user@host# wildcard delete <statement-path> <identifier> <regular-expression>
```

NOTE: When you use the wildcard command to remove related configuration items, the regular expression must be the final statement.

If the JUNOS software matches more than eight related items, the CLI displays only the first eight items.

Example: Deleting Interfaces from the Configuration

Delete multiple T1 interfaces in the range from t1-0/0/0:0 through t1-0/0/0:23:

```
user@host# wildcard delete interfaces t1-1/3/0:.*
  matched: t1-1/3/0:0
  matched: t1-1/3/0:1
  matched: t1-1/3/0:2
Delete 3 objects? [yes,no] (no) no
```

Example: Deleting Routes from the Configuration

Delete static routes in the range from 172.0.0.0 to 172.255.0.0:

```
user@host# wildcard delete routing-options static route 172.*
  matched: 172.16.0.0/12
  matched: 172.16.14.0/24
  matched: 172.16.100.0/24
  matched: 172.16.128.0/19
  matched: 172.16.160.0/24
  matched: 172.17.12.0/23
  matched: 172.17.24.0/23
  matched: 172.17.28.0/23
  ...
Delete 13 objects? [yes,no] (no)
```

Copying a Statement in the Configuration

When you have many statements in a configuration that are similar, you can add one statement, then make copies of that statement. Copying a statement duplicates that statement and the entire hierarchy of statements configured under that statement. Copying statements is useful when you are configuring many physical or logical interfaces of the same type.

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

```
user@host# copy existing-statement to new-statement
```

Immediately after you have copied a portion of the configuration, the configuration might not be valid. You must check the validity of the new configuration, and if necessary, modify either the copied portion or the original portion for the configuration to be valid.

Example: Copying a Statement in the Configuration

After you have created one virtual connection (VC) on an interface, copy its configuration to create a second VC:

```
[edit interfaces]
user@host# show
at-1/0/0 {
    description "PAIX to MAE West"
    encapsulation atm-pvc;
      unit 61 {
        point-to-point;
        vci 0.61;
        family inet {
          address 10.0.1.1/24;
        }
      }
    }
}
[edit interfaces]
user@host# edit at-1/0/0
[edit interfaces at-1/0/0]
user@host# copy unit 61 to unit 62
[edit interfaces at-1/0/0]
user@host# show
description "PAIX to MAE West"
encapsulation atm-pvc;
    unit 61 {
      point-to-point;
      vci 0.61;
      family inet {
        address 10.0.1.1/24;
      }
    }
    unit 62 {
      point-to-point;
      vci 0.61;
      family inet {
        address 10.0.1.1/24;
      }
    }
}
```

Renaming an Identifier

When modifying a configuration, you can rename an identifier that is already in the configuration. You can do this either by deleting the identifier (using the delete command) and then adding the renamed identifier (using the set and edit commands), or you can rename the identifier using the rename configuration mode command:

```
user@host# rename <statement-path> identifier1 to identifier2
```

Example: Renaming an Identifier

Change the Network Time Protocol (NTP) server address to 10.0.0.6:

```
[edit]
user@host# rename system network-time server 10.0.0.7 to server 10.0.0.6
```

Inserting a New Identifier

When configuring the router, you can enter most statements and identifiers in any order. Regardless of the order in which you enter the configuration statements, the CLI always displays the configuration in a strict order. However, there are a few cases where the ordering of the statements matters because the configuration statements create a sequence that is analyzed in order.

For example, in a routing policy or firewall filter, you define terms that are analyzed sequentially. Also, when you create a named path in dynamic MPLS, you define an ordered list of the transit routers in the path, starting with the first transit router and ending with the last one.

To modify a portion of the configuration in which the statement order matters, use the insert configuration mode command:

```
user@host# insert <statement-path> identifier1 (before | after) identifier2
```

If you do not use the insert command, but instead simply configure the identifier, it is placed at the end of the list of similar identifiers.

Examples: Inserting a New Identifier

Insert policy terms in a routing policy configuration. Note that if you do not use the insert command, but rather just configure another term, the added term is placed at the end of the existing list of terms.

```
[edit]
user@host# show
policy-options {
    policy-statement statics {
      term term1 {
        from {
          route-filter 192.168.0.0/16 orlonger;
          route-filter 224.0.0.0/3 orlonger;
        }
        then reject;
      }
      term term2 {
        from protocol direct;
        then reject;
      }
      term term3 {
        from protocol static;
        then reject;
      }
      term term4 {
        then accept;
      }
    }
}
[edit]
user@host# rename policy-options policy-statement statics term term4 to term term6
[edit]
user@host# set policy-options policy-statement statics term term4 from protocol local
[edit]
user@host# set policy-options policy-statement statics term term4 then reject
[edit]
user@host# set policy-options policy-statement statics term term5 from protocol aggregate
[edit]
user@host# set policy-options policy-statement statics term term5 then reject
[edit]
user@host# insert policy-options policy-statement statics term term4 after term term3
[edit]
user@host# insert policy-options policy-statement statics term term5 after term term4
```

```
[edit]
user@host# show policy-options policy-statement statics
term term1 {
    from {
        route-filter 192.168.0.0/16 orlonger;
        route-filter 224.0.0.0/3 orlonger;
    }
    then reject;
}
term term2 {      # reject direct routes
    from protocol direct;
    then reject;
}
term term3 {        # reject static routes
    from protocol static;
    then accept;
}
term term4 {       #reject local routes
    from protocol local;
    then reject;
}
term term5 {      #reject aggregate routes
    from protocol aggregate;
    then reject;
}
term term6 {   #accept all other routes
    then accept;
}
```

Insert a transit router in a dynamic MPLS path:

```
[edit protocols mpls path ny-sf]
user@host# show
1.1.1.1;
2.2.2.2;
3.3.3.3 loose;
4.4.4.4 strict;
6.6.6.6;
[edit protocols mpls path ny-sf]
user@host# insert 5.5.5.5 before 6.6.6.6
[edit protocols mpls path ny-sf]
user@host# set 5.5.5.5 strict
[edit protocols mpls path ny-sf]
user@host# show
1.1.1.1;
2.2.2.2;
3.3.3.3 loose;
4.4.4.4 strict;
5.5.5.5 strict;
6.6.6.6;
```

Running an Operational Mode CLI Command from Configuration Mode

At times, you might need to display the output of an operational mode show or other command while configuring the software. While in configuration mode, you can execute a single operational mode command by issuing the configuration mode run command and specifying the operational mode command:

```
[edit]
user@host# run operational-mode-command
```

Example: Running an Operational Mode CLI Command from Configuration Mode

Display the priority value of the Virtual Router Redundancy Protocol (VRRP) master router while you are modifying the VRRP configuration for a backup router:

```
[edit interfaces ge-4/2/0 unit 0 family inet vrrp-group 27]
user@host# show
virtual-address [ 192.168.1.15 ];
[edit interfaces ge-4/2/0 unit 0 family inet vrrp-group 27]
user@host# run show vrrp detail
Physical interface: ge-5/2/0, Unit: 0, Address: 192.168.29.10/24
  Interface state: up, Group: 10, State: backup
  Priority: 190, Advertisement interval: 3, Authentication type: simple
  Preempt: yes, VIP count: 1, VIP: 192.168.29.55
  Dead timer: 8.326, Master priority: 201, Master router: 192.168.29.254
[edit interfaces ge-4/2/0 unit 0 family inet vrrp-group 27]
user@host# set priority ...
```

Displaying Configuration Mode Command History

In configuration mode, you can display a list of the recent commands you issued while in configuration mode. To do this, use the run show cli history command:

```
user@host> configure
...
[edit]
user@host# run show cli history
      12:40:08 – show
      12:40:17 – edit protocols
      12:40:27 – set isis
      12:40:29 – edit isis
      12:40:40 – run show cli history
[edit protocols isis]
user@host#
```

By default, this command displays the last 100 commands issued in the CLI. If you specify a number with the command, it displays that number of recent commands. For example:

```
user@host# run show cli history 3
      12:40:08 – show
      12:40:17 – edit protocols
      12:40:27 – set isis
```

Verifying a Configuration

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

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

If the commit check command finds an error, a message indicates the location of the error.

Committing a Configuration

To save software configuration changes to the configuration database and activate the configuration on the router, use the commit configuration mode command:

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

The configuration is checked for syntax errors. If the syntax is correct, the configuration is activated and becomes the current, operational router configuration.

You can issue the commit command from any hierarchy level.

These sections discuss how to commit configurations:

- Committing a Configuration and Exiting Configuration Mode on page 256
- Activating a Configuration but Requiring Confirmation on page 257
- Scheduling a Commit on page 258
- Monitoring the Commit Process on page 259
- Adding a Comment to Describe the Committed Configuration on page 260

If the configuration contains syntax errors, a message indicates the location of the error and the configuration is not activated. The error message has the following format:

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

For example:

```
[edit firewall filter login-allowed term allowed from]
    'icmp-type [ echo-request echo-reply ];'
      keyword 'echo-reply' unrecognized
```

You must correct the error before recommitting the configuration. To return quickly to the hierarchy level where the error is located, copy the path from the first line of the error and paste it at the configuration mode prompt at the [edit] hierarchy level.

When you commit a configuration, you commit the entire configuration in its current form. If more than one user is modifying the configuration, committing it saves and activates the changes of all the users.

After you commit the configuration and are satisfied that it is running successfully, you should issue the request system snapshot command to back up the new software onto the /altconfig file system. If you do not issue the request system snapshot command, the configuration on the alternate boot drive will be out of sync with the configuration on the primary boot drive.

The request system snapshot command backs up the root file system to /altroot, and /config to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue this command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Committing a Configuration and Exiting Configuration Mode

To save software configuration changes, activate the configuration on the router, and exit configuration mode, use the commit and-quit configuration mode command. This command succeeds only if the configuration contains no errors.

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

Activating a Configuration but Requiring Confirmation

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

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

```
[edit]
user@host# commit confirmed
commit confirmed will be automatically rolled back in 10 minutes unless confirmed
commit complete
[edit]
user@host#
```

To keep the new configuration active, enter a commit or commit check command within 10 minutes of the commit confirmed command. For example:

```
[edit]
user@host# commit confirmed
commit confirmed will be automatically rolled back in 10 minutes unless confirmed
commit complete
[edit]
user@host#
```

If the commit is not confirmed within a certain amount of time (10 minutes by default), the JUNOS software automatically rolls back to the previous configuration and a broadcast message is sent to all logged-in users.

To show when a rollback is scheduled after a commit confirmed command, enter the show system commit command. For example:

```
user@host# show commit confirmed
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

Like the commit command, the commit confirmed command verifies the configuration syntax and reports any errors. If there are no errors, the configuration is activated and begins running on the router.

Figure 6 on page 258 illustrates how the commit confirmed command works.

Figure 6: Confirm a Configuration

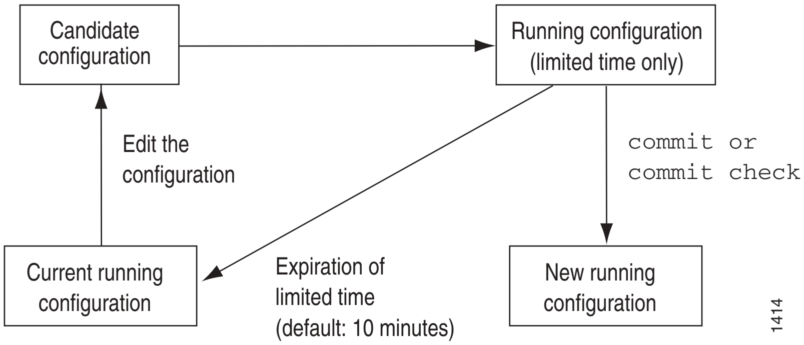
Candidate configuration
Running configuration (limited time only)
Edit the configuration
commit or commit check
Current running configuration
Expiration of limited time (default: 10 minutes)
New running configuration
1414

To change the amount of time before you have to confirm the new configuration, specify the number of minutes when you issue the command:

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

Scheduling a Commit

You can schedule when you want your candidate configuration to become active. To save software configuration changes and activate the configuration on the router at a future time or upon reboot, use the commit at configuration mode command, specifying reboot or a future time at the [edit] hierarchy level:

```
[edit]
user@host # commit at <string>
```

string is reboot or the future time to activate the configuration changes. You can specify time in two formats:

- A time value in the form *hh*:*mm* [:ss] (hours, minutes, and optionally seconds)—Commit the configuration at the specified time, which must be in the future but before 11:59:59 PM on the day the commit at configuration command is issued. Use 24-hour time for the *hh* value; for example, 04:30:00 is 4:30:00 AM, and 20:00 is 8:00 PM. The time is interpreted with respect to the clock and time zone settings on the router.

- A date and time value in the form *yyyy-mm-dd hh*:*mm* [:ss] (year, month, date, hours, minutes, and, optionally, seconds)—Commit the configuration at the specified day and time, which must be after the commit at command is issued. Use 24-hour time for the *hh* value. For example, 2003-08-21 12:30:00 is 12:30 PM on August 21, 2003. The time is interpreted with respect to the clock and time zone settings on the router.

  Enclose the *string* value in quotation marks (“). For example, commit at “18:00:00”. For date and time, include both values in the same set of quotation marks. For example, commit at "2005-03-10 14:00:00".

A “commit check” is performed when you issue the commit at configuration mode command. If the result of the check is successful, then the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit can be performed until the scheduled commit is completed.

NOTE: If the JUNOS software fails before the configuration changes become active, all configuration changes are lost.

You cannot issue the commit at configuration command after you issue the request system reboot command.

You cannot issue the request system reboot command once you schedule a commit operation for a specific time in the future.

You cannot commit a configuration when a scheduled commit is pending. For information about how to cancel a scheduled configuration by means of the clear command, see the *JUNOS System Basics and Services Command Reference*.

Monitoring the Commit Process

To monitor the commit process, use the display detail command after the pipe with the commit command:

```
user@host# commit | display detail
```

For example:

```
[edit]
user@host# commit | display detail
2003-09-22 15:39:39 PDT: exporting juniper.conf
2003-09-22 15:39:39 PDT: setup foreign files
2003-09-22 15:39:39 PDT: propagating foreign files
2003-09-22 15:39:39 PDT: complete foreign files
2003-09-22 15:39:40 PDT: copying configuration to juniper.data+
2003-09-22 15:39:40 PDT: dropping unchanged foreign files
2003-09-22 15:39:40 PDT: daemons checking new configuration
2003-09-22 15:39:41 PDT: commit wrapup...
2003-09-22 15:39:42 PDT: activating '/var/etc/ntp.conf'
2003-09-22 15:39:42 PDT: activating '/var/etc/kmd.conf'
2003-09-22 15:39:42 PDT: activating '/var/db/juniper.data'
2003-09-22 15:39:42 PDT: notifying daemons of new configuration
2003-09-22 15:39:42 PDT: signaling 'Firewall daemon', pid 24567, signal 1,
    status 0
2003-09-22 15:39:42 PDT: signaling 'Interface daemon', pid 24568, signal 1,
    status 0
2003-09-22 15:39:43 PDT: signaling 'Routing protocol daemon', pid 25679,
    signal 1, status 0
2003-09-22 15:39:43 PDT: signaling 'MIB2 daemon', pid 24549, signal 1, status
    0
2003-09-22 15:39:43 PDT: signaling 'NTP daemon', pid 37863, signal 1, status 0
2003-09-22 15:39:43 PDT: signaling 'Sonet APS daemon', pid 24551, signal 1,
    status 0
2003-09-22 15:39:43 PDT: signaling 'VRRP daemon', pid 24552, signal 1, status
    0
2003-09-22 15:39:43 PDT: signaling 'PFE daemon', pid 2316, signal 1, status 0
```

```
2003-09-22 15:39:43 PDT: signaling 'Traffic sampling control daemon', pid
24553, signal 1, status 0
2003-09-22 15:39:43 PDT: signaling 'IPSec Key Management daemon', pid
24556, signal 1, status 0
2003-09-22 15:39:43 PDT: signaling 'Forwarding UDP daemon', pid 2320, signal
    1, status 0
commit complete
```

Adding a Comment to Describe the Committed Configuration

You can include a comment that describes changes to the committed configuration. To add a comment that describes the changes to the committed configuration, include the commit comment statement. The comment can be as long as 512 bytes and you must type it on a single line.

```
[edit]
user@host # commit comment <comment-string>
```

comment-string is the text of the comment.

NOTE: You cannot include a comment with the commit check command.

To add a comment to the commit command, include the comment statement after the commit command:

```
[edit]
user@host# commit comment “add user joe”
commit complete
[edit]
user@host#
```

To add a comment to the commit confirmed command, include the comment statement after the commit confirmed command:

```
[edit]
user@host# commit confirmed comment “add customer to port 27”
commit confirmed will be automatically rolled back in 10 minutes unless confirmed
commit complete
[edit]
user@host#
```

To view these commit comments, issue the show system commit operational mode command.

Synchronizing Routing Engines

If your router has two Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the other by issuing the commit synchronize command. The Routing Engine on which you execute this command (requesting Routing Engine) copies and loads its candidate configuration to the other (responding Routing Engine). Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.

For example, if you are logged in to re1 (requesting Routing Engine) and you want re0 (responding Routing Engine) to have the same configuration as re1, issue the commit synchronize command on re1. re1 copies and loads its candidate configuration to re0. Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, re1's candidate configuration is activated and becomes the current operational configuration on both Routing Engines.

NOTE: When you issue the commit synchronize command, you must use the groups re0 and re1. For information about how to use the apply groups statement, see “Applying a Configuration Group” on page 619.

The responding Routing Engine must be running JUNOS release 5.0 or higher.

For information about issuing the commit synchronize command on a routing matrix, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

To synchronize a Routing Engine's current operational configuration file with the other, log in to the Routing Engine from which you want to synchronize and issue the commit synchronize command:

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

NOTE: You can also add the commit synchronize statement at the [edit system] hierarchy level so that a commit command automatically invokes a commit synchronize command by default. For more information, see “Configuring Multiple Routing Engines to Synchronize Configurations Automatically” on page 383.

Example: Using Apply Groups re0 and re1

The following example shows apply groups re0 and re1 with some configuration data that might be different on re0 and re1:

```
re0 {
    system {
        host-name my_router_RE0;
    }
    interfaces {
        fxp0 {
            unit 0 {
                family inet {
                    address 192.168.15.49/24;
                }
                family iso;
            }
        }
    }
}
re1 {
    system {
        host-name my_router_RE1;
    }
    interfaces {
        fxp0 {
            unit 0 {
                family inet {
                    address 192.168.15.50/24;
                }
                family iso;
            }
        }
    }
}
```

Example: Setting Apply Groups re0 and re1

The following example sets the apply groups re0 and re1:

```
[edit]
user@host# set apply-groups [re0 re1]
[edit]
user@host#
```

Saving a Configuration to a File

You might want to save the configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the save configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

By default, the configuration is saved to a file in your home directory, which is on the flash drive. For information about specifying the filename, see "Specifying Filenames and URLs" on page 360.

Loading a Configuration

You can create a file, copy the file to the local router, and then load the file in to the CLI. After you have loaded the file, you can commit it to activate the configuration on the router, or you can edit the configuration interactively using 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 load configuration mode command:

```
[edit]
user@host# load (merge | override | patch | replace | update) filename relative)
```

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

```
[edit]
user@host# load (merge | override | patch | replace | update) terminal relative)
[Type ^D to end input]
```

To replace an entire configuration, specify the override option at any level of the hierarchy.

An override operation discards the current candidate configuration and loads the configuration in *filename* or the one that you type at the terminal. When you use the override option and commit the configuration, all system processes reparse the configuration.

To replace only the configuration that has changed, specify the update option at any level of the hierarchy. An update operation compares the current configuration and the current candidate configuration, and loads only the changes between these configurations in *filename* or the one that you type at the terminal. When you use the update operation and commit the configuration, the JUNOS software attempts to notify the smallest set of system processes that are affected by the configuration change.

To combine the current configuration and the configuration in *filename* or the one that you type at the terminal, specify the merge option. 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.

To replace portions of a configuration, specify the replace option. For this operation to work, you must include replace: tags in the file or configuration you type at the terminal. The software searches for the replace: tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the replace operation adds to the configuration the statements marked with the replace: tag.

To use the merge, replace, or update option without specifying the full hierarchy level, specify the relative option. For example:

```
[edit system]
user@host# show static-host-mapping
bob sysid 987.654.321ab

[edit system]
user@host# load replace terminal relative
{Type ^D at a new line to end input]
replace: static-host-mapping {
    bob sysid 0123.456.789bc;
}
load complete

[edit system]
user@host# show static-host-mapping
bob sysid 0123.456.789bc;
```

To change part of the configuration with a patch file and mark only those parts as changed, specify the patch option.

If, in an override or merge operation, you specify a file or type text that contains replace: tags, the replace: tags are ignored, and the override or merge operation is performed.

If you are performing a replace operation and the file you specify or text you type does not contain any replace: tags, the replace operation is effectively equivalent to a merge operation. This might be useful 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.

For information about specifying the filename, see “Specifying Filenames and URLs” on page 360.

To copy a configuration file from another network system to the local router, you can use the SSH and telnet utilities, as described in the *JUNOS System Basics and Services Command Reference*.

Examples: Load a Configuration from a File

Figure 7: Example 1: Load a Configuration from a File

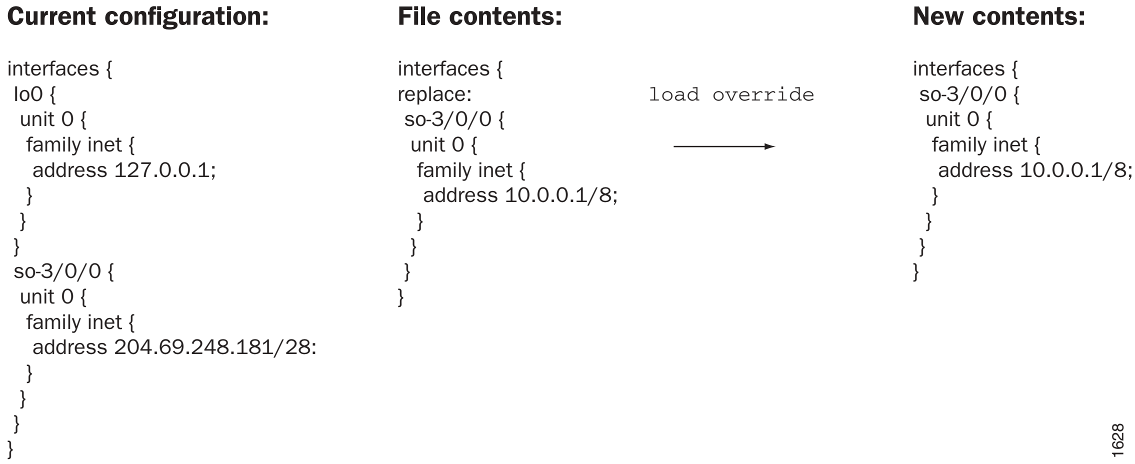

Current configuration:

```
interfaces {
 lo0 {
  unit 0 {
   family inet {
    address 127.0.0.1;
   }
  }
 }
 so-3/0/0 {
  unit 0 {
   family inet {
    address 204.69.248.181/28:
   }
  }
 }
}
```

File contents:

```
interfaces {
replace:
 so-3/0/0 {
  unit 0 {
   family inet {
    address 10.0.0.1/8;
   }
  }
 }
}
```

```
load override
```

New contents:

```
interfaces {
 so-3/0/0 {
  unit 0 {
   family inet {
    address 10.0.0.1/8;
   }
  }
 }
}
```

Figure 8: Example 2: Load a Configuration from a File

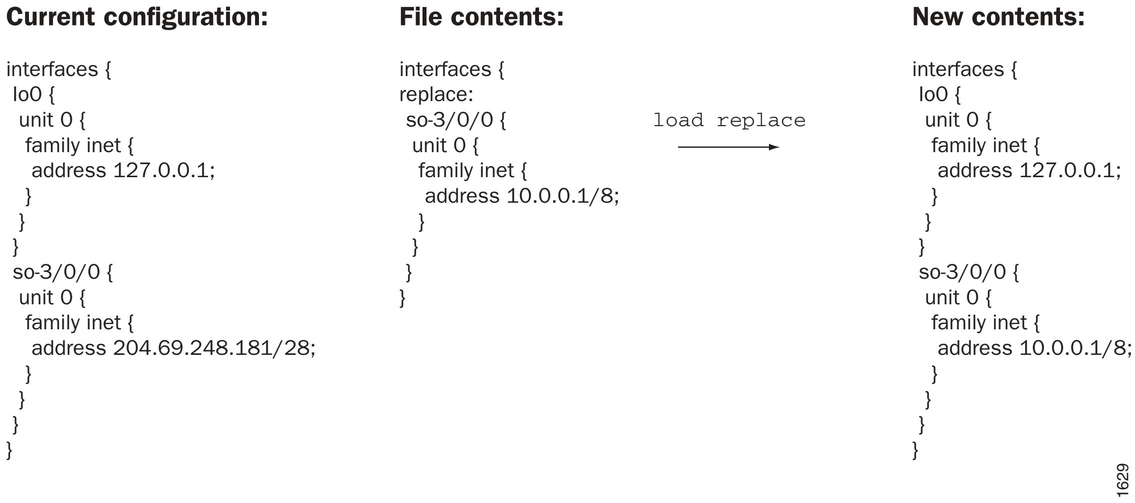

Current configuration:

```
interfaces {
 lo0 {
  unit 0 {
   family inet {
    address 127.0.0.1;
   }
  }
 }
 so-3/0/0 {
  unit 0 {
   family inet {
    address 204.69.248.181/28;
   }
  }
 }
}
```

File contents:

```
interfaces {
replace:
 so-3/0/0 {
  unit 0 {
   family inet {
    address 10.0.0.1/8;
   }
  }
 }
}
```

```
load replace
```

New contents:

```
interfaces {
 lo0 {
  unit 0 {
   family inet {
    address 127.0.0.1;
   }
  }
 }
 so-3/0/0 {
  unit 0 {
   family inet {
    address 10.0.0.1/8;
   }
  }
 }
}
```

Figure 9: Example 3: Load a Configuration from a File

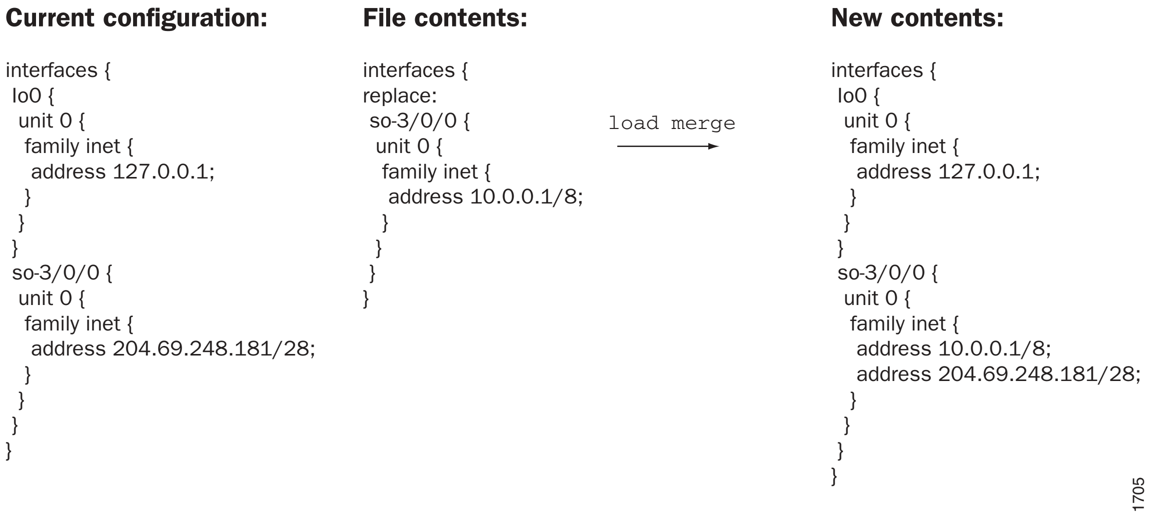

Current configuration:

```
interfaces {
 lo0 {
  unit 0 {
   family inet {
    address 127.0.0.1;
   }
  }
 }
 so-3/0/0 {
  unit 0 {
   family inet {
    address 204.69.248.181/28;
   }
  }
 }
}
```

File contents:

```
interfaces {
replace:
 so-3/0/0 {
  unit 0 {
   family inet {
    address 10.0.0.1/8;
   }
  }
 }
}
```

New contents:

```
interfaces {
 lo0 {
  unit 0 {
   family inet {
    address 127.0.0.1;
   }
  }
 }
 so-3/0/0 {
  unit 0 {
   family inet {
    address 10.0.0.1/8;
    address 204.69.248.181/28;
   }
  }
 }
}
```

Figure 10: Example 4: Load a Configuration from a File

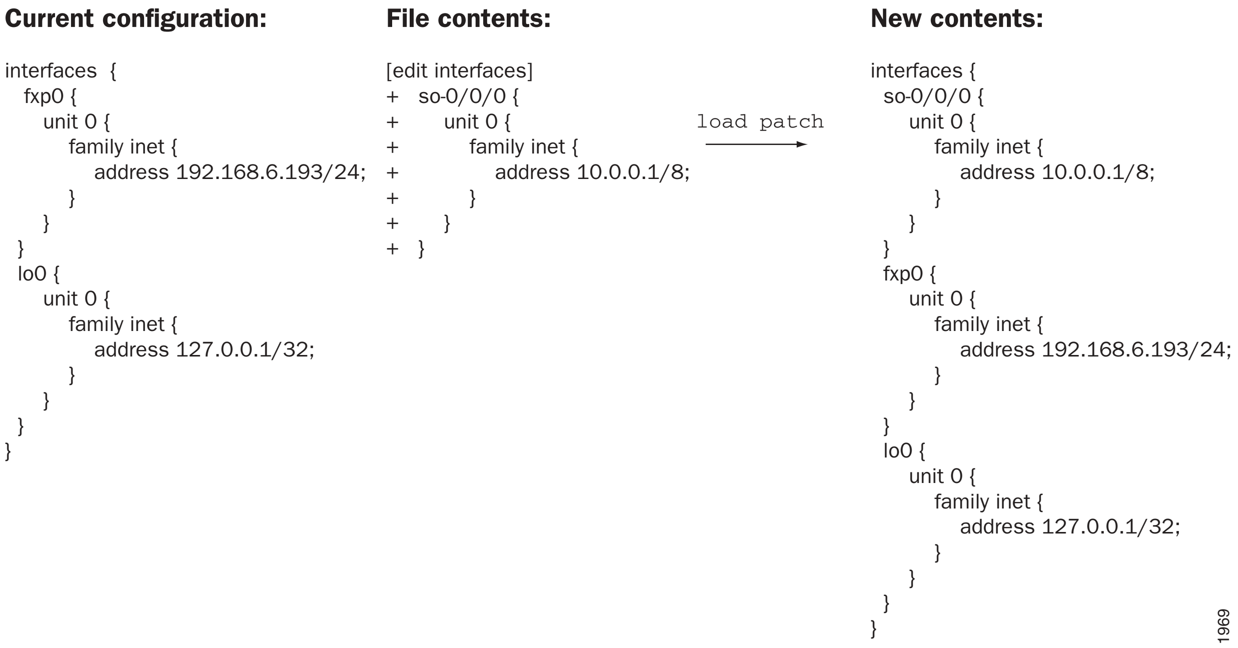

Current configuration:

```
interfaces  {
   fxp0 {
      unit 0 {
          family inet {
              address 192.168.6.193/24;
          }
      }
  }
  lo0 {
      unit 0 {
          family inet {
              address 127.0.0.1/32;
          }
      }
  }
}
```

File contents:

```
[edit interfaces]
+   so-0/0/0 {
+       unit 0 {
+           family inet {
+               address 10.0.0.1/8;
+           }
+       }
+   }
```

New contents:

```
interfaces {
  so-0/0/0 {
      unit 0 {
          family inet {
              address 10.0.0.1/8;
          }
      }
  }
  fxp0 {
      unit 0 {
          family inet {
              address 192.168.6.193/24;
          }
      }
  }
  lo0 {
      unit 0 {
          family inet {
              address 127.0.0.1/32;
          }
      }
  }
}
```

Returning to a Previously Committed Configuration

To return to the most recently committed configuration and load it into configuration mode without activating it, use the rollback configuration mode command:

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

To activate the configuration that you loaded, use the commit command:

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

To return to a configuration prior to the most recently committed one, include the number in the rollback command:

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

number can be a number in the range from 0 through 49. The most recently saved configuration is number 0 (which is the default configuration to which the system returns), and the oldest saved configuration is number 49.

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

```
[edit]
user@host# rollback ?
Possible completions:
<[Enter]> Execute this command
<number> Numeric argument
0 2001-02-27 12:52:10 PST by abc via cli
1 2001-02-26 14:47:42 PST by cde via cli
2 2001-02-14 21:55:45 PST by fgh via cli
3 2001-02-10 16:11:30 PST by hij via cli
4 2001-02-10 16:02:35 PST by klm via cli
| Pipe through a command
[edit]
```

For more information about configuration versions, see “How the Configuration Is Stored” on page 221.

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 “Configuring Access Privilege Levels” on page 402.

Example: Returning to a Previously Committed Version of the Configuration

Return to and activate the version stored in the file juniper.conf.3:

```
[edit]
user@host# rollback 3
load complete
[edit]
user@host# commit
```

Creating and Returning to a Rescue Configuration

If a user accidently commits a configuration that denies management access to the router, you must either connect a console to the router or invoke a *rescue* configuration. Using a rescue configuration is the recommended method.

A rescue configuration is one that you know allows management access to the router. You save a rescue configuration for recovery purposes — in case it is necessary to restore a valid, nondefault configuration. Once you have created a rescue configuration, you can return to it at any time.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the request system configuration rescue save command:

```
user@host> request system configuration rescue save
user@host>
```

To return to the rescue configuration, use the rollback rescue configuration mode command:

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

To activate the rescue configuration that you have loaded, use the commit command:

```
[edit]
user@host# rollback rescue
load complete
[edit]
user@host# commit
```

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time using either the rollback command, issue the request system configuration rescue save command:

```
user@host> request system configuration rescue save
user@host>
```

To delete an existing rescue configuration, issue the request system configuration rescue delete command:

```
user@host> request system configuration rescue delete
user@host>
```

For more information about the request system configuration rescue delete and request system configuration rescue save commands, see the *JUNOS System Basics and Services Command Reference*.

Configuration Mode Error Messages

If you do not type an option for a statement that requires one, a message indicates the type of information expected.

In this example, you need to type an area number to complete the command:

```
[edit]
user@host# set protocols ospf area<Enter>
                                 ^
syntax error, expecting <identifier>.
```

In this example, you need to type a value for the hello interval to complete the command:

```
[edit]
user@host# set protocols ospf area 45 interface so-0/0/0
            hello-interval<Enter>
                              ^
syntax error, expecting <data>
```

If you have omitted a required statement at a particular hierarchy level, when you attempt to move from that hierarchy level or when you issue the show command in configuration mode, a message indicates which statement is missing. For example:

```
[edit system login user phil]
user@host# up
Warning: missing mandatory statement: 'class'
[edit system login]
user@host# show
user phil {
      full-name “Phil James”;
        # Warning: missing mandatory statement(s): 'class'
}
```

Deactivating and Reactivating Statements and Identifiers in a Configuration

In a configuration, you can deactivate statements and identifiers so that they do not take effect when you issue the commit command. Any deactivated statements and identifiers are marked with the inactive: tag. They remain in the configuration, but are not activated when you issue a commit command.

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

```
deactivate (statement | identifier)
```

To reactivate a statement or identifier, use the activate configuration mode command:

```
activate (statement | identifier)
```

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

In some portions of the configuration hierarchy, you can include a disable statement to disable functionality. One example is disabling an interface by including the disable statement at the [edit interface *interface-name*] hierarchy level. When you deactivate a statement, that specific object or property is completely ignored and is not applied at all when you issue a commit command. When you disable a functionality, it is activated when you issue a commit command but is treated as though it is down or administratively disabled.

Examples: Deactivating and Reactivating Statements and Identifiers in a Configuration

Deactivate an interface in the configuration:

```
[edit interfaces]
user@host# show
at-5/2/0 {
    traceoptions {
      traceflag all;
    }
    atm-options {
      vpi 0 maximum-vcs 256;
    }
    unit 0 {
...
[edit interfaces]
user@host# deactivate at-5/2/0
[edit interfaces]
user@host# show
inactive: at-5/2/0 {
    traceoptions {
      traceflag all;
    }
...
```

Reactivate the interface:

```
[edit interfaces]
user@host# activate at-5/2/0
[edit interfaces]
user@host# show
at-5/2/0 {
    traceoptions {
        traceflag all;
    }
...
```

Adding Comments in a Configuration

You can include comments in a configuration to describe any statement in the configuration. You can add comments interactively in the CLI and by editing the ASCII configuration file.

When you add comments in configuration mode, they are associated with a statement at the current level. Each statement can have one single-line comment associated with it. Before you can associate a comment with a statement, the statement must exist. The comment is placed on the line preceding the statement.

To add comments to a configuration, use the annotate configuration mode command:

```
user@host# annotate statement "comment-string"
```

statement is the configuration statement to which you are attaching the comment; it must be at the current hierarchy level. If a comment for the specified *statement* already exists, it is deleted and replaced with the new comment.

comment-string is the text of the comment. The comment text can be any length, and you must type it on a single line. If the comment contains spaces, you must enclose it in quotation marks. In the comment string, you can include the comment delimiters /* */ or #. If you do not specify any, the comment string is enclosed with the /* */ comment delimiters.

To delete an existing comment, specify an empty comment string:

```
user@host# annotate statement ""
```

When you edit the ASCII configuration file and add comments, they can be one or more lines and must precede the statement they are associated with. If you place the comments in other places in the file, such as on the same line following a statement or on a separate line following a statement, they are removed when you use the load command to open the configuration into the CLI.

When you include comments in the configuration file directly, you can format comments in the following ways:

- Start the comment with a /* and end it with a */. The comment text can be on a single line or can span multiple lines.
- Start the comment with a # and end it with a new line (carriage return).

If you add comments with the annotate command, you can view the comments within the configuration by entering the show configuration mode command or the show configuration operational mode command.

When configuring interfaces, you can add comments about the interface by including the description statement at the [edit interfaces *interface-name*] hierarchy level. Any comments you include appear in the output of the show interfaces commands. For more information about the description statement, see the *JUNOS Network Interfaces Configuration Guide*.

Examples: Including Comments in Configurations

Add comments to a configuration:

```
[edit]
user@host# show
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/0 {
                hello-interval 5;
            }
        }
    }
}
[edit]
user@host# edit protocols ospf
[edit protocols ospf]
user@host# set area 0.0.0.0
user@host# annotate area 0.0.0.0 "Backbone area configuration added June 15,
1998"
[edit protocols ospf]
user@host# edit area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@host# annotate interface so0 "Interface from router sj1 to router sj2"
[edit protocols ospf area 0.0.0.0]
user@host# top
```

```
[edit]
user@host# show
protocols {
    ospf {
      /* Backbone area configuration added June 15, 1998 */
      area 0.0.0.0 {
        /* Interface from router sj1 to router sj2 */
        interface so-0/0/0 {
          hello-interval 5;
        }
      }
    }
}
[edit]
user@host#
```

The following excerpt from a configuration example illustrates how to enter comments in a configuration file:

```
/* This comment goes with routing-options */
routing-options {
    /* This comment goes with routing-options traceoptions */
    traceoptions {
      /* This comment goes with routing-options traceoptions tracefile */
      tracefile rpd size 1m files 10;
      /* This comment goes with routing-options traceoptions traceflag task */
      traceflag task;
      /* This comment goes with routing-options traceoptions traceflag general
*/
      traceflag general;
      }
    autonomous-system 10458;    /* This comment is dropped */
}
routing-options {
    rib-groups {
      ifrg {
        import-rib [ inet.0 inet.2 ];
        /* A comment here is dropped */
      }
      dvmrp-rib {
        import-rib inet.2;
        export-rib inet.2;
        /* A comment here is dropped */
      }
    /* A comment here is dropped */
    }
/* A comment here is dropped */
}
```

Having Multiple Users Configure the Software

Up to 32 users can be in configuration mode simultaneously, and they all can be making changes to the configuration. All changes made by all users are visible to everyone editing the configuration—the changes become visible as soon as the user presses the Enter key at the end of a command that changes the configuration, such as set, edit, or delete.

When any of the users editing the configuration issues a commit command, all changes made by all users are checked and activated.

Example: Using the CLI to Configure the Router

This section walks through an example of creating a simple configuration, illustrating how to use the CLI to create, display, and modify the software configuration for your system. The example used in this section creates the following configuration:

```
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/0 {
                hello-interval 5;
                dead-interval 20;
            }
            interface so-0/0/1 {
                hello-interval 5;
                dead-interval 20;
            }
        }
    }
}
```

Shortcut

You can create this entire configuration with two commands:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/0 hello-interval 5
dead-interval 20
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/1 hello-interval 5
dead-interval 20
```

Longer Configuration Example

The remainder of this section provides a longer example of creating the OSPF configuration. In the process, it illustrates how to use the different features of the CLI.

First, you enter configuration mode by issuing the configure top-level command:

```
user@host> configure
entering configuration mode
[edit]
user@host#
```

The prompt in braces shows that you are in configuration edit mode, at the top of the hierarchy. If you want to create the above configuration, you start by editing the protocols ospf statements:

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

Now, add the OSPF area:

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

Next, add the first interface:

```
[edit protocols ospf area 0.0.0.0]
user@host# edit interface so0
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host#
```

You now have four nested statements. Next, set the hello and dead intervals. Note that command completion (enter a tab or space) and context-sensitive help (type a question mark) are always available.

```
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# set ?
Possible completions:
+ apply-groups         Groups from which to inherit configuration data
> authentication-key   Authentication key
  dead-interval        Dead interval (seconds)
  disable              Disable OSPF on this interface
  hello-interval       Hello interval (seconds)
  interface-type       Type of interface
  metric               Interface metric (1..65535)
> neighbor             NBMA neighbor
  passive              Do not run OSPF, but advertise it
  poll-interval        Poll interval for NBMA interfaces
  priority             Designated router priority
  retransmit-interval  Retransmission interval (seconds)
  transit-delay        Transit delay (seconds)
  transmit-interval    OSPF packet transmit interval (milliseconds)
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# set hello-interval 5
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# set dead-interval 20
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host#
```

You can see what is configured at the current level with the show command:

```
[edit protocols ospf area 0.0.0.0 interface so-o]
user@host# show
hello-interval 5;
dead-interval 20;
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host#
```

You are finished at this level, so back up a level and take a look at what you have so far:

```
[edit protocols ospf area 0.0.0.0 interface so-0/0/0]
user@host# up
[edit protocols ospf area 0.0.0.0]
user@host# show
interface so-0/0/0 {
    hello-interval 5;
    dead-interval 20;
}
[edit protocols ospf area 0.0.0.0]
user@host#
```

The interface statement appears because you have moved to the area statement.

Now, add the second interface:

```
[edit protocols ospf area 0.0.0.0]
user@host# edit interface so-0/0/1
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# set hello-interval 5
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# set dead-interval 20
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# up
[edit protocols ospf area 0.0.0.0]
user@host# show
interface so-0/0/0 {
    hello-interval 5;
    dead-interval 20;
}
interface so-0/0/1 {
    hello-interval 5;
    dead-interval 20;
}
[edit protocols ospf area 0.0.0.0]
user@host#
```

Now, back up to the top level and see what you have:

```
[edit protocols ospf area 0.0.0.0]
user@host# top
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/0 {
          hello-interval 5;
          dead-interval 20;
        }
        interface so-0/0/1 {
          hello-interval 5;
          dead-interval 20;
        }
      }
    }
}
[edit]
user@host#
```

This configuration now contains the statements you want. Before committing it, which activates the configuration, verify that the configuration is correct:

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

Now you can commit the configuration to activate it on the router:

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

Suppose you decide to use different dead and hello intervals on interface so-0/0/1. You can make changes to the configuration. You can go directly to the appropriate hierarchy level by typing the full hierarchy path to the statement you want to edit.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0 interface so-0/0/1
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# show
hello-interval 5;
dead-interval 20;
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# set hello-interval 7
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# set dead-interval 28
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# top
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/0 {
          hello-interval 5;
          dead-interval 20;
        }
        interface so-0/0/1 {
          hello-interval 7;
          dead-interval 28;
        }
      }
    }
}
[edit]
user@host#
```

If you change your mind and decide not to run OSPF on the first interface, you can delete the statement:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@host# delete interface so-0/0/0
[edit protocols ospf area 0.0.0.0]
user@host# top
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/1 {
          hello-interval 7;
          dead-interval 28;
        }
      }
    }
}
[edit]
user@host#
```

Note that everything inside of the statement you deleted was deleted with it. You could eliminate the entire OSPF configuration by simply entering delete protocols ospf while at the top level.

Suppose you decide to use the default values for the hello and dead intervals on your remaining interface, but you want OSPF to run on that interface:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0 interface so-0/0/1
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# delete hello-interval
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# delete dead-interval
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# top
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/1;
      }
    }
}
[edit]
user@host#
```

You can set multiple statements at the same time as long as they are all part of the same hierarchy (the path of statements from the top inward, as well as one or more statements at the bottom of the hierarchy). Doing this can reduce considerably the number of commands you must enter. For example, if you want to go back to the original hello and dead interval timers on interface so-0/0/1, you can enter:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0 interface so-0/0/1
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# set hello-interval 5 dead-interval 20
[edit protocols ospf area 0.0.0.0 interface so-0/0/1]
user@host# exit
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/1 {
          hello-interval 5;
          dead-interval 20;
        }
      }
    }
}
[edit]
user@host#
```

You also can recreate the other interface, as you had it before, with only a single entry:

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface so-0/0/1 hello-interval 5
dead-interval 20
[edit]
user@host# show
protocols {
    ospf {
      area 0.0.0.0 {
        interface so-0/0/0 {
          hello-interval 5;
          dead-interval 20;
        }
        interface so-0/0/1 {
          hello-interval 5;
          dead-interval 20;
        }
      }
    }
}
```

Additional Details About Specifying Statements and Identifiers

This section provides more detailed information about specifying statements and identifiers in configuration mode:

- Specifying Statements on page 281
- Performing CLI Type-Checking on page 283

Specifying Statements

This section provides more detailed information about CLI container and leaf statements so that you can better understand how the CLI displays them in a configuration and how you must specify them when creating ASCII configuration files.

Statements are shown one of two ways, either with braces or without:

- Statement name and identifier, with one or more lower-level statements enclosed in braces:

  ```
  < statement-name > < identifier > {
      statement;
      additional-statements;
  }
  ```

- Statement name, identifier, and a single identifier:

  ```
  < statement-name > < identifier > identifier;
  ```

The *statement-name* is the name of the statement. In the configuration example shown in the previous section, ospf and area are statement names.

The *identifier* is a name or other string that uniquely identifies an instance of a statement. The identifier is used when a statement can be specified more than once in a configuration. In the configuration example shown in the previous section, the identifier for the area statement is 0 and the identifier for the interface statement is so-0/0/0.

When specifying a statement, you must specify either a statement name or an identifier, or both, depending on the statement hierarchy.

You specify identifiers in one of the following ways:

- *identifier*—The *identifier* is a flag, which is a single keyword.
- *identifier value*—The *identifier* is a keyword, and the *value* is a required option variable.
- *identifier* [*value1 value 2 value3* ...]—The *identifier* is a set that accepts multiple values. The brackets are required when you specify a set of identifiers; however, they are optional when you specify only one identifier.

The following examples illustrate how statements and identifiers are specified in the configuration:

```
protocol {                    # Top-level statement (statement-name).
    ospf {                    # Statement under "protocol" (statement-name).
      area 0.0.0.0 {          # OSPF area "0.0.0.0" (statement-name identifier),
        interface so-0/0/0 {# which contains an interface named "so-0/0/0."
          hello-interval 25;# Identifier and value (identifier-name value).
          priority 2;         # Identifier and value (identifier-name value).
          disable;            # Flag identifier (identifier-name).
        }
        interface so-0/0/1;# Another instance of "interface," named so-0/0/1,
      }                       # this instance contains no data, so no braces
    }                         # are displayed.
}
policy-options {              # Top-level statement (statement-name).
    term term1 {              # Statement under "policy-options"
                              # (statement-name value).
      from {                  # Statement under "term" (statement-name).
        route-filter 10.0.0.0/8 orlonger reject;# One identifier ("route-filter") with
        route-filter 127.0.0.0/8 orlonger reject;# multiple values.
        route-filter 128.0.0.0/16 orlonger reject;
        route-filter 149.20.64.0/24 orlonger reject;
        route-filter 172.16.0.0/12 orlonger reject;
        route-filter 191.255.0.0/16 orlonger reject;
      }
      then {                  # Statement under "term" (statement-name).
        next term;            # Identifier (identifier-name).
      }
    }
}
```

When you create an ASCII configuration file, you can specify statements and identifiers in one of the following ways. However, each statement has a preferred style, and the CLI uses that style when displaying the configuration in response to a configuration mode show command.

- Statement followed by identifiers:

```
    statement-name identifier-name [...] identifier-name value [...];
```

- Statement followed by identifiers enclosed in braces:

```
    statement-name {
        identifier-name;
        [...]
        identifier-name value;
        [...]
    }
```

- For some repeating identifiers, you can use one set of braces for all the statements:

```
statement-name {
    identifier-name value1;
    identifier-name value2;
}
```

Performing CLI Type-Checking

When you specify identifiers and values, the CLI expects to receive specific types of input and performs type-checking to verify that the data you entered is in the correct format. For example, for a statement in which you must specify an IP address, the CLI checks that you entered an address in a valid format. If you have not, an error message indicates what you were expected to type. Table 12 lists the data types the CLI checks.

Table 12: CLI Configuration Input Types

Data Type	Format	Examples
Physical interface name (used in the edit interfaces hierarchy)	*type-fpc/pic/port*	**Correct:** so-0/0/1 **Incorrect:** so-0
Full interface name	*type-fpc/pic/port<:channel>.logical*	**Correct:** so-0/0/1.0 **Incorrect:** so-0/0/1
Full or abbreviated interface name (used in places other than the edit interfaces hierarchy)	*type-<fpc</pic/port>><<:channel>.logical>*	**Correct:** so, so-1, so-1/2/3:4.5
IP address	0x*hex-bytes* *octet<.octet<.octet.<octet>>>*	**Correct:** 1.2.3.4, 0x01020304, 128.8.1, 128.8 **Sample translations:** 1.2.3 becomes 1.2.3.0 0x01020304 becomes 1.2.3.4 0x010203 becomes 0.1.2.3
IP address (destination prefix) and prefix length	0x*hex-bytes</length>* *octet<.octet<.octet.<octet>>></length>*	**Correct:** 10/8, 128.8/16, 1.2.3.4/32, 1.2.3.4 **Sample translations:** 1.2.3 becomes 1.2.3.0/32 0x01020304 becomes 1.2.3.4/32 0x010203 becomes 0.1.2.3/32 default becomes 0.0.0.0/0

Data Type	Format	Examples
International Organization for Standardization (ISO) address	*hex-nibble<hex-nibble ...>*	**Correct:** 47.1234.2345.3456.00, 47123423453456.00, 47.12.34.23.45.34.56.00 **Sample translations:** 47123456 becomes 47.1234.56 47.12.34.56 becomes 47.1234.56 4712.3456 becomes 47.1234.56
OSPF area identifier (ID)	0x*hex-bytes* *octet<.octet<.octet.<octet>>>* d*ecimal-number*	**Correct:** 54, 0.0.0.54, 0x01020304, 1.2.3.4 **Sample translations:** 54 becomes 0.0.0.54 257 becomes 0.0.1.1 128.8 becomes 128.8.0.0 0x010203 becomes 0.1.2.3

Chapter 9

Summary of CLI Environment Commands

The following sections explain each of the command-line interface (CLI) environment commands. The commands are organized alphabetically.

set cli complete-on-space

Syntax

```
set cli complete-on-space (off | on);
```

Release Information Command introduced before JUNOS Release 7.4.

Description Configure the keys to use for command completion.

Default When you type a space or tab, the CLI performs command completion.

Options off—Allow only a tab to be used for command completion.

on—Allow either a space or a tab to be used for command completion.

Sample Output

```
user@host> set cli com<Space>
user@host> set cli complete-on-space off
user@host> set cli com<Tab>
user@host> set cli complete-on-space on
user@host>
```

Usage Guidelines See “Setting Command Completion” on page 215.

Required Privilege Level view

set cli directory

Syntax set cli directory *directory*

Release Information Command introduced before JUNOS Release 7.4.

Description Set the current working directory.

Options *directory* is the pathname of working directory.

Usage Guidelines See “Setting the CLI Directory” on page 214.

Required Privilege Level view

set cli idle-timeout

Syntax set cli idle-timeout <*minutes*>

Release Information Command introduced before JUNOS Release 7.4.

Description Set the maximum time that an individual session can be idle before the user is logged off the router. The session times out after remaining at the CLI operational mode prompt for the specified time. The session can time out while monitoring log files.

Default If you do not issue this command, and the user’s login class does not specify this value, the user is never forced off the system after extended idle times.

Options *minutes*—Maximum idle time
Range: 0 through 100,000 minutes. Setting it to 0 disables the timeout.

Usage Guidelines See “Setting the Idle Timeout” on page 215.

Required Privilege Level view

See Also idle-timeout on page 557

set cli prompt

Syntax set cli prompt *string*

Release Information Command introduced before JUNOS Release 7.4.

Description Set the prompt to display within the CLI.

Default user@host>

Options *string*—CLI prompt. To include spaces in the prompt, enclose the string in quotation marks.

Sample Output

```
user@host> set cli prompt "cli% "
cli%
```

Usage Guidelines See “Setting the CLI Prompt” on page 214.

Required Privilege Level view

set cli restart-on-upgrade

Syntax set cli restart-on-upgrade (off | on)

Release Information Command introduced before JUNOS Release 7.4.

Description For an individual session, set the CLI to prompt you to restart the router after upgrading the software.

Default The CLI prompts you to restart, unless the screen length has been set to 0.

Options off—Disables the prompt.

on—Enables the prompt.

Usage Guidelines See “Setting the CLI to Prompt after a Software Upgrade” on page 215.

Required Privilege Level view

set cli screen-length

Syntax set cli screen-length *lines*

Release Information Command introduced before JUNOS Release 7.4.

Description Set the number of lines of text that the screen can display.

Options *lines*—Number of lines on the screen.

Range: 0 through 100,000 columns
Default: 24 lines

Sample Output

```
user@host> set cli screen-length 66
Screen length is set to 66
user@host>
```

Usage Guidelines See “Setting the Screen Length” on page 214.

Required Privilege Level view

set cli screen-width

Syntax set cli screen-width *width*

Release Information Command introduced before JUNOS Release 7.4.

Description Set the number of characters that the screen can display on a single line.

Options *width*—Number of columns on the screen.

Range: 0 through 100,000 columns
Default: 80 columns

Sample Output

```
user@host> set cli screen-width 40
Screen width set to 40
user@host>
```

Usage Guidelines See “Setting the CLI Prompt” on page 214.

Required Privilege Level view

set cli terminal

Syntax set cli terminal *terminal-type*

Release Information Command introduced before JUNOS Release 7.4.

Description Set the terminal type.

Options *terminal-type*—Type of terminal that is connected to the port.
Values: ansi, vt100, small-xterm, xterm
Default: The terminal type is unknown.

Usage Guidelines See "Setting the Terminal Type" on page 214.

Required Privilege Level view

set cli timestamp

Syntax set cli timestamp [format *timestamp-format* | disable]

Release Information Command introduced before JUNOS Release 7.4.

Description Set a timestamp for CLI output. By default, no timestamp appears in CLI output. If you issue this command, a date and time appear at the top of the command output in the format you specify.

Default If you enter the command without specifying a timestamp format, the default timestamp format is *Mmm dd hh:mm:ss* (for example, Feb 08 17:25:44).

Options format *timestamp-format*—Sets the data and time format for the timestamp. The timestamp format you specify can include the following placeholders in any order:

- %m—Two-digit month
- %d—Two-digit date
- %T—Six-digit hour, minute, and seconds

Enclose the format in single quotation marks ('). For example, set cli timestamp format '%m%d' results in a month, date timestamp such as Jun 30.

disable—Removes the timestamp from the CLI.

Usage Guidelines See "Setting the CLI Timestamp" on page 215.

Required Privilege Level view

set date

Syntax set date *YYYYMMDDhhmm.ss*

Release Information Command introduced before JUNOS Release 7.4.

Description Set the current date and time on the router.

Options *YYYYMMDDhhmm.ss*—Date and time to set. *YYYY* is the four-digit year, *MM* is the two-digit month, *DD* is the two-digit date, *hh* is the two-digit hour, *mm* is the two-digit minute, and ss is the two-digit second. At a minimum, you must specify the two-digit minute. All other parts of the date and time are optional.

Usage Guidelines See “Setting the Current Date and Time” on page 200.

Required Privilege Level view

See Also ntp on page 570, time-zone on page 602

set date ntp

Syntax set date ntp <*ntp-server*>

Release Information Command introduced before JUNOS Release 7.4.

Description Use a Network TIme Protocol (NTP) server to synchronize the current date and time setting on the router.

You do not need to reboot the router when you use the set date ntp command.

Options none—Uses the system NTP server list.

ntp-server—IP address of one or more NTP servers to query. When querying more than one server, the IP addresses are enclosed in quotation marks using the format
“*ip-address ip-address*”, for example, "200.49.40.1 129.127.28.4".

Usage Guidelines See “Setting the Date and Time from NTP Servers” on page 200.

Required Privilege Level view

Sample Output: set date ntp (to query one server)

```
user@host> set date ntp 172.17.12.9
14 Sep 22:00:50 ntpdate[20603]: step time server 172.17.12.9 offset 0.000461 sec
```

Sample Output: set date ntp (to query two servers)

```
user@host> set date ntp "200.49.40.1 129.127.28.4"
10 Feb 13:50:21 ntpdate[794]: step time server 129.127.28.4 offset 0.000163 sec
```

set date ntp source-address

Syntax set date ntp source-address *source-address*

Release Information Command introduced before JUNOS Release 7.4.

Description Specify a source address that is used by the JUNOS software to contact the remote NTP sever.

Options *source-address*—A valid IP address configured on one of the router interfaces.

Usage Guidelines See “Setting the Source Address to Contact the NTP Server” on page 201.

Required Privilege Level view

See Also source-address on page 594

show cli

Syntax show cli

Release Information Command introduced before JUNOS Release 7.4.

Description Display information about how the CLI environment is configured.

Sample Output

```
user@host> show cli
CLI screen length set to 60
CLI screen width set to 80
CLI complete-on-space set to on
user@host>
```

Usage Guidelines See “Displaying CLI Settings” on page 216.

Required Privilege Level view

show cli history

Syntax show cli history <*count*>

Release Information Command introduced before JUNOS Release 7.4.

Description List recent commands that you issued in the CLI and the time they were issued.

If you issue the run show cli history command from configuration mode, the command lists the most recent configuration mode commands that you issued and the time they were issued.

Options *count*—(Optional) Number of recent commands to display.

Range: 0 through 65,535

Default: 100

Sample Output

```
user@host> show cli history
  12:33:39 – configure
  12:42:52 – show cli history
  12:43:02 – show interfaces terse
  12:43:14 – show interfaces lo0
  12:43:20 – show bgp
  12:43:28 – show bgp next-hop-database
  12:43:32 – show cli history
user@host> configure
...
[edit]
user@host# run show cli history
  12:40:08 – show
  12:40:17 – edit protocols
  12:40:27 – set isis
  12:40:29 – edit isis
  12:40:40 – run show cli history
[edit protocols isis]
user@host#
```

Usage Guidelines See “Displaying CLI Command History” on page 201 and “Displaying Configuration Mode Command History” on page 254.

Required Privilege Level view

Chapter 10

Summary of CLI Configuration Mode Commands

The following sections explain each of the command-line interface (CLI) configuration mode commands. The commands are organized alphabetically.

activate

Syntax activate (*statement* | *identifier*)

Release Information Command introduced before JUNOS Release 7.4.

Description Remove the inactive: tag from a statement, effectively adding the statement or identifier back to the configuration. Statements or identifiers that have been activated take effect when you next issue the commit command.

Options *identifier*—Identifier from which you are removing the inactive tag. It must be an identifier at the current hierarchy level.

statement—Statement from which you are removing the inactive tag. It must be a statement at the current hierarchy level.

Usage Guidelines See “Deactivating and Reactivating Statements and Identifiers in a Configuration” on page 270.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also deactivate on page 297

annotate

Syntax annotate *statement* "*comment-string*"

Release Information Command introduced before JUNOS Release 7.4.

Description Add comments to a configuration. You can add comments only at the current hierarchy level.

Any comments you add appear only when you view the configuration by entering the show command in configuration mode or the show configuration command in operational mode.

Options *comment-string*—Text of the comment. You must enclose it in quotation marks. In the comment string, you can include the comment delimiters /* */ or #. If you do not specify any, the comment string is enclosed with the /* */ comment delimiters. If a comment for the specified *statement* already exists, it is deleted and replaced with the new comment.

statement—Statement to which you are attaching the comment.

Usage Guidelines See "Adding Comments in a Configuration" on page 271.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also See the description statement in the *JUNOS Network Interfaces Configuration Guide*.

commit

Syntax commit <<*at* <"string">> <and-quit> <check> <comment <"*comment-string*">> <confirmed> <display detail> <*minutes*> <synchronize>

Release Information Command introduced before JUNOS Release 7.4.

Description Commit the set of changes to the database and cause the changes to take operational effect.

Options at <"string">—(Optional) Save software configuration changes and activate the configuration at a future time, or upon reboot.

string is reboot or the future time to activate the configuration changes. Enclose the *string* value (including reboot) in quotation marks ("). You can specify time in two formats:

- A time value in the form *hh*:*mm* [:ss] (hours, minutes, and optionally seconds)— Commit the configuration at the specified time, which must be in the future but before 11:59:59 PM on the day the commit at configuration command is issued. Use 24-hour time for the *hh* value; for example, 04:30:00 is 4:30:00 AM, and 20:00 is 8:00 PM. The time is interpreted with respect to the clock and time zone settings on the router.

- A date and time value in the form *yyyy-mm-dd hh*:*mm* [:ss] (year, month, date, hours, minutes, and, optionally, seconds)—Commit the configuration at the specified day and time, which must be after the commit at command is issued. Use 24-hour time for the *hh* value. For example, 2003-08-21 12:30:00 is 12:30 PM on August 21, 2003. The time is interpreted with respect to the clock and time zone settings on the router.

For example, commit at "18:00:00". For date and time, include both values in the same set of quotation marks. For example, commit at "2005-03-10 14:00:00".

A *commit check* is performed when you issue the commit at configuration mode command. If the result of the check is successful, then the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit can be performed until the scheduled commit is completed.

NOTE: If the JUNOS software fails before the configuration changes become active, all configuration changes are lost.

You cannot issue the commit at configuration command when there is a pending reboot.

You cannot issue the request system reboot command once you schedule a commit operation for a specific time in the future.

You cannot commit a configuration when a scheduled commit is pending. For information about how to use the clear command to cancel a scheduled configuration, see the *JUNOS System Basics and Services Command Reference*.

and-quit—(Optional) Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

check—(Optional) Verify the syntax of the configuration, but do not activate it.

comment < "*comment-string*" > —(Optional) Add a comment that describes the committed configuration. The comment can be as long as 512 bytes and must be typed on a single line. You cannot include a comment with the commit check command. Enclose *comment-string* in quotation marks ("). For example, commit comment "Includes changes recommended by SW Lab".

confirmed <*minutes*>—(Optional) Require that the commit be confirmed within the specified amount of time. To confirm a commit, enter either a commit or commit check command. If the commit is not confirmed within the time limit, the configuration rolls back automatically to the precommit configuration and a broadcast message is sent to all logged-in users. To show when a rollback is scheduled, enter the show system commit command.
Range: 1 through 65,535 minutes
Default: 10 minutes

display detail—(Optional) Monitors the commit process.

synchronize—(Optional) If your router has two Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the other by issuing the commit synchronize command. The Routing Engine on which you execute this command (request Routing Engine) copies and loads its candidate configuration to the other (responding Routing Engine). Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.

NOTE: When you issue the commit synchronize command, you must use the apply-groups re0 and re1 commands. For information about how to use groups, see “Applying a Configuration Group” on page 619.

The responding Routing Engine must use JUNOS Release 5.0 or later.

Usage Guidelines See “Verifying a Configuration” on page 255, “Committing a Configuration” on page 255, “Scheduling a Commit” on page 258, “Synchronizing Routing Engines” on page 261, “Monitoring the Commit Process” on page 259, and “Adding a Comment to Describe the Committed Configuration” on page 260.

Required Privilege Level configure—To enter configuration mode.

copy

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

Release Information Command introduced before JUNOS Release 7.4.

Description Make a copy of an existing statement in the configuration.

Options *existing-statement*—Statement to copy.

new-statement—Copy of the statement.

Usage Guidelines See “Copying a Statement in the Configuration” on page 249.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

deactivate

Syntax deactivate (*statement* | *identifier*)

Release Information Command introduced before JUNOS Release 7.4.

Description Add the inactive: tag to a statement, effectively commenting out the statement or identifier from the configuration. Statements or identifiers marked as inactive do not take effect when you issue the commit command.

Options *identifier*—Identifier to which you are adding the inactive: tag. It must be an identifier at the current hierarchy level.

statement—Statement to which you are adding the inactive: tag. It must be a statement at the current hierarchy level.

Usage Guidelines See “Deactivating and Reactivating Statements and Identifiers in a Configuration” on page 270.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also activate on page 293, delete on page 297

delete

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

Release Information Command introduced before JUNOS Release 7.4.

Description Delete a statement or identifier. All subordinate statements and identifiers contained within the specified statement path are deleted with it.

Deleting a statement or an identifier effectively “unconfigures” or disables the functionality associated with that statement or identifier.

If you do not specify *statement-path* or *identifier*, the entire hierarchy starting at the current hierarchy level is removed.

Options *statement-path*—(Optional) Path to an existing statement or identifier. Include this if the statement or identifier to be deleted is not at the current hierarchy level.

identifier—(Optional) Name of the statement or identifier to delete.

Usage Guidelines See “Removing a Statement from the Configuration” on page 246.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also deactivate on page 297

edit

Syntax edit *statement-path*

Release Information Command introduced before JUNOS Release 7.4.

Description Move inside the specified statement hierarchy. If the statement does not exist, it is created.

You cannot use the edit command to change the value of identifiers. You must use the set command.

Options *statement-path*—Path to the statement.

Usage Guidelines See "Creating and Modifying the Configuration" on page 235.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also set on page 304

exit

Syntax exit <configuration-mode>

Release Information Command introduced before JUNOS Release 7.4.

Description Exit the current level of the statement hierarchy, returning to the level prior to the last edit command, or exit from configuration mode. The quit and exit commands are synonyms.

Options none—Return to the previous edit level. If you are at the top of the statement hierarchy, exit configuration mode.

configuration-mode—(Optional) Exit from configuration mode.

Usage Guidelines See "Moving Among Levels of the Hierarchy" on page 238.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also top on page 308, up on page 308

help

Syntax help < (apropos | topic | reference) <*string*> >

Release Information Command introduced before JUNOS Release 7.4.

Description Display help about available configuration statements.

Options apropos—Display all hierarchy levels containing the statement.

reference—Display summary information for the statement.

string—String or regular expression matching configuration statements for which you need help.

topic—(Optional) Display usage guidelines for the statement.

Entering the help command without an option provides introductory information on how to use the help command.

Usage Guidelines See “Getting Help Based on a String in a Statement Name” on page 233.

Required Privilege Level configure—To enter configuration mode.

insert

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

Release Information Command introduced before JUNOS Release 7.4.

Description Insert an identifier in to an existing hierarchy.

Options after—Place *identifier1* after *identifier2*.

before—Place *identifier1* before *identifier2*.

identifier1—Existing identifier.

identifier2—New identifier to insert.

statement-path—(Optional) Path to the existing identifier.

Usage Guidelines See “Inserting a New Identifier” on page 251.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

load

Syntax load (patch | merge | override | replace | update) (*filename* | terminal) relative

Release Information Command introduced before JUNOS Release 7.4.

Description Load a configuration from an ASCII configuration file or from terminal input. Your current location in the configuration hierarchy is ignored when the load operation occurs.

Options *filename*—Name of the file to load. For information about specifying the filename, see "Specifying Filenames and URLs" on page 360.

merge—Combine the configuration that is currently shown in the CLI and the configuration in *filename*.

override—Discard the entire configuration that is currently shown in the CLI and load the entire configuration in *filename*.

patch—Change part of the configuration and mark only those parts as changed.

replace—Look for a replace: tag in *filename*, delete the existing statement of the same name, and replace it with the configuration in *filename*.

relative—Use the merge or replace option without specifying the full hierarchy level.

terminal—Use the text you type at the terminal as input to the configuration. Type Ctrl-D to end terminal input.

update—Replace only the configuration that has changed in the *filename.*

Usage Guidelines See "Loading a Configuration" on page 263.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

quit

Syntax quit <configuration-mode>

Release Information Command introduced before JUNOS Release 7.4.

Description Exit the current level of the statement hierarchy, returning to the level prior to the last edit command, or exit from configuration mode. The quit and exit commands are synonyms.

Options none—Return to the previous edit level. If you are at the top of the statement hierarchy, exit configuration mode.

configuration-mode—(Optional) Exit from configuration mode.

Usage Guidelines See “Moving Among Levels of the Hierarchy” on page 238.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also top on page 308, up on page 308

rename

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

Release Information Command introduced before JUNOS Release 7.4.

Description Rename an existing configuration statement or identifier.

Options *identifier1*—Existing identifier to rename.

identifier2—New name of identifier.

statement-path—(Optional) Path to an existing statement or identifier.

Usage Guidelines See “Renaming an Identifier” on page 251.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

rollback

Syntax rollback <*number*>

Release Information Command introduced before JUNOS Release 7.4.

Description Return to a previously committed configuration. The software saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the commit configuration command.

The currently operational JUNOS software configuration is stored in the file juniper.conf, and the last three committed configurations are stored in the files juniper.conf.1, juniper.conf.2, and juniper.conf.3. These four files are located in the directory /config, which is on the router's flash drive. The remaining 46 previous committed configurations, the files juniper.conf.4 through juniper.conf.49, are stored in the directory /var/db/config, which is on the router's hard disk.

Options none—Return to the most recently saved configuration.

number—Configuration to return to.
Range: 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
Default: 0

rescue—Return to the rescue configuration.

Usage Guidelines See "Returning to a Previously Committed Configuration" on page 267 and "Creating and Returning to a Rescue Configuration" on page 268

Required Privilege Level rollback—To roll back to configurations other than the one most recently committed.

run

Syntax run *command*

Release Information Command introduced before JUNOS Release 7.4.

Description Run a top-level CLI command without exiting from configuration mode.

Options *command*—CLI top-level command.

Usage Guidelines See "Running an Operational Mode CLI Command from Configuration Mode" on page 254.

Required Privilege Level configure—To enter configuration mode.

save

Syntax save *filename*

Release Information Command introduced before JUNOS Release 7.4.

Description Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

When saving a file to a remote system, the software uses the scp/ssh protocol.

Options *filename*—Name of the saved file. You can specify a filename in one of the following ways:

- *filename*—File in the user's home directory (the current directory) on the local flash drive.
- *path/filename*—File on the local flash drive.
- /var/*filename* or /var/*path*/*filename*—File on the local hard disk.
- a:*filename* or a:*path*/*filename*—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- *hostname:/path/filename, hostname:filename, hostname:path/filename,* or scp://*hostname*/*path*/*filename*—File on an scp/ssh client. This form is not available in the worldwide version of the JUNOS software. The default path is the user's home directory on the remote system. You can also specify *hostname* as *username@hostname*.
- ftp://*hostname*/*path*/*filename*—File on an FTP server. You can also specify *hostname* as *username@hostname* or *username*:*password@hostname*. The default path is the user's home directory. To specify an absolute path, the path must start with the string %2F; for example, ftp://*hostname*/%2F*path*/*filename*. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed:

```
user@host> file copy ftp://username@ftp.hostname.net//filename
file copy ftp.hostname.net: Not logged in.

user@host> file copy ftp://username:prompt@ftp.hostname.net//filename
Password for username@ftp.hostname.net:
```

- http://*hostname*/*path*/*filename*—File on a Hypertext Transfer Protocol (HTTP) server. You can also specify *hostname* as *username@hostname* or *username*:*password@hostname*. If a password is required and you omit it, you are prompted for it.
- re0:*/path/filename* or re1:*/path/filename*—File on a local Routing Engine.

Usage Guidelines See "Saving a Configuration to a File" on page 263.

Required Privilege Level configure—To enter configuration mode.

set

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

Release Information Command introduced before JUNOS Release 7.4.

Description Create a statement hierarchy and set identifier values. This is similar to edit except that your current level in the hierarchy does not change.

Options *identifier*—Name of the statement or identifier to set.

statement-path—(Optional) Path to an existing statement hierarchy level. If that hierarchy level does not exist, it is created.

Usage Guidelines See "Creating and Modifying the Configuration" on page 235.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also edit on page 298

show

Syntax show <*statement-path*> <*identifier*>

Release Information Command introduced before JUNOS Release 7.4.

Description Display the current configuration.

Options none—Display the entire configuration at the current hierarchy level.

identifier—(Optional) Display the configuration for the specified identifier.

statement-path—(Optional) Display the configuration for the specified statement hierarchy path.

Usage Guidelines See "Displaying the Current Configuration" on page 241.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See the following sections:

- show | display inheritance defaults on page 305
- show | display set on page 305
- show | display set relative on page 306
- show groups junos-defaults on page 307

show | display inheritance defaults

Syntax show | display inheritance defaults <| grep compress>

Release Information Command introduced before JUNOS Release 7.4.

Description Display the JUNOS software defaults that have been applied to the configuration.

Usage Guidelines See “Using JUNOS Default Groups” on page 639 and “Compressing the Current Configuration File” on page 383.

Options | grep compress—Display information about the compression of the current operational configuration.

Required Privilege Level view

Sample Output

```
user@host# show system ports | display inheritance defaults
## ## 'console' was inherited from group 'junos-defaults'
## 'vt100' was inherited from group 'junos-defaults'
## console type vt100;
```

show | display set

Syntax show | display set

Release Information Command introduced before JUNOS Release 7.4.

Description Display the configuration as a series of configuration mode commands required to recreate the configuration from the top level of the hierarchy as set commands

Usage Guidelines See “Displaying set Commands from the Configuration” on page 243.

Required Privilege Level view

Sample Output

```
user@host# show | display set
set interfaces fe-0/0/0 unit 0 family inet address 192.168.1.230/24
set interfaces fe-0/0/0 unit 0 family iso
set interfaces fe-0/0/0 unit 0 family mpls
set interfaces fe-0/0/0 unit 1 family inet address 10.0.0.1/8
deactivate interfaces fe-0/0/0 unit 1
```

show | display set relative

Syntax show | display set relative

Release Information Command introduced before JUNOS Release 7.4.

Description Display the configuration as a series of configuration mode commands required to recreate the configuration from the current hierarchy level.

Usage Guidelines See “Displaying set Commands from the Configuration” on page 243.

Required Privilege Level view

Sample Output

```
[edit interfaces fe-0/0/0]
user@host# show
unit 0 {
    family inet {
      address 192.107.1.230/24;
    }
    family iso;
    family mpls;
}
inactive: unit 1 {
    family inet {
      address 10.0.0.1/8;
    }
}
user@host# show | display set relative
set unit 0 family inet address 192.107.1.230/24
set unit 0 family iso
set unit 0 family mpls
set unit 1 family inet address 10.0.0.1/8
deactivate unit 1
```

show groups junos-defaults

Syntax show groups junos-defaults

Release Information Command introduced before JUNOS Release 7.4.

Description Display the full set of available preset statements from the JUNOS software default group.

Usage Guidelines See “Using JUNOS Default Groups” on page 639.

Required Privilege Level view

Sample Output

```
user@host# show groups junos-defaults
groups {
    junos-defaults {
        applications {
            #
            # File Transfer Protocol
            #
            application junos-ftp {
                application-protocol ftp;
                protocol tcp;
                destination-port 21;
            }
            #
            # Trivial File Transfer Protocol
            #
            application junos-tftp {
                application-protocol tftp;
                protocol udp;
                destination-port 69;
            }
            #
            # RPC port mapper on TCP
            #
            application junos-rpc-portmap-tcp {
                application-protocol rpc-portmap;
                protocol tcp;
                destination-port 111;
            }
            #
            # RPC port mapper on UDP
            #
        }
    }
}
```

status

Syntax status

Release Information Command introduced before JUNOS Release 7.4.

Description Display the users currently editing the configuration.

Usage Guidelines See “Displaying Users Currently Editing the Configuration” on page 245.

Required Privilege Level configure—To enter configuration mode.

top

Syntax top <*configuration-command*>

Release Information Command introduced before JUNOS Release 7.4.

Description Return to the top level of configuration command mode, which is indicated by the [edit] banner.

Option *configuration-command*—Issue configuration mode commands from the top of the hierarchy.

Usage Guidelines See “Moving Among Levels of the Hierarchy” on page 238 and “Issuing Relative Configuration Commands” on page 240.

Required Privilege Level configure—To enter configuration mode.

See Also exit on page 298, up on page 308

up

Syntax up <*number*> <*configuration-command*>

Release Information Command introduced before JUNOS Release 7.4.

Description Move up one level in the statement hierarchy.

Options none—Move up one level in the configuration hierarchy.

number—(Optional) Move up the specified number of levels in the configuration hierarchy.

configuration-command—Issue configuration mode commands from a location higher in the hierarchy.

Usage Guidelines See “Moving Among Levels of the Hierarchy” on page 238 and “Issuing Relative Configuration Commands” on page 240.

Required Privilege Level configure—To enter configuration mode.

See Also exit on page 298, top on page 308

wildcard

Syntax wildcard delete <*statement-path*> <*identifier*> <*regular-expression*>

Release Information Command introduced before JUNOS Release 7.4.

Description Delete a statement or identifier. All subordinate statements and identifiers contained within the specified statement path are deleted with it.

Deleting a statement or an identifier effectively "unconfigures" or disables the functionality associated with that statement or identifier.

If you do not specify *statement-path* or *identifier*, the entire hierarchy starting at the current hierarchy level is removed.

Options delete—Delete several related configuration items simultaneously, such as channelized interfaces or static routes, by using a single command and regular expressions.

statement-path—(Optional) Path to an existing statement or identifier. Include this if the statement or identifier to be deleted is not at the current hierarchy level.

identifier—(Optional) Name of the statement or identifier to delete.

Usage Guidelines See "Using Regular Expressions to Remove Related Configuration Items" on page 248.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

Chapter 11

Summary of CLI Operational Mode Commands

The following sections explain each of the command-line interface (CLI) operational mode commands. The commands are organized alphabetically.

clear

Syntax clear (arp | bfd | bgp | chassis | cli | firewall | igmp | ilmi | interfaces | ipsec | ipv6 | isis | ldp | log | mpls | msdp | multicast | ospf | ospf3 | pgm | pim | pppoe | rip | route | rsvp | snmp | system | vpls | vrrp)

Release Information Command introduced before JUNOS Release 7.4.

Description Clear statistics and protocol database information.

Usage Guidelines Most clear commands are discussed in the *JUNOS System Basics and Services Command Reference*. The following clear commands are discussed in the *JUNOS Interfaces Command Reference*: clear ike, clear ilmi, clear interfaces, clear ipsec, clear pppoe

Required Privilege Level clear

configure

Syntax configure | configure exclusive | configure private

Release Information Command introduced before JUNOS Release 7.4.

Description Enter configuration mode.

Usage Guidelines See “Entering Configuration Mode” on page 222.

Required Privilege Level configure

file

Syntax file (archive | checksum | compare | copy | delete | list | rename | show)

Release Information Command introduced before JUNOS Release 7.4.

Description Copy files to and from the router, compare files, or delete a file on a local router.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

See Also file archive on page 312

Required Privilege Level maintenance

file archive

Syntax file archive <compress> source *source* destination *destination*

Release Information Command introduced before JUNOS Release 7.4.

Description Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.

Options compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix .tgz.

source—Source of the original file or files. Specify the source as a URL or filename, as described in "How to Specify Filenames and URLs" on page 46.

destination—Destination of the archived file or files. Specify the destination as a URL or filename, as described in "How to Specify Filenames and URLs" on page 46. The JUNOS software adds one of the following suffixes if the destination filename does not already have it:

- For archived files—the suffix .tar
- For archived and compressed files—the suffix .tgz

Required Privilege Level maintenance

Sample Output: file archive for multiple files The following sample command archives all the message files in the local directory /var/log/messages as the single file messages-archive.tar in the same directory:

```
user@host> file archive source /var/log/messages* destination
/var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

Sample Output: file archive for one file The following sample command archives one message file in the local directory /var/log/messages as the single file messages-archive.tar in the same directory:

```
user@host> file archive source /var/log/messages destination
/var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

Sample Output: file archive compress

The following sample command archives and compresses all the message files in the local directory /var/log/messages as the single file messages-archive.tgz in the same directory:

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

help

Syntax help <(topic | reference | syslog | tip (cli <*number*>))>

Release Information Command introduced before JUNOS Release 7.4.

Description Display help about available configuration statements or general information on getting help.

Options Display all hierarchy levels containing the statement.

reference—Display summary information for the statement.

syslog—Display system log messages.

tip cli <*number*> —Display CLI tips that are associated with a number.

topic—(Optional) Display usage guidelines for the statement.

Entering the help command without an option provides introductory information on using the help and ? commands.

Usage Guidelines See “Getting Help Based on a String in a Statement Name” on page 185, “Displaying Tips About CLI Commands” on page 185, and “Configuring Tips” on page 412.

Required Privilege Level None.

monitor

Syntax monitor (interface | label-switched-path | list | traffic start | stop |)

Release Information Command introduced before JUNOS Release 7.4.

Description Monitor a log file or interface traffic in real time.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference.*

Required Privilege Level Depends on the specific command.

mtrace

Syntax mtrace *source* <from-source | monitor | to-gateway>

Release Information Command introduced before JUNOS Release 7.4.

Description Display trace information about a multicast path.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level view

ping

Syntax ping (host | atm | mpls | mpls l2circuit | mpls l2vpn | mpls l3vpn)

Release Information Command introduced before JUNOS Release 7.4.

Description Check the reachability of network hosts.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level network

| (pipe)

Syntax | (compare | count | display (changed | commit-scripts | detail | display set | inheritance | xml) | except *pattern* | find *pattern* | hold | last *lines* | match *pattern* | no-more | request message (all | *account@terminal*) resolve <full-names> | save *filename* | trim *columns*)

Release Information Command introduced before JUNOS Release 7.4.
display commit-scripts option added in JUNOS Release 7.4.

Description Filter the output of an operational mode or a configuration mode command.

Options compare (filename | rollback *n*)—(Configuration mode only, and only with the show command) Compare configuration changes with another configuration file.

count—Display the number of lines in the output.

display—Display additional information about the configuration contents.

- changed—Tag changes with junos:changed attribute (XML only).

- commit-scripts—(Configuration mode only) Display all statements that are in a configuration, including statements that were generated by transient changes. For more information, see the *JUNOS Configuration Scripting Guide*.

- detail—(Configuration mode only) Display configuration data detail.

- inheritance <brief | default | groups | terse>—(Configuration mode only) Display inherited configuration data and source group.
- set—Display the configuration as a series of configuration mode commands required to recreate the configuration.
- xml—(Operational mode only) Display the command output as JUNOScript (Extensible Markup Language [XML]) tags.

except *pattern*—Ignore text matching a regular expression when searching the output. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.

find *pattern*—Display the output starting at the first occurrence of text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks (" ").

last *lines*—Display the last number of lines you want to view from the end of the configuration. For more information about how to display the last number of lines from the end of the configuration, see "Displaying the Lines You Want to View from the End of the Output" on page 196.

hold—Hold text without exiting the –More– prompt.

match *pattern*—Search for text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.

no-more—Display output all at once rather than one screen at a time.

resolve—Convert IP addresses into Domain Name System (DNS) names. Truncates to fit original size unless full-names specified. To prevent the names from being truncated, use the full-names option.

request message (all | *account@terminal*)—Display command output on the terminal of a specific user logged in to your router, or on the terminals of all users logged in to your router.

save *filename*—Save the output to a file or URL. For information about specifying the filename, see "Specifying Filenames and URLs" on page 360.

trim *columns*—Trim specified number of columns from the start line.

Usage Guidelines See "Filtering Command Output" on page 190 and "Moving Among Levels of the Hierarchy" on page 238.

quit

Syntax quit

Release Information Command introduced before JUNOS Release 7.4.

Description Exit from the CLI to a UNIX shell.

Required Privilege Level shell and maintenance

See Also start on page 321

request

Syntax request <chassis | ipsec switch | message | mpls | routing-engine | security | services | system | flow-collector | support information>

Release Information Command introduced before JUNOS Release 7.4.

Description Stop or reboot router components, switch between primary and backup components, display messages, and display system information.

NOTE: If your routing platform contains two Routing Engines and you want to shut the power off to the routing platform or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded) and then the master Routing Engine. To halt a Routing Engine, issue the request system halt command. You can also halt both Routing Engines at the same time by issuing the request system halt both-routing-engines command.

If you want to reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) and then the master Routing Engine.

CAUTION: Halt the backup Routing Engine before you remove it or shut off the power to the router; otherwise you might need to reinstall the JUNOS software.

NOTE: If you reboot the TX Matrix platform, all the T640 master Routing Engines connected to the TX Matrix platform reboot. If you halt both Routing Engines on a TX Matrix platform, all the T640s Routing Engines connected to the TX Matrix platform are also halted.

NOTE: If you insert a Flexible PIC Concentrator (FPC) into your routing platform, you may need to issue the request chassis fpc command (or press the online button) to bring the FPC online. This applies to FPCs in M20, M40, M40e, M160, M320, and T-series platforms. For command usage, see the request chassis fpc command description in the *JUNOS System Basics and Services Command Reference*.

Usage Guidelines Most request commands are discussed in the *JUNOS System Basics and Services Command Reference*. The following request commands are discussed in the *JUNOS Interfaces Command Reference*: request ipsec switch and request services.

Required Privilege Level maintenance

request security certificate

See the following sections:

- request security certificate (for CA Certificate) on page 317
- request security certificate (for Signed Certificate) on page 318

request security certificate (for CA Certificate)

Syntax request security certificate enroll filename *filename* ca-file *ca-file* ca-name *ca-name* url *url*

Release Information Command introduced before JUNOS Release 7.4.

Description Obtain a certificate from a certificate authority (CA). The results are saved in the specified file in the /var/etc/ikecert directory.

Options ca-file *ca-file*—Name of certificate authority profile in configuration.

ca-name *ca-name*—Name of certificate authority

filename *filename*—File that stores the public key certificate.

url *url*—Certificate authority URL.

Required Privilege Level maintenance

Sample Output

```
user@host> request security certificate enroll filename ca_verisign ca-file
verisign ca-name juniper.net url
http://pilotonsiteipsec.verisign.com/cgi-bin/pkiclient.exe
URL: http://pilotonsiteipsec.verisign.com/cgi-bin/pkiclient.exe CA name:
juniper.net CA file: verisign Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
the key management process (kmd) log file at /var/log/kmd. <--------------
```

request security certificate (for Signed Certificate)

Syntax request security certificate enroll filename *filename* subject *subject* alternative-subject *alternative-subject* certificate-authority *certificate-authority* key-file *key-file* domain-name *domain-name*

Release Information Command introduced before JUNOS Release 7.4.

Description Obtain a signed certificate from a certificate authority. The signed certificate validates the CA and the owner of the certificate. The results are saved in a specified file to the /var/etc/ikecert directory.

Options alternative-subject *alternative-subject*—A tunnel source address.

certification-authority *certificate-authority*—Name of certification authority profile in configuration.

domain-name *domain-name*—Fully qualified domain name.

filename *filename*—Name of file that stores the certificate.

key-file *key-file*—File containing a local private key.

subject *subject*—A distinguished name (DN), which consists of a set of components—for example, an organization (O), an organization unit (OU), a country (C), and a locality (L).

Required Privilege Level maintenance

Sample Output:

```
user@host> request security certificate enroll filename host.crt subject
c=uk,o=london alternative-subject 10.50.1.4 certification-authority verisign
key-file host-1.prv domain-name host.juniper.net
CA name: juniper.net CA file: ca_verisign
local pub/private key pair: host.prv
subject: c=uk,o=london domain name: host.juniper.net
alternative subject: 10.50.1.4
Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
the key management process (kmd) log file at /var/log/kmd. <--------------
```

restart

Syntax

```
restart
<adaptive-services | chassis-control | class-of-service | disk-monitoring |
ecc-error-logging | firewall | interface-control | kernel-replication | l2tpd-service |
mib-process | network-access-service | pgm | pic-services-logging | pppoe |
remote-operations | routing <logical-router logical-router-name> | sampling |
service-deployment | snmp | web-management>
<gracefully | immediately | soft>
```

Release Information Command introduced before JUNOS Release 7.4.

Description Restart router software processes on all platforms (with the exception of routing matrixes and J-series Services Routers).

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level reset

restart (Routing Matrix)

Syntax

```
restart
<adaptive-services | chassis-control | class-of-service | disk-monitoring |
ecc-error-logging | firewall | interface-control | kernel-replication | l2tpd-service |
mib-process | network-access-service | pgm | pic-services-logging | pppoe |
remote-operations | routing <logical-router logical-router-name> | sampling |
service-deployment | snmp | web-management>
<all | all-lcc | lcc number>
<gracefully | immediately | soft>
```

Release Information Command introduced before JUNOS Release 7.4.

Description Restart router software processes on a routing matrix.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference.*

Required Privilege Level reset

restart (J-series Services Routers)

Syntax restart
<adaptive-services | chassis-control | class-of-service | dhcp | firewall |
interface-control | l2tpd-service | mib-process | network-access-service | pgm | pppoe |
remote-operations | routing <logical-router *logical-router-name*> | sampling |
service-deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>

Release Information Command introduced before JUNOS Release 7.4.

Description Restart router software processes on J-series Services Routers.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level reset

restart routing

Syntax restart routing <logical-router *logical-router-name*> <gracefully | immediately | soft>

Release Information Command introduced before JUNOS Release 7.4.

Description Restart the JUNOS software process on a logical router.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level reset

set

Syntax set (chassis | cli | date)

Release Information Command introduced before JUNOS Release 7.4.

Description Configure chassis, CLI properties, and the router's date and time.

Usage Guidelines See "Controlling the CLI Environment" on page 213 and "Setting the Current Date and Time" on page 200. For information about setting chassis properties, see the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level view

show

Syntax show (accounting | aps | arp | as-path | bfd | bgp | chassis | class-of-service | cli | configuration | connections | dvmrp | dynamic-tunnels | firewall | helper | host | igmp | ike | ilmi | interfaces | ipsec | ipv6 | isis | l2circuit | l2vpn | ldp | link-management | log | mld | mpls | msdp | multicast | ntp | ospf | ospf3 | passive-monitoring | pfe | pgm | pim | policer | policy | pppoe | rip | ripng | route | rsvp | sap | services | snmp | system | task | ted | version | vpls | vrrp)

Release Information Command introduced before JUNOS Release 7.4.

Description Show information about all aspects of the software, including interfaces and routing protocols.

Usage Guidelines Most show commands are discussed in the *JUNOS System Basics and Services Command Reference*. The following show commands are discussed in the *JUNOS Interfaces Command Reference*: show aps, show ike, show ilmi, show interfaces, show ipsec, show passive-monitoring, show pppoe, show services, and show vrrp.

Required Privilege Level Depends on the specific command.

ssh

Syntax ssh

Release Information Command introduced before JUNOS Release 7.4.

Description Open an SSH shell connected to another host.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level network

start

Syntax start shell

Release Information Command introduced before JUNOS Release 7.4.

Description Start a UNIX shell on the router.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference*.

Required Privilege Level shell and maintenance

telnet

Syntax telnet

Release Information Command introduced before JUNOS Release 7.4.

Description Establish a telnet connection to another host.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference.*

Required Privilege Level network

test

Syntax test (configuration | interface | msdp | policy)

Release Information Command introduced before JUNOS Release 7.4.

Description Run various diagnostic debugging commands.

Usage Guidelines Most of the test commands are discussed in the *JUNOS System Basics and Services Command Reference*. The test interface command is discussed in the *JUNOS Interfaces Command Reference*.

Required Privilege Level Depends on the specific command.

traceroute

Syntax traceroute

Release Information Command introduced before JUNOS Release 7.4.

Description Trace the route to a remote host.

Usage Guidelines See the *JUNOS System Basics and Services Command Reference.*

Required Privilege Level network

update

Syntax update

Release Information Command introduced before JUNOS Release 7.4.

Description Update private candidate configuration with a copy of the most recently committed configuration, including your private changes.

Usage Guidelines See “Updating the Configure Private Configuration” on page 227.

NOTE: The update command is available only when you are in configure private mode.

Part 3

Software Installation and Upgrade

Chapter 12

Installation Overview

Your router comes with JUNOS software installed on it. When you power on the router, all software starts automatically. You simply need to configure the software and the router will be ready to participate in the network.

The software is installed on the router's flash drive (a nonrotating drive) and hard disk (a rotating disk). A copy of the software also is provided on removable media, either an LS-120 floppy disk or a PC card, which can be inserted into the router's drive or card slot. Normally, when you power on the router, it runs the copy of the software that is installed on the flash drive.

You might want to upgrade the router software as new features are added or software problems are fixed. You normally obtain new software by downloading the images onto your router or onto another system on your local network. Then you install the software upgrade on the router's flash and hard disks. You can also copy the software onto the removable media.

If the software on the flash drive, hard disk, or removable media becomes damaged, you can reinstall the software onto those devices.

This chapter discusses the following concepts and terminology related to installing and upgrading the JUNOS software:

JUNOS Software Distribution

This section discusses the following topics:

- Software Release Names on page 326
- Package Names on page 326

Software Release Names

A JUNOS software release has a name in the following format:

JUNOS-*m.nZx*

m.n is two integers that represent the software release number; *m* denotes the major release number.

Z is a capital letter that indicates the type of software release. In most cases, it is an R, to indicate that this is released software. If you are involved in testing prereleased software, this letter might be an A (for alpha-level software), B (for beta-level software), or I (a capital letter I; for internal, test, or experimental versions of software).

x represents the version of the major software release.

The following is an example of a software release name:

JUNOS-7.0R1

Package Names

A *package* is a collection of files that make up a software component.

NOTE: All JUNOS software is delivered in signed packages that contain digital signatures, Secure Hash Algorithm (SHA-1), and Message Digest 5 (MD5) checksums. Which checksum is used depends on the software version:

- Digital signatures are used when you upgrade or downgrade between JUNOS Release 7.0 and a later version.
- The SHA-1 checksum is used when you upgrade or downgrade between JUNOS Release 6.4 and a later version.
- The MD5 checksum is used when you upgrade or downgrade from JUNOS Release 6.3 or earlier to a later version.

You can upgrade the packages individually.

A software package has a name in the following format:

package-name-release.tgz or *package-name-release*-signed.tgz

package-name is the name of the package. Examples are jroute (the routing package) and jkernel (the operating system package).

Each JUNOS software release consists of a set of software packages whose names contain the package name and the software release version, and includes the following components:

- Kernel and network tools package, which contains the operating system
- Base package, which contains additions to the operating system
- Routing package, which contains the software that runs on the Routing Engine
- Encryption package, which contains security software (domestic version)
- Packet Forwarding Engine software package
- J-Web package which contains the graphical user interface software for M-series and T-series routers
- Documentation package, which contains the documentation for the software

release is the software release number; for example, 7.4 R1 or 7.3 R1.5.

The following are examples of package names:

```
jroute-7.4R1-signed.tgz
jkernel-7.4R1-signed.tgz
jpfe-7.4R1-signed.tgz
```

The JUNOS software contains two additional packages: jinstall and jbundle. These packages contain the complete JUNOS software.

- jinstall—A package used to upgrade from JUNOS Release 6.*x* to 7.*x* or when the software becomes damaged. If you upgrade from 6.*x* to 7.*x* using jinstall, use jbundle for subsequent upgrades or downgrades. The jinstall package completely reinstalls the software. It rebuilds the JUNOS file system only and retains configuration information from the previous version. However, logs and other types of auxiliary information may be erased during installation. For more information about how to use jinstall when the software becomes damaged, see "Reinstalling Software Using jinstall" on page 349.
- jbundle—A package used to downgrade from Release 7.4. jbundle is also used to upgrade or downgrade between minor versions of the JUNOS software. jbundle modifies the smallest set of files needed to change to the new software version.

  The jweb package is not included in the jinstall and jbundle software bundles.

NOTE: You cannot use the jbundle package to upgrade from JUNOS 5.*x* to JUNOS 7.*x*. For more information about how to upgrade from Release 5.*x* to 7.*x* or downgrade from Release 7.*x* to 5.*x*, see "Upgrading to Release 7.x from Release 5.x" on page 353.

Two sets of JUNOS software packages are provided: one for customers in the United States and Canada and another for other customers. The worldwide version does not include any high encryption capabilities for data leaving the router. Otherwise, the two packages are identical.

Storage Media

The router has three forms of storage media:

- Flash drive, which is a nonrotating drive. When a new router is shipped from the factory, the JUNOS software is preinstalled on the flash drive.
- Hard disk, which is a rotating drive. When a new router is shipped from the factory, the JUNOS software is preinstalled on the hard disk. This drive also is used to store system log files and diagnostic dump files.
- Removable media, either a PC card or an LS-120 floppy disk. The removable media that ships with each router contains a copy of the JUNOS software.

Table 13 specifies the router's device names. The device names are displayed when the router boots.

Table 13: Device Names

Device	Flash Drive	Hard Disk	Removable Media
Routing Engine 200 (RE-M40) (CLI Name = RE1)	none	ad0	ad1
Routing Engine 333 (CLI Name = RE2)	ad0	ad1	ad3
Routing Engine 600 (CLI Name = RE3)	ad0	ad1	ad3
Routing Engine 1600 (CLI Name = RE4)	ad0	ad1	ad3 and ad4
Routing Engine 400 (CLI Name = RE5)	ad0	ad1	ad3

Boot Devices

There are three devices from which the router boots: the flash drive, the hard disk, or a removable medium. Typically, the router boots from the flash drive. The disk from which the router boots is called the *primary boot device*, and the other disk is the *alternate boot device.* The primary boot device is generally the flash drive, and the alternate boot device is generally the hard disk.

NOTE: If the router boots from an alternate boot device, a yellow alarm lights the LED on the router's craft interface. Some routers have text and LED displays.

For information about chassis conditions that trigger alarms, see "Chassis Conditions That Trigger Alarms" on page 805.

Boot Sequence

Normally, the router boots from the flash drive. If it fails, the router attempts to boot from the hard disk, which is the alternate boot device.

If a removable medium is installed when the router boots, the router attempts to boot the image on it. If the booting fails, the router tries the flash drive and then the hard disk.

When the router boots from the storage media (flash drive, hard disk, or removable media) it expands its search in the /config directory of the routing platform for the following files in the following order: juniper.conf (the main configuration file), rescue.conf (the rescue configuration file), and juniper.conf.1 (the first rollback configuration file). When the search finds the first configuration file that can be loaded properly, the file loads and the search ends. If none of the files can be loaded properly, the routing platform does not function properly.

NOTE: To reinstall the JUNOS software, you boot the router from the removable media. Do not insert the removable media during normal operations. The router does not operate normally when it is booted from the removable media.

If the router boots from an alternate boot device, the JUNOS software displays a message indicating this when you log in to the router. For example, this message shows that the software booted from the hard disk (/dev/ad2s1a):

```
login: username
Password: password
Last login: date on terminal

— JUNOS 7.4 R1 built date
—
— NOTICE: System is running on alternate media device (/dev/ad2s1a).
```

The default boot order for the M7i Internet router is different from other Juniper Networks routers, because the default configuration of the Routing Engine on the M7i router does not include an internal compact flash disk.

If the Routing Engine does not have an internal compact flash disk, two copies of the JUNOS software are preinstalled on the router: one on a PC card that can be inserted into the slot in the Routing Engine faceplate, and one on a rotating hard disk in the Routing Engine. When the router boots, it first attempts to access the software image on the PC card. If a PC card is not inserted into the Routing Engine or the attempt otherwise fails, the router tries the hard disk.

If the Routing Engine has an internal compact flash disk, three copies of the JUNOS software are preinstalled on the router. When the router boots, it first attempts to access the image on the PC card. If a PC card is not inserted into the Routing Engine or the attempt otherwise fails, the router next tries the flash disk, and finally the hard disk.

Installing JUNOS-FIPS

JUNOS-FIPS has special installation and configuration requirements. Installation procedures include downloading the FIPS software package from www.juniper.net. For more information, see the *JUNOS-FIPS Configuration Guide*

Verifying PIC Combinations

On Juniper Networks routing platforms, you can typically install any combination of Physical Interface Cards (PICs) on a single Enhanced Flexible PIC Concentrator (FPC) or two PIC slots served by a single Layer 2/Layer 3 Packet Processing application-specific integrated circuit (ASIC). However, on some routers, there are some combinations of PICs that cannot be installed together.

During software installation, the configuration checker in the installation program checks for invalid PIC combinations. If the configuration checker finds an invalid combination of PICs, the installation process stops and displays a message similar to the following:

```
user@host> request system software add jbundle-7.4R1.x-domestic-signed.tgz
50306991 bytes transferred in 6.0 seconds (8.05 MBps)
Package contains jbundle-7.4-R1.x-domestic-signed.tgz; renaming...
Installing package
'/var/tmp/jbundle-7.4R1.x-domestic-signed.tgz'... Verified MD5 checksum of
jbundle-7.4R1.x-domestic-signed.tgz
Adding jbundle...
Available space: 36416 require: 28135
FPC slot 4 contains invalid combination of PICS
    PIC slot 1 - 4 port E3 Intelligent Queuing PIC ID
    PIC slot 2 - Single GigEther Intelligent Queuing PIC ID
    PIC slot 3 - Single ATM-II OC12 IQ PIC ID
WARNING: This installation attempt will be aborted. If you
WARNING: wish to force the installation despite these warnings
WARNING: you may use the 'force' option on the command line.
pkg_add: package /var/tmp/jbundle-7.4R1.x-domestic-signed.tgz fails
requirements - not
installed
pkg_add: install script returned error status
```

The configuration checker has the following limitations:

- If a PIC is offline when you upgrade the router with new software, the configuration checker cannot detect invalid PIC combinations and cannot warn about them.
- If you specify the force option when you upgrade the JUNOS software, the configuration checker warns about the invalid PIC combination and the software installation continues. However, after rebooting, one or more PICs on the affected router may fail to initialize.
- The configuration checker looks for combinations of three invalid PICs. If an Enhanced FPC contains four invalid PICs, the script generates multiple warnings to remove an invalid PIC from that Enhanced FPC.

If you install a PIC into a router already running JUNOS software, you can identify the presence of invalid PIC combinations from messages in the system logging (syslog) file:

```
Feb 6 17:57:40 CE1 feb BCHIP 0: uCode overflow - needs 129 inst space to
load b3_atm2_LSI_decode for stream 12
Feb 6 17:57:41 CE1 chassisd[2314]: CHASSISD_IFDEV_DETACH_PIC:
ifdev_detach_pic(0/3)
Feb 6 17:57:41 CE1 feb BCHIP 0: binding b3_atm2_LSI_decode to stream 12
failed
Feb 6 17:57:41 CE1 feb PFE: can not bind B3 ucode prog b3_atm2_LSI_decode to
FPC 0: stream 12
```

For more information about checking for unsupported PIC combinations, see the corresponding PIC guide for your router, the *JUNOS Release Notes*, and *Technical Support Bulletin PSN-2004-12-002* on the Juniper Networks Support Web site at http://www.juniper.net/support/.

Available Disk Space

During software installation or upgrades, the installation program may fail if your router has a shortage of disk space. If a disk space error occurs, use one or more of the following options to complete the installation:

- Use the request system storage cleanup command to delete unnecessary files and increase storage space on the router (J-series Services Routers only).

- Specify the unlink option when you use the request system software add command to install the JUNOS software.

- Download the software packages you need from the Juniper Networks Support Web site, http://www.juniper.net/support/. The download program provides intelligent disk space management to enable installation.

For more information on the request system storage cleanup command and the request system software add command, see the *JUNOS System Basics and Services Command Reference*.

Chapter 13

Configuring the Software Initially

You can configure the router from a system console connected to the router's console port or by using telnet to access the router remotely.

Before you configure the software for the first time, you need the following information:

- Name of the machine
- Machine's domain name
- IP address and prefix length information for router's management Ethernet interface
- IP address of a default router
- IP address of a Domain Name System (DNS) server
- Password for the user "root"

To configure the software for the first time, follow these steps:

1. Power on the router. The JUNOS software boots automatically.
2. Log in as the user root. There is no password.
3. Start the command-line interface (CLI):

   ```
   root# cli
   root@>
   ```

4. Enter configuration mode:

   ```
   cli> configure
   [edit]
   root@#
   ```

5. Configure the name of the machine. If the name includes spaces, enclose the entire name in quotation marks (" ").

   ```
   [edit]
   root@# set system host-name host-name
   ```

6. Configure the machine's domain name:

   ```
   [edit]
   root@# set system domain-name domain-name
   ```

7. Configure the IP address and prefix length for the router's management Ethernet interface:

   ```
   [edit]
   root@# set interfaces fxp0 unit 0 family inet address address/prefix-length
   ```

8. Configure the IP address of a default router. This system is called the backup router because it is used only while the routing protocol process is not running.

   ```
   [edit]
   root@# set system backup-router address
   ```

9. Configure the IP address of a DNS server:

   ```
   [edit]
   root@# set system name-server address
   ```

10. Set the root password, entering either a clear-text password that the system will encrypt, a password that is already encrypted, or an SSH public key string.

    To enter a clear-text password, use the following command to set the root password:

    ```
    [edit]
    root@# set system root-authentication plain-text-password
    New password: type password
    Retype new password: retype password
    ```

    To enter a password that is already encrypted, use the following command to set the root password:

    ```
    [edit]
    root@# set system root-authentication encrypted-password
    encrypted-password
    ```

    To enter an SSH public string, use the following command to set the root password:

    ```
    [edit]
    root@# set system root-authentication ssh-rsa key
    ```

NOTE: JUNOS-FIPS has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If JUNOS-FIPS is installed, you cannot configure passwords unless they meet this standard.

11. Optionally, display the configuration statements:

```
[edit]
root@ show
system {
      host-name host-name;
      domain-name domain.name;
      backup-router address;
      root-authentication {
         (encrypted-password "password" | public-key);
         ssh-rsa "public-key";
         ssh-dsa "public-key";
      }
      name-server {
         address;
}
interfaces {
      fxp0 {
         unit 0 {
            family inet {
               address address;
            }
         }
      }
}
```

12. Commit the configuration, which activates the configuration on the router:

```
[edit]
root@# commit
```

13. If you want to configure additional properties at this time, remain in configuration mode and add the necessary configuration statements. Then commit the changes to activate them on the router:

```
[edit]
root@host-name# commit
```

14. When you have completed configuring the router, exit from configuration mode:

```
[edit]
root@host-name# exit
root@host-name>
```

15. After you have installed the software on the router, committed the configuration, and are satisfied that the new configuration is successfully running, you should issue the request system snapshot command to back up the new software to the /altconfig file system. If you do not issue the request system snapshot command, the configuration on the alternate boot drive will be out of sync with the configuration on the primary boot drive.

    The request system snapshot command causes the root file system to be backed up to /altroot, and /config to be backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue this command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Chapter 14

Reinstalling the Software Using the Install Media

If any of the software becomes damaged, you can reinstall it. This chapter discusses the following topics related to reinstalling the JUNOS software:

- Preparing to Reinstall the JUNOS Software on page 337
- Reinstalling the JUNOS Software on page 338
- Reconfiguring the JUNOS Software on page 338

Preparing to Reinstall the JUNOS Software

Before you install the JUNOS software, you must do the following:

1. Copy the existing configuration in the file /config/juniper.conf from the router to another system or to removable media. You also might want to copy any backup configurations (the files named /config/juniper.conf.*n*, where *n* is a number from 0 through 9). To copy the files, use the file copy command.
2. Have available the removable medium that shipped with the router (also called a boot floppy). If you do not have a boot floppy, contact customer support.

Reinstalling the JUNOS Software

To reinstall the JUNOS software, follow these steps:

1. Insert the removable medium into the router.

NOTE: You can store a configuration on install media such as an (LS-120 floppy disk or PC card). For more information about storing a configuration on an install media, see "Copying a Configuration to a PC Card or LS-120 Floppy Disk" on page 347.

2. Reboot the router, either by power-cycling it or by issuing the request system reboot command from the command-line interface (CLI).

3. When the software asks the following question, type **y**:

```
WARNING: The installation will erase the contents of your disk. Do you wish
to continue (y/n)?
```

4. The router then copies the software from the removable medium onto your system, occasionally displaying status messages. Copying the software can take up to 10 minutes.

5. Remove the removable medium when prompted. The router then reboots from the primary boot device on which the software was just installed. When the reboot is complete, the router displays the login prompt.

Reconfiguring the JUNOS Software

After you have reinstalled the software, you must copy the router's configuration files back to the router. (You also can configure the router from scratch, as described in "Configuring the Software Initially" on page 333.) However, before you can copy the configuration files, you must establish network connectivity.

To reconfigure the software, follow these steps:

1. Log in as root. There is no password.

2. Start the CLI:

```
root# cli
root@>
```

3. Enter configuration mode:

```
cli> configure
[edit]
root@#
```

4. Configure the name of the machine. If the name includes spaces, enclose the entire name in quotation marks (" ").

   ```
   [edit]
   root@# set system host-name host-name
   ```

5. Configure the machine's domain name:

   ```
   [edit]
   root@# set system domain-name domain-name
   ```

6. Configure the IP address and prefix length for the router's management Ethernet interface:

   ```
   [edit]
   root@# set interfaces fxp0 unit 0 family inet address address/prefix-length
   ```

7. Configure the IP address of a default router. This system is called the backup router because it is used only while the routing protocol process is not running.

   ```
   [edit]
   root@# set system backup-router address
   ```

8. Configure the IP address of a Domain Name System (DNS) server:

   ```
   [edit]
   root@# set system name-server address
   ```

9. Set the root password, entering either a clear-text password that the system will encrypt, a password that is already encrypted, or an SSH public key string.

   To enter a clear-text password, use the following command to set the root password:

   ```
   [edit]
   root@# set system root-authentication plain-text-password
   New password: type password
   Retype new password: retype password
   ```

   To enter a password that is already encrypted, use the following command to set the root password:

   ```
   [edit]
   root@# set system root-authentication encrypted-password
   encrypted-password
   ```

   To enter an SSH public string, use the following command to set the root password:

   ```
   [edit]
   root@# set system root-authentication ssh-rsa key
   ```

NOTE: JUNOS-FIPS has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If JUNOS-FIPS is installed on the router, you cannot configure passwords unless they meet this standard.

10. Commit the changes:

    ```
    [edit]
    root@# commit
    ```

11. Exit from configuration mode:

    ```
    [edit]
    root@# exit
    root@>
    ```

12. To check that the router has network connectivity, issue a ping command to a system on the network:

    ```
    root@> ping address
    ```

    If there is no response, reboot the router.

13. Copy the existing configuration and any backup configurations back to the router. Place the files in the /config directory. To copy the files, use the file copy command.

14. Load and activate the desired configuration:

    ```
    root@> configure
    [edit]
    root@# load merge /config/filename or load replace /config/filename
    [edit]
    root@# commit
    ```

15. After you have installed the software on the router, committed the configuration, and are satisfied that the new configuration is successfully running, you should issue the request system snapshot command to back up the new software to the /altconfig file system. If you do not issue the request system snapshot command, the configuration on the alternate boot drive will be out of sync with the configuration on the primary boot drive.

    The request system snapshot command causes the root file system to be backed up to /altroot, and /config to be backed up to /altconfig. The root and /config file systems are on the router's flash drive and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue this command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Chapter 15

Upgrading Software Packages

Each JUNOS software release consists of the following software packages:

- jkernel—Operating system package
- jbase—Additions to the operating system
- jroute—Software that runs on the Routing Engine
- jpfe—Software that runs on the Packet Forwarding Engine
- jdocs—Documentation for the software
- jcrypto—Encryption software (in domestic software only)
- jweb—Software that runs the J-Web graphical user interface.

The JUNOS software contains two additional packages: jbundle and jinstall. These packages contain the complete JUNOS software.

- jinstall—A package used to upgrade between major versions of the JUNOS software (for example, JUNOS Release 6.*x* to 7.*x*) or when the software becomes damaged. If you upgrade to 7.*x* using jinstall, use jbundle for subsequent upgrades or downgrades. The jinstall package completely reinstalls the software. It rebuilds the JUNOS file system only but retains configuration information from the previous version. However, logs and other types of auxiliary information may be erased during installation. For more information about how to use jinstall when the software becomes damaged, see “Reinstalling Software Using jinstall” on page 349.

- jbundle—A package used to downgrade from Release 7.*x*. jbundle is also used to upgrade or downgrade between minor versions of the JUNOS software. jbundle modifies the smallest set of files needed to change to the new software version.

jbundle cannot be used to upgrade from JUNOS 5.*x* to JUNOS 7.*x*. For more information about how to upgrade from Release 5.*x* to 7.*x*, see “Upgrading to Release 7.x from Release 5.x” on page 353.

The jweb package is not included in the jinstall and jbundle software bundles.

NOTE: The J-series Services Routers use a new software bundle, junos-jseries, for software upgrades. You cannot install the jinstall or jbundle software bundles on the J-series Services Routers. Similarly, you cannot install the junos-jseries bundle on M-series or T-series routing platforms.

To determine which packages are running on the router and to get information about these packages, use the show version command at the top level of the command line interface (CLI).

This chapter discusses the following topics:

- Upgrading All Software Packages on page 342
- Upgrading Individual Software Packages on page 345
- Copying a Configuration to a PC Card or LS-120 Floppy Disk on page 347

Upgrading All Software Packages

To upgrade all software packages, follow these steps:

1. Download the software packages you need from the Juniper Networks Support Web site, http://www.juniper.net/support/. Choose either Canada and U.S. Version or Worldwide Version.

   To download the software packages, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks Web site, https://www.juniper.net/registration/Register.jsp.

NOTE: We recommend that you upgrade all software packages out-of-band using the console because in-band connections can be lost during the upgrade process.

2. Back up the currently running and active file system so that you can recover to a known, stable environment in case something goes wrong with the upgrade:

   ```
   user@host> request system snapshot
   ```

   The root file system is backed up to /altroot, and /config is backed up to /altconfig. The root and /config file systems are on the router’s flash drive, and the /altroot and /altconfig file systems are on the router’s hard disk.

NOTE: After you issue the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

3. Copy each software package to the router. We recommend that you copy them to the /var/tmp directory, which is on the rotating medium (hard disk) and is a large file system.

```
user@host> file copy ftp://username:prompt@ftp.hostname.net/
  filename /var/tmp/filename
```

4. Add the new software package:

```
user@host>  request system software add validate
/var/tmp/jinstall-7.x-package-name-signed.tgz
Checking compatibility with configuration Initializing...
Using jbase-7.x-package-name
Using /var/tmp/jinstall-7.x-package-name.signed.tgz
Verified jinstall-7.x-package-name.tgz signed by
PackageDevelopment_0 Using
/var/validate/tmp/jinstall-signed/jinstall-7.x-package-name.tgz
Using /var/validate/tmp/jinstall/jbundle-7.x-package-name.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.x-package-name.tgz
Using /var/validate/tmp/jbundle/jkernel-7.x-package-name.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.x-package-name.tgz
Using /var/validate/tmp/jbundle/jpfe-7.x-package-name.tgz
Using /var/validate/tmp/jbundle/jdocs-7.x-package-name.tgz
Using /var/validate/tmp/jbundle/jroute-7.x-package-name.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package
'/var/tmp/jinstall-7.x-package-name-signed.tgz'...
Verified jinstall-7.x-package-name-signed.tgz signed by
PackageDevelopment_0
Pre-checking requirements for jinstall...
Auto-deleting old jinstall...
Deleting saved config files...
Deleting bootstrap installer...
Adding jinstall...

WARNING:    This package will load JUNOS 7.x software.
WARNING:    It will save JUNOS configuration files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. It will attempt to preserve dumps
WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.
Saving the config files...
Installing the bootstrap installer...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY.
Use the
WARNING:    'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.
```

```
Saving package file in
/var/sw/pkg/jinstall-7.x-package-name-signed.tgz...
Saving state for rollback...
```

package-name is the full URL to the file. *release-number* is the major software release number; for example, 7.4 R1.

NOTE: When you upgrade or downgrade JUNOS software, we recommend that you include the validate option to the request system software add command to check that the candidate software is compatible with the current configuration. By default, when you add a package with a different release number, the validation check is done automatically.

We do not recommend using the no-validate option to suppress validation. For more information about this command, see the *JUNOS System Basics and Services Command Reference.*

5. Reboot the router to start the new software:

   ```
   user@host> request system reboot
   ```

6. After you have upgraded or downgraded the software and are satisfied that the new software is successfully running, issue the request system snapshot command to back up the new software:

   ```
   user@host> request system snapshot
   ```

   The root file system is backed up to /altroot, and /config is backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Upgrading Individual Software Packages

To upgrade an individual JUNOS software package, follow these steps:

1. Download the software packages you need from the Juniper Networks Support Web site, http://www.juniper.net/support/. Choose either Canada and U.S. Version or Worldwide Version.

   To download the software packages, you must have a service contract and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks Web site, https://www.juniper.net/registration/Register.jsp.

NOTE: We recommend that you upgrade all individual software packages out-of-band using the console or fxp0 interface because in-band connections can be lost during the upgrade process.

2. Back up the currently running and active file system so that you can recover to a known, stable environment in case something goes wrong with the upgrade:

   ```
   user@host> request system snapshot
   ```

   The root file system is backed up to /altroot, and /config is backed up to /altconfig. The root and /config file systems are on the router’s flash drive, and the /altroot and /altconfig file systems are on the router’s hard disk.

NOTE: After you issue the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

3. Copy each software package to the router. You might want to copy them to the /var/tmp directory, which is on the rotating media (hard disk) and is a large file system.

   ```
   user@host> file copy ftp://username:prompt@ftp.hostname.net/
     filename /var/tmp/filename
   ```

4. Add the new software package:

   ```
   user@host> request system software add
   /var/tmp/package-name-signed.tgz
   Checking available free disk space...11200k available, 6076k suggested.
   ```

   package-name is the full URL to the file.

   The system might display the following message:

   ```
   pkg_delete: couldn’t entirely delete package
   ```

   This message indicates that someone manually deleted or changed an item that was in a package. You do not need to take any action; the package is still properly deleted.

If you are upgrading more than one package at the same time, add jbase first and the routing software package jroute later. If you are using this procedure to upgrade all packages at once, add them in the following order:

```
user@host> request system software add /var/tmp/jbase-release-signed.tgz
user@host> request system software add
/var/tmp/jkernel-release-signed.tgz
user@host> request system software add /var/tmp/jpfe-release-signed.tgz
user@host> request system software add /var/tmp/jdocs-release-signed.tgz
user@host> request system software add /var/tmp/jweb-release-signed.tgz
user@host> request system software add
/var/tmp/jroute-release-signed.tgz
user@host> request system software add
/var/tmp/jcrypto-release-signed.tgz
```

5. Reboot the router to start the new software:

   ```
   user@host> request system reboot
   ```

6. After you have upgraded or downgraded the software and are satisfied that the new software is successfully running, issue the request system snapshot command to back up the new software.

   ```
   user@host> request system snapshot
   ```

   The root file system is backed up to /altroot, and /config is backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Installing the J-Web Package

As an alternative to entering CLI commands, JUNOS also supports a J-Web graphical user interface (GUI). The J-Web user interface allows you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser with Hypertext Transfer Protocol (HTTP) or HTTP over Secure Sockets Layer (HTTPS) enabled.

The J-Web user interface comes standard on J-series Services Routers. It is provided as an optional, licensed software package (jweb package) on M-series and T-series routers. The jweb package is not included in jinstall and jbundle software bundles. It must be installed separately. To install the package on M-series and T-series routers, follow the procedure described in "Upgrading Individual Software Packages" on page 345.

Because the J-Web package is bundled separately from other packages, it is possible to have a version mismatch between J-Web and other JUNOS software packages you have installed.

To check for a version mismatch, use the show system alarms CLI command. The version number must match exactly or a system alarm appears. For example, if you have the 7.4R1.2 jroute package installed and the 7.4R1.1 jweb package installed, an alarm is activated. For more information on the show system alarms command, see the *JUNOS System Basics and Services Command Reference*.

J-Web supports weak (56-bit) encryption by default. This enables international customers to install J-Web and use HTTPS connections for J-Web access. Domestic customers can also install the jcrypto strong encryption package. This package automatically overrides the weak encryption.

Copying a Configuration to a PC Card or LS-120 Floppy Disk

You can copy a router configuration to a PC card or an LS-120 floppy disk from a workstation and then load it onto a router.

To copy a router configuration to a PC card or an LS-120 floppy disk, follow these steps:

1. Insert the PC card or LS-120 floppy disk into a workstation that supports a “DOS/FAT” file system.

2. Mount the PC card or LS-120 floppy disk DOS partition on your UNIX workstation. (This is not necessary for a Windows workstation.)

3. Copy the desired router configuration to the PC card or LS-120 floppy disk as juniper.conf (or juniper.conf.gz, if the configuration is in a compressed format).

4. Unmount the PC card or LS-120 floppy disk from your UNIX workstation. (This is not necessary for a Windows workstation.)

5. Remove the PC card or LS-120 floppy disk.

For information about how to load a configuration from a PC card or LS-120 floppy disk, see, “Reinstalling the JUNOS Software” on page 338.

Chapter 16

Reinstalling Software Using jinstall

If the software becomes damaged, you might want to reinstall it using jinstall. The jinstall package completely reinstalls the software. This package rebuilds the JUNOS file system only but retains configuration information from the previous version. However, logs and other types of auxiliary information may be erased during installation.

To completely reinstall the software using jinstall, follow these steps:

1. Download the software packages you need from the Juniper Networks Support Web site, http://www.juniper.net/support/. Choose either the U.S. and Canada Version or the Worldwide Version.

   To download the software packages, you must have a service contract and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks Web site, https://www.juniper.net/registration/Register.jsp.

NOTE: We recommend that you upgrade and downgrade software packages out-of-band using the console because in-band connections can be lost during the downgrade or upgrade process.

2. Back up the currently running and active file system so that you can recover to a known, stable environment in case something goes wrong with the upgrade:

   ```
   user@host> request system snapshot
   ```

   The root file system is backed up to /altroot, and /config is backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue this command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

3. Copy the jinstall package to the router. You might want to copy them to the /var/tmp directory, which is on the rotating media (hard disk) and is a large file system.

   ```
   user@host> file copy ftp://username:prompt@ftp.hostname.net/
     filename /var/tmp/filename
   ```

4. Add the jinstall package:

```
user@host> request system software add /var/tmp/
jinstall-7.x-package-name-signed.tgz

Installing package
'/var/tmp/jinstall-7.x-package-name-signed.tgz...
Verified jinstall-7.x-package-name-signed.tgz signed by PackageDevelopment_0
Adding jinstall...

WARNING:    This package will load JUNOS 7.x software.
WARNING:    It will save JUNOS configuration files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. It will attempt to preserve dumps
WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.

Saving the config files...
Installing the bootstrap installer...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY.
WARNING:   Use the 'request system reboot' command when software installation
WARNING:   is complete. To abort the installation, do not reboot your system,
WARNING:   instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in
/var/sw/pkg/jinstall-7.x-package-name-signed.tgz...
Saving state for rollback...
```

NOTE: The installation process removes most stored files (except log, juniper.conf, and SSH files) on the router, such as configuration templates and shell scripts. To preserve these files, copy them to another system before upgrading or downgrading the software.

5. If desired, add the optional jweb package. For more information, see "Installing the J-Web Package" on page 346.

6. Reboot the router to load the JUNOS software:

```
root@host> request system reboot
Reboot the system? [yes,no] (no) yes
Shutdown NOW!
```

NOTE: You must reboot to load the JUNOS software. To reboot, issue the request system reboot command when you are done installing the software.

To abort the installation, do not reboot your system; instead, issue the request system software delete jinstall command when you are done installing the software.

All the software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The router then reboots from the primary boot device on which the software was just installed. When the reboot is complete, the router displays the login prompt.

7. Log in and verify the version of software running after the router reboots. Issue the show log message or show version command.

8. After you have upgraded or downgraded the software and are satisfied that the new software is successfully running, issue the request system snapshot command to back up the new software.

   The request system snapshot command causes the root file system to be backed up to /altroot, and /config to be backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

NOTE: You cannot issue the request system software rollback command to return to the previously installed software after using a jinstall package.

To return to the previously installed software, use the jinstall package that corresponds with the previously installed software.

Chapter 17

Upgrading to Release 7.*x* from Release 5.*x*

This chapter explains how to use the JUNOS 7.*x* jinstall package to upgrade to JUNOS Release 7.*x* from Release 5.*x*. After you upgrade to 7.*x*, you can use the jbundle package for subsequent upgrades or downgrades. 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 (from outside the United States).

NOTE: Before upgrading, you should 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:

```
user@host> request system snapshot
```

The installation process removes most stored files on the router, such as configuration templates and shell scripts. The only exceptions are log, juniper.conf, and SSH files. To preserve the stored files, copy them to another system before upgrading or downgrading the router.

To upgrade to or downgrade from JUNOS Release 7.*x*, follow these steps:

1. Download the software packages you need from the Juniper Networks Support Web site, http://www.juniper.net/support/. Choose either the United States and Canada version or the Worldwide version.

   To download the software packages, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks Web site, https://www.juniper.net/registration/Register.jsp.

NOTE: When installing software using the jinstall package, we recommend that you access the router by means of the console serial management port. We recommend that you upgrade and downgrade software packages out-of-band using the console because in-band connections can be lost during the downgrade or upgrade process.

If the console serial management port cannot be used, you can connect to the router through the management Ethernet port fxp0. However, important messages that might be generated during the upgrade process will not be available.

2. Log in to the Juniper Networks Web site.

3. Download the jinstall package to your local host.

4. Copy the jinstall package to the router. You might want to copy it to the /var/tmp directory, which is on the rotating media (hard disk) and is a large file system.

```
user@host> file copy ftp://username:prompt@ftp.hostname.net/
filename /var/tmp/filename
```

NOTE: When upgrading or downgrading JUNOS software, we recommend that you use the validate option with the request system software add command to check that the candidate software is compatible with the current configuration. By default, when you add a package with a different release number, the validation check is done automatically.

We do not recommend using the no-validate option to suppress validation. For more information about this command, see the *JUNOS System Basics and Services Command Reference*.

5. Add the jinstall package:

```
user@host> request system software add /var/tmp/jinstall-package-name
NOTICE: Validating configuration against jinstall-package-name
Checking compatibility with configuration
Initializing...
Using jbase-7.x-package-name
Using /var/tmp/jinstall-7.x-package-name
Verified MD5 checksum of
/var/validate/tmp/jinstall/jinstall-7.x-package-name
Using
/var/validate/tmp/jinstall-signed/jinstall-7.x-package-name
Using /var/validate/tmp/jinstall/jbundle-7.x-package-name
Using /var/validate/tmp/jbundle/jbase-7.x-package-name
Using /var/validate/tmp/jbundle/jkernel-7.x-package-name
Using /var/validate/tmp/jbundle/jcrypto-7.x-package-name
Using /var/validate/tmp/jbundle/jpfe-7.x-package-name
Using /var/validate/tmp/jbundle/jdocs-7.x-package-name
Using /var/validate/tmp/jbundle/jroute-7.x-package-name
Validating against /config/juniper.conf
mgd: commit complete
Validation succeeded
Installing package
'/var/tmp/jinstall-7.x-package-name'...
Verified MD5 checksum of jinstall-7.x-package-name
Auto-deleting old jinstall...
Deleting saved config files...
Deleting bootstrap installer...
Adding jinstall...
```

```
WARNING:    This package will load JUNOS software release-number.
WARNING:    It will save JUNOS configuration files, log files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. This is the pre-installation stage
WARNING:    and all the software is loaded when you reboot the system.

Saving the config files...
Installing the bootstrap installer...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use
WARNING    the 'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in /var/sw/pkg/ jinstall-package-name...
Saving state for rollback...
```

NOTE: A package is installed only if the Message Digest 5 (MD5) checksum within it matches the MD5 hash recorded in its corresponding .md5 file. (For example, jinstall-7.*x*R1-export-signed.tgz contains jinstall-7.*x*R1-export.tgz and jinstall-7.*x*R1-export.tgz.md5. The jinstall-7.*x*R1-export.tgz package will only be installed if it matches the MD5 hash recorded in jinstall-7.*x*R1-export.tgz.md5.) For information about packages, see "Package Names" on page 326

NOTE: The jbundle package cannot be used to upgrade from JUNOS Release 5.*x* to JUNOS Release 7.*x*.

Use the jinstall package to completely reinstall the software. This package rebuilds the file system but retains configuration information from the previous version. retains configuration information from the previous version. (Logs and other types of auxiliary information may be erased during installation.)

6. If desired, add the optional jweb package. For more information, see "Installing the J-Web Package" on page 346.

7. Reboot the router to load the JUNOS software:

   ```
   user@host> request system reboot
   Reboot the system? [yes,no] (no) yes
   Shutdown NOW!
   ```

NOTE: To quit the installation, do not reboot your system; instead, issue the request system software delete jinstall command when software installation completes.

All the software is loaded when you issue the reboot command. Loading the software can take 5 to 10 minutes.

8. 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.

NOTE: After you add a JUNOS 7.*x* jinstall package, you cannot issue the request system software rollback command to return to the previously installed software. To downgrade to another supported release, follow the instructions for upgrading, but replace the JUNOS 7.*x* jinstall package with one labeled for the appropriate release.

9. Log in and verify the version of software running after the router reboots. Issue the show log message or show version command.

10. After you have upgraded or downgraded the software and are satisfied that the new software is successfully running, issue the request system snapshot command to back up the new software.

    The request system snapshot command causes the root file system to be backed up to /altroot, and /config to be backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard disk.

NOTE: After you issue the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

To return to the previous version of JUNOS software, you must reinstall that software following these upgrade or downgrade procedures.

Part 4
System Management

Chapter 18

System Management Overview

The JUNOS software provides a variety of parameters that allow you to configure system management functions, including the router's hostname, address, and domain name; the addresses of Domain Name System (DNS) servers; user login accounts, including user authentication and the root-level user account; time zones and Network Time Protocol (NTP) properties; and properties of the router's auxiliary and console ports.

This chapter discusses the following topics, which provide background information related to configuring system management:

- Specifying IP Addresses, Network Masks, and Prefixes on page 360
- Specifying Filenames and URLs on page 360
- Directories on the Router on page 361
- Tracing and Logging Operations on page 362
- Configuring Protocol Authentication on page 364
- Configuring User Authentication on page 364

Specifying IP Addresses, Network Masks, and Prefixes

Many statements in the JUNOS software configuration include an option to specify an IP address or route prefix. In this manual, this option is represented in one of the following ways:

- *network/prefix-length*—Network portion of the IP address, followed by a slash and the destination prefix length (previously called the subnet mask). For example, 10.0.0.1/8.
- *network*—IP address. For example, 10.0.0.2.
- *destination-prefix/prefix-length*—Route prefix, followed by a slash and the destination prefix length. For example, 192.168.1.10/32.

You enter all IP addresses in classless mode. You can enter the IP address with or without a prefix length, in standard dotted notation (for example, 1.2.3.4), or hexadecimal notation as a 32-bit number in network-byte order (for example, 0x01020304). If you omit any octets, they are assumed to be zero. Specify the prefix length as a decimal number in the range from 1 through 32.

Specifying Filenames and URLs

In some command-line interface (CLI) commands and configuration statements—including file copy, file archive, load, save, set system login user *username* authentication *load-key-file*, and request system software add—you can include a filename. On a routing matrix, you can include chassis information; for example, lcc0, lcc0-re0, or lcc0-re1, as part of the file name. A routing matrix is a multichassis architecture composed of one TX Matrix platform, to which you can connect from one to four T640 routing nodes. For more information about the routing matrix, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852 and the *TX Matrix Platform Hardware Guide*.

You can specify a filename or URL in one of the following ways:

- *filename*—File in the user's current directory on the local flash drive. You can use wildcards to specify multiple source files or a single destination file. Wildcards are not supported in Hypertext Transfer Protocol (HTTP) or FTP. For more information about how to use wildcards, see "Using Wildcards" on page 624.

NOTE: Wildcards are supported only by the file (compare | copy | delete | list | rename | show) commands. When you issue the file show command with a wildcard, it must resolve to one filename.

- *path/filename*—File on the local flash disk.
- /var/*filename* or /var/*path/filename*—File on the local hard disk. You can also specify a file on a local Routing Engine for a specific T640 routing node on a routing matrix:

```
user@host> file delete lcc0-re0:/var/tmp/junk
```

- a:*filename* or a:*path/filename*—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.

- *hostname:/path/filename, hostname:filename, hostname:path/filename,* or scp://*hostname/path/filename*—File on an scp/ssh client. This form is not available in the worldwide version of the JUNOS software. The default path is the user's home directory on the remote system. You can also specify *hostname* as *username@hostname*.

- ftp://*hostname/path/filename*—File on an FTP server. You can also specify *hostname* as *username@hostname* or *username*:*password@hostname*. The default path is the user's home directory. To specify an absolute path, the path must start with %2F; for example, ftp://*hostname*/%2F*path/filename*. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed:

```
user@host> file copy ftp://username@ftp.hostname.net//filename
file copy ftp.hostname.net: Not logged in.
user@host> file copy ftp://username:prompt@ftp.hostname.net//filename
Password for username@ftp.hostname.net:
```

- http://*hostname/path/filename*—File on an HTTP server. You can also specify *hostname* as *username@hostname* or *username*:*password@hostname*. If a password is required and you omit it, you are prompted for it.

- re0:/*path/filename* or re1:/*path/filename*—File on a local Routing Engine.You can also specify a file on a local Routing Engine for a specific T640 routing node on a routing matrix:

```
user@host> show log lcc0-re1:chassisd
```

Directories on the Router

JUNOS software files are stored in the following directories on the router:

- /altconfig—When you back up the currently running and active file system partitions on the router to standby partitions using the request system snapshot command, the /config directory is backed up to /altconfig. Normally, the /config directory is on the flash drive and /altconfig is on the hard disk.

- /altroot—When you back up the currently running and active file system partitions on the router to standby partitions using the request system snapshot command, the root file system (/) is backed up to /altroot. Normally, the root directory is on the flash drive and /altroot is on the hard disk.

- /config—This directory is located on the primary boot device, that is, on the drive from which the router booted (generally the flash drive, device wd0). This directory contains the current operational router configuration and the last three committed configurations, in the files juniper.conf, juniper.conf.1, juniper.conf.2, and juniper.conf.3, respectively.

- /var—This directory is always located on the hard disk (device wd2). It contains the following subdirectories:

  - /var/home—Contains users' home directories, which are created when you create user access accounts. For users using SSH authentication, their .ssh file, which contains their SSH key, is placed in their home directory. When a user saves or loads a configuration file, that file is loaded from their home directory unless the user specifies a full pathname.

  - /var/db/config—Up to six additional previous versions of committed configurations, which are stored in the files juniper.conf.4 through juniper.conf.9.

  - /var/log—Contains system log and tracing files.

  - /var/tmp—Contains core files. The software saves the current core file (0) and the four previous core files, which are numbered from 1 through 4 (from newest to oldest).

Each router ships with removable media (device wfd0) that contains a backup copy of the JUNOS software.

Tracing and Logging Operations

Tracing and logging operations allow you to track events that occur in the router—both normal router operations and error conditions—and to track the packets that are generated by or passed through the router. The results of tracing and logging operations are placed in files in the /var/log directory on the router.

Logging operations use a system logging mechanism similar to the UNIX syslogd utility to record systemwide, high-level operations, such as interfaces going up or down and users logging in to or out of the router. You configure these operations by using the syslog statement at the [edit system] hierarchy level, as described in "Configuring System Log Messages" on page 427, and by using the options statement at the [edit routing-options] hierarchy level, as described in the *JUNOS Routing Protocols Configuration Guide*.

Tracing operations record more detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You configure tracing operations using the traceoptions statement. You can define tracing operations in different portions of the router configuration:

- Global tracing operations—Define tracing for all routing protocols. You define these tracing operations at the [edit routing-options] hierarchy level of the configuration. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

- Protocol-specific tracing operations—Define tracing for a specific routing protocol. You define these tracing operations in the [edit protocol] hierarchy when configuring the individual routing protocol. Protocol-specific tracing operations override any equivalent operations that you specify in the global traceoptions statement. If there are no equivalent operations, they supplement the global tracing options. If you do not specify any protocol-specific tracing, the routing protocol inherits all the global tracing operations.

- Tracing operations within individual routing protocol entities—Some protocols allow you to define more granular tracing operations. For example, in Border Gateway Protocol (BGP), you can configure peer-specific tracing operations. These operations override any equivalent BGP-wide operations or, if there are no equivalents, supplement them. If you do not specify any peer-specific tracing operations, the peers inherit, first, all the BGP-wide tracing operations and, second, the global tracing operations.

- Interface tracing operations—Define tracing for individual router interfaces and for the interface process itself. You define these tracing operations at the [edit interfaces] hierarchy level of the configuration as described in the *JUNOS Network Interfaces Configuration Guide*.

Configuring Protocol Authentication

Some interior gateway protocols (IGPs)—Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and Routing Information Protocol (RIP)—and Resource Reservation Protocol (RSVP) allow you to configure an authentication method and password. Neighboring routers use the password to verify the authenticity of packets sent by the protocol from the router or from a router interface. The following authentication methods are supported:

- Simple authentication (IS-IS, OSPF, and RIP)—Uses a simple text password. The receiving router uses an authentication key (password) to verify the packet. Because the password is included in the transmitted packet, this method of authentication is relatively insecure. We recommend that you *not* use this authentication method.

- MD5 and HMAC-MD5 (IS-IS, OSPF, RIP, and RSVP)—Message Digest 5 (MD5) creates an encoded checksum that is included in the transmitted packet. HMAC-MD5, which combines HMAC authentication with MD5, adds the use of an iterated cryptographic hash function. With both types of authentication, the receiving router uses an authentication key (password) to verify the packet. HMAC-MD5 authentication is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.

In general, authentication passwords are text strings consisting of a maximum of 16 or 255 letters and digits. Characters can include any ASCII strings. If you include spaces in a password, enclose all characters in quotation marks (“ ”).

JUNOS-FIPS has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If JUNOS-FIPS is installed on the router, you cannot configure passwords unless they meet this standard.

Configuring User 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 in to the router.

RADIUS and TACACS+ are authentication methods for validating users who attempt to access the router using telnet. They are both 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 be both a RADIUS and TACACS+ client, and you can also configure authentication passwords in the JUNOS configuration file. You can prioritize the methods to configure the order in which the software tries the different authentication methods when verifying user access.

Chapter 19

System Management Configuration Statements

To configure system management, you can include the following statements in the configuration:

```
system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
            tacplus {
                server {
                    server-address {
                        port port-number;
                        secret password;
                        single-connection;
                        timeout seconds;
                    }
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
```

```
arp {
  passive-learning;
  aging-timer minutes;
}
authentication-order [ authentication-methods ];
backup-router address <destination destination-address>;
building name;
commit synchronize;
(compress-configuration-files | no-compression-configuration-files);
default-address-selection;
dump-device (compact-flash | remove-compact | usb);
diag-port-authentication (encrypted-password “password ” |
      plain-text-password);
domain-name domain-name;
domain-search [domain-list];
host-name host-name;
internet-options address <destination destination-address>;
internet-options {
  path-mtu-discovery;
  source-port upper-limit <upper-limit>;
  source-quench;
}
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
  lata service-area;
  latitude degrees;
  longitude degrees;
  npa-nxx number;
  postal-code postal-code;
  rack number;
  vcoord vertical-coordinate;
}
login {
  announcement text;
  class class-name {
    allow-commands "regular-expression";
    allow-configuration “regular-expression”;
    deny-commands "regular-expression";
    deny-configuration “regular-expression”;
    idle-timeout minutes;
    login-tip;
    permissions [ permissions ];
  }
  message text;
  passwords {
    change-type (set-transitions | character-set);
    format (md5 | sha1 | des);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
  }
  user username {
    full-name complete-name;
    uid uid-value;
    class class-name;
    authentication {
      (encrypted-password "password" | plain-text-password);
      ssh-rsa "public-key";
      ssh-dsa "public-key";
    }
  }
}
login-tip number;
mirror-flash-on-disk;
name-server {
  address;
}
no-redirects;
ntp {
  authentication-key key-number type type value password;
  boot-server (NTP) address;
  broadcast <address> <key key-number> <version value> <ttl value>;
  broadcast-client;
  multicast-client <address>;
  peer address <key key-number> <version value> <prefer>;
  source-address source-address;
  server address <key key-number> <version value> <prefer>;
  trusted-key [ key-numbers ];
}
ports {
  auxiliary {
    type terminal-type;
  }
pic-console-authentication {
  encrypted-password encrypted-password;
  plain-text-password;
  console {
    insecure;
    log-out-on-disconnect;
    type terminal-type;
  }
}
processes {
  adaptive-services (enable | disable) failover failover-option;
  alarm-control (enable | disable) failover failover-option;
  chassis-control (enable | disable) failover failover-option;
  class-of-service (enable | disable) failover failover-option;
  craft-control (enable | disable) failover failover-option;
  disk-monitoring (enable | disable) failover failover-option;
  ecc-error-logging (enable | disable) failover failover-option;
  firewall (enable | disable) failover failover-option;
  inet-process (enable | disable) failover failover-option;
  interface-control (enable | disable) failover failover-option;
  kernel-replication (enable | disable) failover failover-option;
  l2tp-service (enable | disable) failover failover-option;
  link-management (enable | disable) failover failover-option;
  mib-process (enable | disable) failover failover-option;
  network-access-service (enable | disable) failover failover-option;
```

```
    ntp (enable | disable) failover failover-option;
    pgm (enable | disable) failover failover-option;
    pic-services-logging (enable | disable) failover failover-option;
    pppoe (enable | disable) failover failover-option;
    redundancy-device (enable | disable) failover failover-option;
    remote-operations (enable | disable) failover failover-option;
    routing (enable | disable) failover failover-option;
    sampling (enable | disable) failover failover-option;
    service-deployment (enable | disable) failover failover-option;
    snmp (enable | disable) failover failover-option;
    timeout seconds;
    watchdog (enable | disable) failover failover-option;
    web-management (enable | disable) failover failover-option;
    }
}
radius-server server-address {
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
scripts {
    commit {
        allow-transients;
        file filename.xsl {
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        traceoptions {
            file filename <files number> <size size>;
            flag flag;
        }
    }
}
services {
    finger {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    ftp {
        <connection-limit limit>;
        <rate-limit limit>;
    }
```

```
   service-deployment {
      servers server-address {
         port port-number;
      }
      source-address source-address;
   }
   ssh {
      root-login (allow | deny | deny-password);
      protocol-version [v1 v2];
      <connection-limit limit>;
      <rate-limit limit>;
   }
   telnet {
      <connection-limit limit>;
      <rate-limit limit>;
   }
   web-management {
      http {
         port port;
      }
      https {
         local-certificate name;
         port port;
      }
   }
   xnm-clear-text {
      <connection-limit limit>;
      <rate-limit limit>;
   }
   xnm-ssl {
      <connection-limit limit>;
      local-certificate name;
      <rate-limit limit>;
   }
}
static-host-mapping {
   host-name {
      alias [ alias ];
      inet [ address ];
      sysid system-identifier ;
   }
}
syslog {
   archive {
      files number ;
      size size ;
      (world-readable | no-world-readable);
   }
   console {
      facility severity ;
   }
```

```
        file filename {
            facility severity;
            explicit-priority;
            archive {
                files number;
                size size;
                (world-readable | no-world-readable);
            }
        }
        host (hostname | other-routing-engine | scc-master) {
            facility severity;
            explicit-priority;
            facility-override facility;
            log-prefix string;
        }
        source-address source-address;
        time-format (year | millisecond | year millisecond);
        user (username | *) {
            facility severity;
        }
    }
    tacplus-options service-name service-name;
    tacplus-server server-address {
        secret password;
        single-connection;
        source-address source-address;
        timeout seconds;
    }
    time-zone (GMThour-offset | time-zone);
}
```

Chapter 20
Configuring Basic System Management

This chapter discusses the following topics:

Configuring the Router's Name and Addresses

For the router, you can do the following:

- Configuring the Router's Name on page 372
- Mapping the Router's Name to IP Addresses on page 372
- Configuring an ISO System Identifier on page 373

An example shows how to configure a router's name, IP address, and system identifier on page 373.

Configuring the Router's Name

To configure the router's name, include the host-name statement at the [edit system] hierarchy level:

```
[edit system]
host-name host-name;
```

Mapping the Router's Name to IP Addresses

To map a router's hostname to one or more IP addresses, include the inet statement at the [edit system static-host-mapping *host-name*] hierarchy level:

```
[edit system]
static-host-mapping {
    host-name {
       inet [ address ];
       alias [ alias ];
    }
}
```

host-name is the name specified by the host-name statement at the [edit system] hierarchy level.

For each host, you can specify one or more aliases.

Configuring an ISO System Identifier

For Intermediate System-to-Intermediate System (IS-IS) to operate on the router, you must configure a system identifier (system ID). The system identifier is commonly the media access control (MAC) address or the IP address expressed in binary-coded decimal (BCD). For more information, see the *JUNOS Routing Protocols Configuration Guide*.

To configure an International Organization for Standardization (ISO) system ID, include the sysid statement at the [edit system static-host-mapping *host-name*] hierarchy level:

```
[edit system]
static-host-mapping {
    host-name {
      sysid system-identifier ;
    }
}
```

host-name is the name specified by the host-name statement at the [edit system] hierarchy level.

system-identifier is the ISO system identifier. It is the 6-byte system ID portion of the IS-IS Network Service Access Point (NSAP). We recommend that you use the host's IP address represented in BCD format. For example, the IP address 192.168.1.77 is 1921.6800.1077 in BCD.

Example: Configuring a Router's Name, IP Address, and System ID

Configure the router's name, map the name to an IP address and alias, and configure a system identifier:

```
[edit]
user@host# set system host-name router-sj1
[edit]
user@host# set system static-host-mapping router-sj1 inet 192.168.1.77
[edit]
user@host# set system static-host-mapping router-sj1 alias sj1
[edit]
user@host# set system static-host-mapping router-sj1 sysid 1921.6800.1077
[edit]
user@host# show
system {
    host-name router-sj1;
    static-host-mapping {
      router-sj1 {
        inet 192.168.1.77;
        alias sj1;
        sysid 1921.6800.1077;
      }
    }
}
```

Configuring the Router's Domain Name

For each router, you should configure the name of the domain in which the router is located. This is the default domain name that is appended to hostnames that are not fully qualified. To configure the domain name, include the domain-name statement at the [edit system] hierarchy level:

```
[edit system]
domain-name domain-name;
```

Example: Configuring the Router's Domain Name

Configure the router's domain name:

```
[edit]
user@host# set system domain-name company.net
[edit]
user@host# show
system {
    domain-name company.net;
}
```

Configuring Which Domains to Search

If your router is included in several different domains, you can configure those domain names to be searched.

To configure more than one domain to be searched, include the domain-search statement at the [edit system] hierarchy level:

```
[edit system]
domain-search [domain-list];
```

The domain list can contain up to 6 domain names, with a total of up to 256 characters.

Example: Configuring Which Domains to Search

Configure two domains to be searched:

```
[edit system]
domain-search [domainone.net domainonealternate.com]
```

Configuring a DNS Name Server

To have the router resolve hostnames into addresses, you must configure one or more Domain Name System (DNS) name servers by including the name-server statement at the [edit system] hierarchy level:

```
[edit system]
name-server {
    address;
}
```

Example: Configuring a DNS Name Server

Configure two DNS name servers:

```
[edit]
user@host# set system name-server 192.168.1.253
[edit]
user@host# set system name-server 192.168.1.254
[edit]
user@host# show
system {
    name server {
      192.168.1.253;
      192.168.1.254;
    }
}
```

Configuring a Backup Router

When the router is booting, the routing protocol process (rpd) is not running; therefore, the router has no static or default routes. To allow the router to boot and to ensure that the router is reachable over the network if the routing protocol process fails to start properly, you configure a backup router (running IP version 4 [IPv4] or IP version 6 [IPv6]), which is a router that is directly connected to the local router (that is, on the same subnet).

To configure a backup router running IPv4, include the backup-router statement at the [edit system] hierarchy level:

```
[edit system]
backup-router address <destination destination-address>;
```

To configure a backup router running IPv6, include the inet6-backup-router statement at the [edit system] hierarchy level:

```
[edit system]
inet6-backup-router address <destination destination-address>;
```

By default, all hosts (default route) are reachable through the backup router. To eliminate the risk of installing a default route in the forwarding table, include the destination option, specifying an address that is reachable through the backup router. Specify the address in the format *network/mask-length* so that the entire network is reachable through the backup router.

When the routing protocols start, the address of the backup router is removed from the local routing and forwarding tables. To have the address remain in these tables, configure a static route for that address by including the static statement at the [edit routing-options] hierarchy level.

Example: Configuring a Backup Router Running IPv4

Configure a backup router and have its address remain in the routing and forwarding tables:

```
[edit]
system {
    backup-router 192.168.1.254 destination 208.197.1.0/24;
}
routing-options {
    static {
      route 208.197.1.0/24 {
        next-hop 192.168.1.254;
        retain;
      }
    }
}
```

Example: Configuring a Backup Router Running IPv6

Configure a backup router running IPv6 and have its address remain in the routing and forwarding tables:

```
[edit]
system {
    backup-router 8:3::1 destination abcd::/48;
}
routing-options {
    rib inet6.0 {
        static {
            route abcd::/48 {
                next-hop 8:3::1;
                retain;
            }
        }
    }
}
```

Configuring Flash Disk Mirroring

You can direct the hard disk to automatically mirror the contents of the compact flash. When you include the mirror-flash-on-disk statement, the hard disk maintains a synchronized mirror copy of the compact-flash contents. Data written to the compact flash is simultaneously updated in the mirrored copy of the hard disk. If the flash drive fails to read data, the hard disk automatically retrieves its mirrored copy of the flash disk.

CAUTION: We recommend that you disable flash disk mirroring when you upgrade or downgrade the router.

You cannot issue the request system snapshot command while flash disk mirroring is enabled.

To configure the mirroring of the compact flash to the hard disk, include the mirror-flash-on-disk statement at the [edit system] hierarchy level:

```
[edit system]
mirror-flash-on-disk;
```

NOTE: After you have enabled or disabled the mirror-flash-on-disk statement, you must reboot the router for your changes to take effect. To reboot, issue the request system reboot command.

Configuring the System Location

To configure the physical location of the system, include the location statement at the [edit system] hierarchy level:

```
[edit system]
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
```

- altitude *feet*—Number of feet above sea level.
- building *name*— The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").
- country-code *code*—Two-letter country code.
- floor *number*—Floor in the building.
- hcoord *horizontal-coordinate*—Bellcore Horizontal Coordinate.
- lata *service-area*—Long-distance service area.
- latitude *degrees*—Latitude in degree format.
- longitude *degrees*—Longitude in degree format.
- npa-nxx *number*—First six digits of the phone number (area code and exchange).
- postal-code *postal-code*—Postal code.
- rack *number*—Rack number.
- vcoord *vertical-coordinate*—Bellcore Vertical Coordinate.

Configuring the Root Password

The JUNOS software is preinstalled on the router. When the router is powered on, it is ready to be configured. Initially, you log in to the router as the user “root” with no password. After you log in, you should configure the root (superuser) password by including the root-authentication statement at the [edit system] hierarchy level:

```
[edit system]
root-authentication {
    (encrypted-password "password"| plain-text-password);
    ssh-dsa "public-key";
    ssh-rsa "public-key";
}
```

If you configure the plain-text-password option, you are prompted to enter and confirm the password:

```
[edit system]
user@host# set root-authentication plain-text-password
New password: type password here
Retype new password: retype password here
```

For information about how to create plain-text passwords, see “Specifying Plain-Text Passwords” on page 17.

To load an SSH key file, enter the load-key-file command. This command loads RSA (SSH version 1) and DSA (SSH version 2) public keys. You can also configure a user to use ssh-rsa and ssh-dsa keys.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the load-key-file statement. To view the SSH keys entries, use the configuration mode show command. For example:

```
[edit system]
user@host# set root-authentication load-key-file my-host:.ssh/identity.pub
.file.19692              |          0 KB |   0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@host# show
root-authentication {
ssh-rsa "1024 35 9727638204084251055468226757249864241630322207
404962528390382038690141584534964170019610608358722961563475
784918273603361276441874265946893207739108344810126831259577
226254616679992783161235004386609158662838224897467326056611
921814895398139655615637862119403276878065381696020274916416
3735913269396344008443 boojum@juniper.net"; # SECRET-DATA
}
```

JUNOS-FIPS has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If JUNOS-FIPS is installed on the router, you cannot configure passwords unless they meet this standard.

Example: Configuring the Root Password

Configure an encrypted password:

```
[edit]
user@host# set system root-authentication encrypted-password
"$1$14c5.$sBopasddsdfs0"
[edit]
user@host# show
system {
    root-authentication {
      encrypted-password "$1$14c5.$sBopasddsdfs0";
    }
}
```

Configure a plain-text password:

```
[edit]
user@host# set system root-authentication plain-text-password
New password: type root password
Retype new password: retype root password
[edit]
user@host# show
system {
    root-authentication {
      encrypted-password "$1$14c5.$sBopasddsdfs0";
    }
}
```

Configuring Special Requirements for Plain-Text Passwords

The JUNOS software has special requirements when you create plain-text passwords on a routing platform. Table 14 on page 381 shows the default requirements.

Table 14: Special Requirements for Plain-Text Passwords

JUNOS Software	JUNOS-FIPS
The password must be between 6 and 128 characters long.	FIPS passwords must be between 10 and 20 characters in length
You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.	You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
Valid passwords must contain at least one change of case or character class.	Passwords must use at least three of the five defined character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters).

JUNOS software supports the following five character classes for plain text passwords:

- Lowercase letters
- Uppercase letters
- Numbers
- Punctuation
- Other special characters

Control characters are not recommended.

To change the requirements for plain-text passwords, include the passwords statement at the [edit system login] hierarchy level:

```
[edit system login]
passwords {
    change-type (set-transitions | character-set);
    format (md5 | sha1 | des);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
}
```

These statements apply to plain-text passwords only, not encrypted passwords.

The change-type statement specifies whether the password is checked for the following:

- The total number of character sets used (character-set)
- The total number of character set changes (set-transitions)

For example, the following password:

```
MyPassWd@2
```

has four character sets (uppercase letters, lowercase letters, special characters, and numbers) and seven character set changes (M–y, y–P, P–a, s–W, W–d, d–@, and @–2).

The change-type statement is optional. If change-type is omitted, JUNOS-FIPS plain-text passwords are checked for character sets and JUNOS plain-text passwords are checked for character set changes.

The minimum-changes statement specifies how many character sets or character set changes are required for the password. This statement is optional. If minimum-changes is not specified, character sets are not checked for JUNOS software. If the change-type statement is configured for character-set, then minimum-changes must be 5 or less because JUNOS software only supports 5 character sets.

The format statement specifies the hash algorithm (md5, sha1 or des) for authenticating plain-text passwords. This statement is optional. For JUNOS software, the default format is md5. For JUNOS-FIPS, only sha1 is supported.

The maximum-length statement specifies the maximum number of characters allowed in a password. This statement is optional. By default JUNOS passwords have no maximum and JUNOS-FIPS passwords must be 20 characters or less. The range is from 20 to 128 characters.

The minimum-length statement specifies the minimum number of characters required for a password. This statement is optional. By default JUNOS passwords must be at least 6 characters long and JUNOS-FIPS passwords must be at least 10 characters long. The range is from 6 to 20 characters.

Changes to password requirements do not take effect until the configuration is committed. When requirements change, only newly created, plain-text passwords are checked; existing passwords are not checked against the new requirements.

The default configuration for JUNOS plain-text passwords is:

```
[edit system login]
passwords {
    change-type character-sets;
    format md5;
    minimum-changes 1;
    minimum-length 6;
}
```

The default configuration for JUNOS-FIPS plain-text passwords is:

```
[edit system login]
passwords {
    change-type set-transitions;
    format sha1;
    maximum-length 20;
    minimum-changes 3;
    minimum-length 10;
}
```

Example: Configuring Special Requirements for Plain-Text Passwords

In this example, the minimum password length is set to 12 characters and the maximum length is set to 22 characters.

```
[edit system login]
passwords {
    minimum-length 12;
    maximum-length 22;
}
```

Configuring Multiple Routing Engines to Synchronize Configurations Automatically

If your router has multiple Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the others by issuing the commit synchronize command.

To make the Routing Engines synchronize automatically whenever a configuration is committed, include the commit synchronize statement at the [edit system] hierarchy level:

```
[edit system]
commit synchronize;
```

The Routing Engine on which you execute the commit command (requesting Routing Engine) copies and loads its candidate configuration to the other (responding) Routing Engines. All Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on all Routing Engines.

Compressing the Current Configuration File

By default, the current operational configuration file is compressed, and is stored in the file juniper.conf.gz, in the /config file system, along with the last three committed versions of the configuration. If you have large networks, the current configuration file might exceed the available space in the /config file system. Compressing the current configuration file allows the file to fit in the file system, typically reducing the size of the file by 90 percent. You might want to compress your current operation configuration files when they reach 3 megabytes (MB) in size.

When you compress the current configuration file, the names of the router's configuration files change. To determine the size of the files in the /config file system, issue the file list /config detail command.

NOTE: We recommend that you use the default setting (compress the router configuration files) to minimize the amount of disk space that they require.

If you do not want to compress the current operational configuration file, include the no-compress-configuration-files statement at the [edit system] hierarchy level:

```
[edit system]
no-compression-configuration-files;
```

Commit the current configuration file to include the no-compression-configuration-files statement. Commit the configuration again to uncompress the current configuration file:

```
[edit system]
user@host# set no-compression-configuration-files
user@host# commit
commit complete
user@host# commit
commit complete
```

To compress the current configuration file, include the compress-configuration-files statement at the [edit system] hierarchy level:

```
[edit system]
compress-configuration-files;
```

Commit the current configuration file to include the compression-configuration-files statement. Commit the configuration again to compress the current configuration file:

```
[edit system]
user@host# set compress-configuration-files
user@host# commit
commit complete
user@host# commit
commit complete
```

For more information about how configurations are stored, see “How the Configuration Is Stored” on page 221.

Chapter 21

Configuring System Authentication

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.

When configuring system authentication, you can do the following:

- Configuring RADIUS Authentication on page 386
- Configuring TACACS+ Authentication on page 388
- Specifying a Source Address for RADIUS and TACACS+ Servers on page 390
- Configuring the Same Authentication Service for Multiple TACACS+ Servers on page 391
- Configuring Template Accounts for RADIUS and TACACS+ Authentication on page 392
- Configuring the Authentication Order on page 395

For examples of configuring system authentication, see "Examples: Configuring System Authentication" on page 396.

Configuring RADIUS Authentication

To use RADIUS authentication on the router, configure information about one or more RADIUS servers on the network by including the radius-server statement at the [edit system] hierarchy level:

```
[edit system]
radius-server server-address {
    accounting-port number;
    port number;
    retry number;
    routing-instance routing-instance-name;
    secret password;
    timeout seconds;
}
```

server-address is the address of the RADIUS server.

You can specify a port number on which to contact the RADIUS server. By default, port number 1812 is used (as specified in RFC 2865).

You must specify a password in the secret statement. Passwords can contain spaces. The secret used by the local router must match that used by the server.

Optionally, you can specify the amount of time that the local router waits to receive a response from a RADIUS server (in the timeout statement) and the number of times that the router attempts to contact a RADIUS authentication server (in the retry statement). By default, the router waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds. By default, the router retries connecting to the server 3 times. You can configure this to be a value in the range from 1 through 10 times.

Optionally, in the routing-instance statement, you can configure a routing instance used to send RADIUS packets to the RADIUS server.

To configure multiple RADIUS servers, include multiple radius-server statements.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the user statement at the [edit system login] hierarchy level, as described in "Configuring Template Accounts for RADIUS and TACACS+ Authentication" on page 392.

You can also configure RADIUS authentication at the [edit access] and [edit access profile] hierarchy level. The JUNOS software uses the following search order to determine which set of servers are used for authentication:

```
[edit access profile profile-name radius-server],
[edit access radius-server server-address],
[edit system radius-server server-address]
```

For more information, see "Configuring Access" on page 647.

Configuring Juniper Networks-Specific RADIUS Attributes

The JUNOS software supports the configuration of Juniper Networks-specific RADIUS attributes. These attributes are known as vendor-specific attributes and are described in RFC 2138, *Remote Authentication Dial In User Service* (*RADIUS*). These Juniper Networks-specific attributes are encapsulated in a RADIUS vendor-specific attribute with the vendor ID set to the Juniper Networks ID number, 2636. Table 15 lists the Juniper Networks-specific attributes you can configure.

Table 15: Juniper Networks-Specific RADIUS Attributes

Name	Description	Type	Length	String
Juniper-Local-User-Name	Indicates the name of the user template used by this user when logging in to a device. This attribute is used only in Access-Accept packets.	1	≥3	One or more octets containing printable ASCII characters.
Juniper-Allow-Commands	Contains an extended regular expression that allows the user to run operational mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	2	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.
Juniper-Deny-Commands	Contains an extended regular expression that denies the user permission to run operation mode commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	3	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.
Juniper-Allow-Configuration	Contains an extended regular expression that allows the user to run configuration mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	4	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.
Juniper-Deny-Configuration	Contains an extended regular expression that denies the user permission to run configuration commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	5	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.

Configuring TACACS+ Authentication

To use TACACS+ authentication on the router, configure information about one or more TACACS+ servers on the network by including the tacplus-server statement at the [edit system] hierarchy level:

```
[edit system]
tacplus-server server-address {
    port port-number;
    secret password;
    single-connection;
    timeout seconds;
}
```

server-address is the address of the TACACS+ server.

port-number is the TACACS+ server port number.

You must specify a secret (password) that the local router passes to the TACACS+ client by including the secret statement. Secrets can contain spaces. The secret used by the local router must match that used by the server.

You can optionally specify the length of time that the local router waits to receive a response from a TACACS+ server by including the timeout statement. By default, the router waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds.

You can optionally have the software maintain one open Transmission Control Protocol (TCP) connection to the server for multiple requests, rather than opening a connection for each connection attempt by including the single-connection statement.

NOTE: Early versions of the TACACS+ server do not support the single-connection option. If you specify this option and the server does not support it, the JUNOS software will be unable to communicate with that TACACS+ server.

To configure multiple TACACS+ servers, include multiple tacplus-server statements.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the user statement at the [edit system login] hierarchy level, as described in "Configuring Template Accounts for RADIUS and TACACS+ Authentication" on page 392.

Configuring Juniper Networks-Specific TACACS+ Attributes

The TACACS attributes listed in Table 16 are specific to Juniper Networks. They are specified in the TACACS+ server configuration file on a per-user basis. The JUNOS software retrieves these attributes through an authorization request of the TACACS+ server after authenticating a user. You do not need to configure these attributes to run JUNOS with TACACS+.

To specify these attributes, include a service statement of the following form in the TACACS+ server configuration file:

```
service = junos-exec {
    local-user-name = <username-local-to-router>
    allow-commands = "<allow-commands-regexp>"
    allow-configuration = "<allow-configuration-regexp>"

    deny-commands = "<deny-commands-regexp>"
    deny-configuration = "<deny-configuration-regexp>"
}
```

This service statement can appear in a user or group statement.

Table 16: Juniper Networks-Specific TACACS+ Attributes

Name	Description	Length	String
local-user-name	Indicates the name of the user template used by this user when logging in to a device.	≥3	One or more octets containing printable ASCII characters.
allow-commands	Contains an extended regular expression that allows the user to run operational mode commands in addition to those commands authorized by the user’s login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.
allow-configuration	Contains an extended regular expression that allows the user to run configuration mode commands in addition to those commands authorized by the user’s login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.
deny-commands	Contains an extended regular expression that denies the user permission to run operational mode commands authorized by the user’s login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.
deny-configuration	Contains an extended regular expression that denies the user permission to run configuration mode commands authorized by the user’s login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression.

Specifying a Source Address for RADIUS and TACACS+ Servers

You can specify which source address the JUNOS software uses when accessing your network to contact an external TACACS+ or RADIUS server for authentication. You can also specify which source address the JUNOS software uses when contacting a TACACS+ server for sending accounting information.

To specify a source address for a TACACS+ server for authentication, include the source-address statement at the [edit system tacplus-server *server-address*] hierarchy level:

```
[edit system tacplus-server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router interfaces.

To specify a source address for a TACACS+ server for system accounting, include the source-address statement at the [edit system accounting destination tacplus server *server-address*] hierarchy level:

```
[edit system accounting destination tacplus server server-address]
source-address source-address;
```

source-address is a valid IP address configured on one of the router interfaces.

To specify a source address for a RADIUS+ server, include the source-address statement at the [edit system radius-server *server-address*] hierarchy level:

```
[edit system radius-server server-address]:
source-address source-address;
```

source-address is a valid IP address configured on one of the router interfaces.

Configuring the Same Authentication Service for Multiple TACACS+ Servers

To configure the same authentication service for multiple TACACS+ servers, include statements at the [edit system tacplus-server] and [edit system tacplus-options] hierarchy levels. For information about how to configure a TACACS+ server at the [edit system tacplus-server] hierarchy level, see, “Configuring TACACS+ Authentication” on page 388.

To assign the same authentication service to multiple TACACS+ servers, include the service-name statement at the [edit system tacplus-options] hierarchy level:

```
[edit system tacplus-options]
service-name service-name;
```

service-name is the name of the authentication service. By default, the service name is set to junos-exec.

Example: Configuring Multiple TACACS+ Servers

Configure the same authentication service for multiple TACACS+ servers:

```
[edit system]
tacplus-server {
    2.2.2.2 secret "$9$2dgoJGDiqP5ZG9A"; ## SECRET-DATA
    3.3.3.3 secret "$9$2dgoJGDiqP5ZG9A.";## SECRET-DATA
}
tacplus-options {
    service-name bob;
}
```

Configuring Template Accounts for RADIUS and TACACS+ Authentication

When you use local password authentication, you must create a local user account for every user who wants to access the system. However, when you are using RADIUS or TACACS+ authentication, you can create single accounts (for authorization purposes) that are shared by a set of users. You create these accounts using the remote and local user template accounts. When a user is using a template account, the command-line interface (CLI) username is the login name; however, the privileges, file ownership, and effective user ID are inherited from the template account.

This section discusses the following topics:

- Using Remote Template Accounts on page 392
- Using Local User Template Accounts on page 393

Using Remote Template Accounts

By default, the JUNOS software uses the remote template accounts when:

- The authenticated user does not exist locally on the router
- The authenticated user's record in the authentication server specifies local user, or the specified local user does not exist locally on the router

To configure the remote template account, include the user remote statement at the [edit system login] hierarchy level and specify the privileges you want to grant to remote users:

```
[edit system login]
user remote {
    full-name "All remote users";
    uid uid-value;
    class class-name;
}
```

To configure different access privileges for users who share the remote template account, include the allow-commands and deny-commands commands in the authentication server configuration file. For information about how to define access privileges on the authentication server, see "Configuring Juniper Networks-Specific RADIUS Attributes" on page 387 and "Configuring Juniper Networks-Specific TACACS+ Attributes" on page 389.

For information about creating user accounts, see "Configuring User Accounts" on page 413. For an example of how to configure a template account, see "Examples: Configuring System Authentication" on page 396.

Using Local User 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 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 JUNOS software, 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 JUNOS software 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.

To configure different access privileges for users who share the local user template account, include the allow-commands and deny-commands commands in the authentication server configuration file. For information about how to configure access privileges on the authentication server, see "Configuring Juniper Networks-Specific RADIUS Attributes" on page 387 and "Configuring Juniper Networks-Specific TACACS+ Attributes" on page 389.

For information about creating user accounts, see "Configuring User Accounts" on page 413. For an example of how to configure a template account, see "Examples: Configuring System Authentication" on page 396.

To configure a local user template, include the user *local-username* statement at the [edit system login] hierarchy level and specify the privileges you want to grant to the local users to whom the template applies:

```
[edit system login]
user local-username {
    full-name "Local user account";
    uid uid-value;
    class class-name;
}
```

Using Local User Template Example

In this example, you configure the sales and engineering local user templates:

```
[edit]
system {
    login {
      user sales {
        uid uid-value;
        class class-name;
      }
      user engineering {
        uid uid-value;
        class class-name;
      }
    }
}
```

Now you configure users on the TACACS+ authentication server:

```
user = simon {
    ...
    service = junos-exec {
      local-user-name = sales
      allow-commands = "configure"
      deny-commands = "shutdown"
    }
}
user = rob {
    ...
    service = junos-exec {
      local-user-name = sales
      allow-commands = "(request system) | (show rip neighbor)"
    deny-commands = "<^clear"
    }
}
user = harold {
    ...
    service = junos-exec {
      local-user-name = engineering
      allow-commands = "monitor | help | show | ping | traceroute"
      deny-commands = "configure"
    }
}
user = jim {
    ...
    service = junos-exec {
      local-user-name = engineering
      allow-commands = "show bgp neighbor"
      deny-commands = "telnet | ssh"
    }
}
```

When the login users Simon and Rob are authenticated, they use the sales local user template. When login users Harold and Jim are authenticated, they use the engineering local user template.

NOTE: Permission bits override allow and deny commands.

Configuring the Authentication Order

If you configure the router to be both a RADIUS and TACACS+ client (by including the radius-server and tacplus-server statements), you can prioritize the order in which the software tries the different authentication methods when verifying that a user can access the router. For each login attempt, the JUNOS software tries the authentication methods in order, starting with the first one, until the password matches.

To configure the authentication order, include the authentication-order statement at the [edit system] hierarchy level:

```
[edit system]
authentication-order [ authentication-methods ];
```

In *authentication-methods*, specify one or more of the following in the preferred order, from first tried to last tried:

- radius—Verify the user using RADIUS authentication services.
- tacplus—Verify the user using TACACS+ authentication services.
- password—Verify the user using the password configured for the user with the authentication statement at the [edit system login user] hierarchy level.

If you do not include the authentication-order statement, users are verified based on their configured passwords.

Example: Removing an Order Set from the Authentication Order

Delete the radius statement from the authentication order:

```
[edit system]
user@host# delete authentication-order radius
```

For more information about how to remove a statement from the configuration, see “Removing a Statement from the Configuration” on page 246.

Example: Inserting an Order Set in the Authentication Order

Insert the tacplus statement after the radius statement:

```
[edit system]
user@host# insert authentication-order tacplus after radius
```

For more information about how to modify a portion of the configuration in which the statement order matters, see “Inserting a New Identifier” on page 251.

Examples: Configuring System Authentication

The following example allows logins only by the individual user Philip, and by users who have been authenticated by a remote RADIUS server. If a user logs in and is not authenticated by the RADIUS server, the user is denied access to the router. However, if the RADIUS server is not available, the user's login name has a local password, and the user enters that password, the user is authenticated (using the password authentication method) and allowed access to the router. For more information about the password authentication method, see "Example: Defaulting to Local User Password Authentication, RADIUS" on page 398.

When Philip tries to log in to the system, if the RADIUS server authenticates him, he is given access and privileges for the super-user class. Local accounts are not configured for other users. When they log in to the system and the RADIUS server authenticates them, they are given access using the same user ID (UID) 9999 and the same privileges for the operator class.

```
[edit]
system {
    authentication-order radius;
    login {
        user philip {
            full-name "Philip";
            uid 1001;
            class super-user;
        user remote {
            full-name "All remote users";
            uid 9999;
            class operator;
        }
    }
}
```

NOTE: For authorization purposes, you can use a template account to create a single account that can be shared by a set of users at the same time. For example, when you create a remote template account, a set of remote users can concurrently share a single UID. For more information about template accounts, see "Configuring Template Accounts for RADIUS and TACACS+ Authentication" on page 392.

Configuring a single remote user template account requires that all users without individual configuration entries share the same class and UID. When you are using RADIUS and telnet or RADIUS and SSH together, you can specify a different template user other than the remote user.

To configure an alternate template user, specify the “User-Name” parameter returned in the RADIUS authentication response packet. Not all RADIUS servers allow you to change this parameter. The following shows a sample JUNOS configuration:

```
[edit]
system {
    authentication-order radius;
    login {
      user philip {
        full-name "Philip";
        uid 1001;
        class super-user;
      }
      user operator {
        full-name "All operators";
        uid 9990;
        class operator;
      }
      user remote {
        full-name "All remote users";
        uid 9999;
        class read-only;
      }
    }
}
```

Assume your RADIUS server is configured with the following information:

- User Philip with password “olympia”
- User Alexander with password “bucephalus” and username “operator”
- User Darius with password “redhead” and username “operator”
- User Roxane with password “athena”

Philip would be given access as a superuser (super-user) because he has his own local user account. Alexander and Darius share UID 9990 and have access as operators. Roxane has no template-user override, so she shares access with all the other remote users, getting read-only access.

Using the Local User Fallback Mechanism

The JUNOS software provides a local user fallback mechanism (password authentication method) that enables users to log in to the router when no TACACS+ or RADIUS authentication servers is available. The following examples illustrate how this mechanism works.

Example: Inserting Password into the Authentication Order

If you specify the following authentication order:

```
[edit]
system authentication-order [tacplus password];
```

the JUNOS software first uses the authentication method TACACS+ to authenticate users when they attempt to log in to the router. The authentication servers are tried in the order specified at the [edit system tacplus-server] hierarchy level. If no TACACS+ authentication server is available, the JUNOS software will try the next authentication method listed, password. The password option also allows users that fail to authenticate with TACACS+ to log in to the router by means of UNIX password authentication.

In effect, this configuration provides a local user fallback mechanism (traditional UNIX password) when all TACACS+ servers are unavailable, but does not restrict authentication to TACACS+ authentication only (all users will be able to try the traditional UNIX password as well).

Example: Defaulting to Local User Password Authentication, TACACS +

If you specify the following authentication order:

```
[edit]
system authentication-order tacplus;
```

and none of the TACACS+ servers configured at the [edit system tacplus-server] hierarchy are available, the JUNOS software will try to use the password authentication method. If a TACACS+ server is available, the JUNOS software will not try to use the password authentication method.

Example: Defaulting to Local User Password Authentication, RADIUS

If you specify the following authentication order:

```
[edit]
system authentication-order radius;
```

and none of the RADIUS servers configured at the [edit system radius-server] hierarchy level are available, the JUNOS software will try to use the password authentication method. If a RADIUS server is available, the JUNOS software will not try to use the password authentication method.

Example: Defaulting to Local User Password Authentication, TACACS + and RADIUS

If you specify the following authentication order:

```
[edit]
system authentication-order [tacplus radius];
```

and no TACACS+ authentication server is available but at least one RADIUS authentication server responds (but fails to authenticate), the JUNOS software will try to use the local user fallback mechanism (password authentication method).

NOTE: If any one authentication method (RADIUS or TACACS+) fails to communicate with all of its configured servers, the JUNOS software will use the local user fallback mechanism (password authentication method).

Chapter 22

Configuring User Access

To configure user access, you do the following:

- Defining Login Classes on page 401
- Configuring User Accounts on page 413
- JUNOS-FIPS Crypto Officer and User Accounts on page 415

For information about how to configure user access by means of SSH, see "Configuring SSH Service" on page 477.

Defining Login Classes

All users who can log in to the router must be in a login class. With login classes, you define the following:

- Access privileges users have when they are logged in to the router
- Commands and statements that users can and cannot specify
- How long a login session can be idle before it times out and the user is logged out.

You can define any number of login classes. You then apply one login class to an individual user account, as described in "Configuring User Accounts" on page 413.

To define a login class and its access privileges, include the class statement at the [edit system login] hierarchy level:

```
[edit system login]
class class-name {
    allow-commands "regular-expression";
    allow-configuration "regular-expression";
    deny-commands "regular-expression";
    deny-configuration "regular-expression";
    idle-timeout minutes;
    no-world-readable;
    permissions [ permissions ];
}
```

Use *class-name* to name the login class. The software contains a few predefined login classes, which are listed in Table 18 on page 404. The predefined login classes cannot be modified.

NOTE: You cannot modify a predefined login class name. If you issue the set command on a predefined class name, the JUNOS software will append -local to the login class name. The following message also appears:

```
warning: '<class-name>' is a predefined class name; changing to '<class-name>-local'
```

NOTE: You cannot issue the rename or copy command on a predefined login class. Doing so results in the following error message:

```
error: target '<classname>' is a predefined class
```

For each login class, you can do the following:

- Configuring Access Privilege Levels on page 402
- Denying or Allowing Individual Commands on page 405
- Configuring the Timeout Value for Idle Login Sessions on page 412
- Configuring Tips on page 412

Configuring Access Privilege Levels

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 privilege level for each command and statement is listed in the summary chapter of the part in which that command or statement is described. The access privileges for each login class are defined by one or more *permission bits*.

To configure access privilege levels, include the permissions statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
permissions [ permissions ];
```

permissions specifies one or more of the permission bits listed in Table 17. Permission bits are not cumulative, so for each class list all the bits needed, including view to display information and configure to enter configuration mode. 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 17: Login Class Permission Bits

Permission Bit	Description
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).
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 level).
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.

Permission Bit	Description
security	Can view security configuration in configuration mode and with the show configuration operational mode command.
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 Simple Network Management Protocol (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 18: Default System Login Classes

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

Example: Configuring Access Privilege Levels

Create two access privilege classes on the router, one for configuring and viewing user accounts only and the second for configuring and viewing SNMP parameters only:

```
[edit]
system {
    login {
      class user-accounts {
        permissions [ configure admin admin-control ];
      }
      class network-mgmt {
        permissions [ configure snmp snmp-control ];
      }
    }
}
```

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 would otherwise be permitted or not allowed by a privilege level specified in the permissions statement. For information about CLI commands, see "CLI Overview" on page 175.

NOTE: The all login class permission bits take precedence over extended regular expressions when a user issues the rollback command.

Users cannot issue the load override command when specifying an extended regular expression. Users can only issue the merge, replace, and patch configuration commands.

This section describes how to define a user's access privileges to individual operational and configuration mode commands. It contains the following topics:

- Specifying Operational Mode Commands on page 405
- Specifying Configuration Mode Commands on page 409

Specifying Operational Mode Commands

You can specify extended regular expressions with the allow-commands and deny-commands statements to define a user's access privileges to individual operational commands. Doing so takes precedence over login class permission bits set for a user. You can include one deny-commands and one allow-commands statement in each login class.

To explicitly allow an individual operational mode command that would otherwise be denied, include the allow-commands statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
allow-commands "regular-expression";
```

To explicitly deny an individual operational mode command that would otherwise be allowed, include the deny-commands statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
deny-commands "regular-expression" ;
```

If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Regular expressions are not case-sensitive.

Use extended regular expressions to specify which operational mode commands are denied or allowed. You specify these regular expressions in the allow-commands and deny-commands statements at the [edit system login class] hierarchy level, or by specifying JUNOS-specific attributes in your TACACS+ or RADIUS authentication server configuration. You must specify that these regular expressions are sent as the value of Juniper vendor-specific attributes. If regular expressions are received during TACACS+ or RADIUS authentication, they override any regular expressions configured on the local router. For information about TACACS+ or RADIUS authentication, see “Configuring User Access” on page 401.

Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2. Table 19 lists common regular expression operators.

Table 19: Common Regular Expression Operators to Allow or Deny Operational Mode Commands

Operator	Match...
|	One of the two terms on either side of the pipe.
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces$" means that the user can issue the show interfaces command but cannot issue show interfaces detail or show interfaces extensive.
[ ]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
( )	A group of commands, indicating an expression to be evaluated; the result is then evaluated as part of the overall expression.

If a regular expression contains a syntax error, authentication fails and the user cannot log in. If a regular expression does not contain any operators, all varieties of the command are allowed. For example, if the following statement is included in the configuration, the user can issue the commands show interfaces detail and show interfaces extensive in addition to showing an individual interface:

```
allow-commands "show interfaces"
```

Example 1: Defining Access Privileges to Individual Operational Mode Commands

The following examples define user access privileges to individual operational mode commands.

If the following statement is included in the configuration and the user does not have the configure login class permission bit, the user can enter configuration mode:

```
[edit system login class class-name]
user@host# set allow-commands configure
```

If the following statement is included in the configuration and the user does not have the configure login class permission bit, the user can enter configuration exclusive mode:

```
[edit system login class class-name]
user@host# set allow-commands "configure exclusive"
```

NOTE: You cannot use runtime variables. In the following example, the runtime variable 1.2.3.4 cannot be used:

```
[edit system login class class-name]
user@host# set deny "show bgp neighbor 1.2.3.4"
```

Example 2: Configuring Access Privileges to Individual Operational Mode Commands

Configure permissions for individual operational mode commands:

```
[edit]
system {
    login {
/*
* This login class has operator privileges and the additional ability to reboot
the router.
*/
      class operator-and-boot {
        permissions [ clear network reset trace view ];
        allow-commands "request system reboot";
      }
/*
* This login class has operator privileges but can't use any commands
beginning with “set”.
*/
      class operator-no-set {
        permissions [ clear network reset trace view ];
        deny-commands "^set";
      }
/*
* This login class has operator privileges and can install software but not view
bgp information.
*/
      class operator-and-install-but-no-bgp {
        permissions [ clear network reset trace view ];
        allow-commands "request system software add";
        deny-commands "show bgp";
      }
    }
}
```

Specifying Configuration Mode Commands

You can specify extended regular expressions with the allow-configuration and deny-configuration attributes to define user access privileges to parts of the configuration hierarchy or individual configuration mode commands. Doing so overrides login class permission bits set for a user. You can also use wildcards to restrict access. When you define access privileges to parts of the configuration hierarchy or individual configuration mode commands, do the following:

- Specify the full paths in the extended regular expressions with the allow-configuration and deny-configuration attributes.
- Enclose parentheses around an extended regular expression that connects two or more terms with the pipe (|) symbol. For example:

```
[edit system login class class-name]
user@host# set deny-configuration "(system login class) | (system services)"
```

NOTE: Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

You cannot define access to keywords such as set, edit, or activate.

For more information about how to use wildcards, see Table 20 on page 410.

To explicitly allow an individual configuration mode command that would otherwise be denied, include the allow-configuration statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
allow-configuration "regular-expression";
```

To explicitly deny an individual configuration mode command that would otherwise be allowed, include the deny-configuration statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
deny-configuration "regular-expression";
```

If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Regular expressions are not case-sensitive.

You can include one deny-configuration and one allow-configuration statement in each login class.

Use extended regular expressions to specify which configuration mode commands are denied or allowed. You specify these regular expressions in the allow-configuration and deny-configuration statements at the [edit system login class] hierarchy level, or by specifying JUNOS-specific attributes in your TACACS+ or RADIUS authentication server's configuration. You must specify that these regular expressions are sent as the value of Juniper vendor-specific attributes. If regular expressions are received during TACACS+ or RADIUS authentication, they override any regular expressions configured on the local router. For information about TACACS+ or RADIUS authentication, see "Configuring User Access" on page 401.

Command regular expressions implement the extended (modern) regular expressions, as defined in POSIX 1003.2. Table 20 lists common regular expression operators.

Table 20: Configuration Mode Commands—Common Regular Expression Operators

Operator	Match...
|	One of the two terms on either side of the pipe.
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces$" means that the user can issue the show interfaces command but cannot issue show interfaces detail or show interfaces extensive.
[ ]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
( )	A group of commands, indicating an expression to be evaluated; the result is then evaluated as part of the overall expression.
*	0 or more terms.
+	One or more terms.
.	Any character except for a space " ".

Example 3: Defining Access Privileges to Individual Configuration Mode Commands

The following examples show how to configure access privileges to individual configuration mode commands.

If the following statement is included in the configuration and the user's login class permission bit is set to all, the user cannot configure telnet parameters:

```
[edit system login class class-name]
user@host# set deny-configuration "system services telnet"
```

If the following statement is included in the configuration and the user's login class permission bit is set to all, the user cannot issue login class commands within any login class whose name begins with "m":

```
[edit system login class class-name]
user@host# set deny-configuration "system login class m.*"
```

If the following statement is included in the configuration and the user's login class permission bit is set to all, the user cannot issue configuration mode commands at the login class or system services hierarchy levels:

```
[edit system login class class-name]
user@host# set deny-configuration "(system login class) | (system services)"
```

If the following statement is included in the configuration and the user's login class permission bit is set to protocols, the user cannot issue login class commands within any login class whose name begins with "m":

```
[edit system login class class-name]
user@host# set deny-configuration "system login class m.*"
```

Example 4: Configuring Access Privileges to Individual Configuration Mode Commands

Configure permissions for individual configuration mode commands:

```
[edit]
system {
    login {
/*
* This login class has operator privileges and the additional ability to issue
commands at the system services hierarchy.
*/
      class only-system-services {
        permissions [ configure ];
        allow-configuration "system services";
      }
/*
* This login class has operator privileges but can't issue any system services
commands.
*/
      class all-except-system-services {
        permissions [ all ];
        deny-configuration "system services";
      }
/*
}
```

Configuring the Timeout Value for Idle Login Sessions

An idle login session is one in which the CLI operational mode prompt is displayed but there is no input from the keyboard. By default, a login session remains established until a user logs out of the router, even if that session is idle. To close idle sessions automatically, you configure a time limit for each login class. If a session established by a user in that class remains idle for the configured time limit, the session automatically closes.

To define the timeout value for idle login sessions, include the idle-timeout statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
idle-timeout minutes;
```

Specify the number of minutes that a session can be idle before it is automatically closed.

If you have configured a timeout value, the CLI displays messages similar to the following when timing out an idle user. It starts displaying these messages 5 minutes before timing out the user.

```
user@host# Session will be closed in 5 minutes if there is no activity.
Warning: session will be closed in 1 minute if there is no activity
Warning: session will be closed in 10 seconds if there is no activity
Idle timeout exceeded: closing session
```

If you configure a timeout value, the session closes after the specified time has elapsed except if the user is running telnet or monitoring interfaces using the monitor interface or monitor traffic command.

Configuring Tips

By default, the tip command is not enabled when a user logs in. To enable tips, include the login-tip statement at the [edit system login class *class-name*] hierarchy level:

```
[edit system login class class-name]
login-tip;
```

Adding this statement enables the tip command for the class specified, provided the user logs in using the CLI. For information about the tip command, see “Displaying Tips About CLI Commands” on page 185.

Configuring User Accounts

User accounts provide one way for users to access the router. (Users can access the router without accounts if you configured RADIUS or TACACS+ servers, as described in “Configuring User Authentication” on page 364.) For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

To create user accounts, include the user statement at the [edit system login] hierarchy level:

```
[edit system login]
user username {
    full-name complete-name;
    uid uid-value;
    class class-name;
    authentication {
      (encrypted-password "password" | plain-text-password);
      ssh-rsa "public-key";
      ssh-dsa "public-key";
    }
}
```

For each user account, you can define the following:

- Username—(Optional) 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—(Optional) If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.

- User identifier (UID)—(Optional) Numeric identifier that is associated with the user account name. The identifier must be in the range from 100 through 64,000 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.

  You must ensure that the UID is unique. However, it is possible to assign the same UID to different users. If you do this, the CLI displays a warning when you commit the configuration, then assigns the duplicate UID.

- User's access privilege—(Required) One of the login classes you defined in the class statement at the [edit system login] hierarchy level, or one of the default classes listed in Table 18 on page 404.
- Authentication method or methods and passwords that the user can use to access the router—(Optional) You can use SSH or a Message Digest 5 (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. For each method, you can specify the user's password. If you configure the plain-text-password option, you are prompted to enter and confirm the password:

  ```
  [edit system]
  user@host# set root-authentication plain-text-password
  New password: type password here
  Retype new password: retype password here
  ```

  For information about how to create a plain-text passwords, see "Specifying Plain-Text Passwords" on page 17.

  For SSH authentication, you can copy the contents of an SSH keys file into the configuration. For information about how to specify filenames, see "Specifying Filenames and URLs" on page 360.

  To load an SSH key file, use the load-key-file command. This command loads RSA (SSH version 1) and DSA (SSH version 2) public keys. You can also configure a user to use ssh-rsa and ssh-dsa keys.

  If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the load-key-file statement. To view the SSH keys entries, use the configuration mode show command. For example:

  ```
  [edit system]
  user@host# set root-authentication load-key-file my-host:.ssh/identity.pub
  .file.19692              |         0 KB |   0.3 kB/s | ETA: 00:00:00 | 100%
  [edit system]
  user@host# show
  root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
  20740496252839038203869014158453496417001961060835872296
  15634757849182736033612764418742659468932077391083448101
  26831259577226254616679992783161235004386609158662838224
  89746732605661192181489539813965561563786211940327687806
  53816960202749164163735913269396344008443
  boojum@juniper.net"; # SECRET-DATA
  }
  ```

An account for the user root is always present in the configuration. You configure the password for root using the root-authentication statement, as described in "Configuring the Root Password" on page 379.

JUNOS-FIPS has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If JUNOS-FIPS is installed on the router, you cannot configure passwords unless they meet this standard.

Example: Configuring User Accounts

Create accounts for four router users, and create an account for the template user “remote.” All users use one of the default system login classes.

```
[edit]
system {
    login {
      user philip {
        full-name “Philip of Macedonia”;
        uid 1001;
        class super-user;
        authentication {
          encrypted-password “$1$poPPeY”;
        }
      }
      user alexander {
        full-name “Alexander the Great”;
        uid 1002;
        class view;
        authentication {
          encrypted-password “$1$14c5.$sBopasdFFdssdfFFdsdfs0”;
          ssh-dsa “8924 37 5678 5678@gaugamela.per”;
        }
      }
      user darius {
        full-name “Darius King of Persia”;
        uid 1003;
        class operator;
        authentication {
          ssh-rsa “1024 37 12341234@ecbatana.per”;
        }
      }
      user anonymous {
        class unauthorized;
      }
      user remote {
        full-name “All remote users”;
        uid 9999;
        class read-only;
      }
    }
}
```

JUNOS-FIPS Crypto Officer and User Accounts

JUNOS-FIPS defines a restricted set of user roles. Unlike JUNOS, which allows a wide range of capabilities to users, FIPS 140-2 defines specific types of users (Crypto Officer, User, and Maintenance). Crypto Officers and FIPS Users perform all FIPS-related configuration tasks and issue all FIPS-related commands. Crypto Officer and FIPS User configurations must follow FIPS 140-2 guidelines. Typically, no user besides a Crypto Officer can perform FIPS-related tasks. For more information, see the *JUNOS-FIPS Configuration Guide*

Crypto Officer User Configuration

JUNOS-FIPS offers finer control of user permissions than those mandated by FIPS 140-2. For FIPS 140-2 conformance, any JUNOS-FIPS user with the secret, security, and maintenance permission bits set is a Crypto Officer. In most cases, the super-user class should be reserved for a Crypto Officer. A FIPS User can be defined as any JUNOS-FIPS user that does not have the secret, security, and maintenance bits set.

FIPS User Configuration

A Crypto Officer sets up FIPS Users. FIPS Users can be granted permissions normally reserved for a Crypto Officer; for example, permission to zeroize the system and individual AS-II FIPS PICs.

Chapter 23

Configuring Time

This chapter discusses the following topics related to configuring time:

- Setting the Time Zone on page 417
- Configuring the Network Time Protocol on page 418

For more information about configuring time, see “Setting the Current Date and Time” on page 200. For more information about setting the date and time for Network Time Protocol (NTP) servers, see “Setting the Date and Time from NTP Servers” on page 200.

Setting the Time Zone

The default local time zone on the router is UTC (Coordinated Universal Time, formerly known as Greenwich Mean Time, or GMT). To modify the local time zone, include the time-zone statement at the [edit system] hierarchy level:

```
[edit system]
time-zone (GMThour-offset | time-zone);
```

You can use the GMT*hour-offset* option to set the time zone relative to UTC (GMT) time. By default, *hour-offset* is 0. You can configure this to be a value in the range from -14 to +12.

You can also specify *time-zone* using the continent and and major city.

For the time zone change to take effect for all processes running on the router, you must reboot the router.

For information about setting the time on the router, see “Setting the Current Date and Time” on page 200.

Examples: Setting the Time Zone

Set the time zone for New York:

```
[edit]
user@host# set system time-zone America/New_York
[edit]
user@host# show
system {
    time-zone America/New_York;
}
```

Set the time zone for Pacific Time:

```
[edit]
user@host# set system time-zone America/Los_Angeles
[edit]
user@host# show
system {
    time-zone America/Los_Angeles;
}
```

Set the time zone relative to UTC (GMT):

```
[edit]
user@host# set system time-zone GMT+2
```

For information about what time zones are available, see time-zone on page 602.

Configuring the Network Time Protocol

NTP provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. NTP uses a returnable-time design in which a distributed subnet of time servers operating in a self-organizing, hierarchical master-slave configuration synchronizes local clocks within the subnet and to national time standards by means of wire or radio. The servers also can redistribute reference time using local routing algorithms and time daemons.

NTP is defined in RFC 1305, *Network Time Protocol (Version 3) Specification, Implementation and Analysis*.

To configure NTP on the router, include the ntp statement at the [edit system] hierarchy level:

```
[edit system]
ntp {
    authentication-key number type type value password;
    boot-server (NTP) address;
    broadcast <address> <key key-number> <version value> <ttl value>;
    broadcast-client;
    multicast-client <address>;
    peer address <key key-number> <version value> <prefer>;
    server address <key key-number> <version value> <prefer>;
    source-address source-address;
    trusted-key [ key-numbers ];
}
```

To configure NTP properties, you can do one or more of the following:

- Configuring the NTP Boot Server on page 419
- Specifying a Source Address for an NTP Server on page 420
- Configuring the NTP Time Server and Time Services on page 420
- Configuring NTP Authentication Keys on page 424
- Configuring the Router to Listen for Broadcast Messages on page 424
- Configuring the Router to Listen for Multicast Messages on page 425

When configuring NTP, you do not actively configure time servers. Rather, all clients also are servers. An NTP server is not believed unless it, in turn, is synchronized to another NTP server—which itself must be synchronized to something upstream, eventually terminating in a high-precision clock.

If the time difference between the local router clock and the NTP server clock is more than 128 milliseconds, but less than 128 seconds, the clocks are slowly stepped into synchronization. However, if the difference is more than 128 seconds, the clocks are not synchronized. You must set the time on the local router so that the difference is less than 128 seconds to start the synchronization process. On the local router, you set the date and time using the set date command. To set the time automatically, use the boot-server statement at the [edit system ntp] hierarchy level, specifying the address of an NTP server.

Configuring the NTP Boot Server

When you boot the router, it issues an ntpdate request, which polls a network server to determine the local date and time. You need to configure a server that the router uses to determine the time when the router boots. Otherwise, NTP will not be able to synchronize to a time server if the server's time appears to be very far off of the local router's time.

To configure the NTP boot server, include the boot-server statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
boot-server (NTP) address;
```

Specify the address of the network server. You must specify an address, not a hostname.

Specifying a Source Address for an NTP Server

For IP version 4 (IPv4), you can specify that if the NTP server configured at the [edit system ntp] hierarchy level is contacted on one of the loopback interface addresses, then the reply will always use a specific source address. This is useful for controlling which source address NTP will use to access your network when it is either responding to an NTP client request from your network or when it itself is sending NTP requests to your network.

To configure the specific source address that the reply will always use, and the source address that requests initiated by NTP server will use, include the source-address statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
source-address source-address;
```

source-address is a valid IP address configured on one of the router interfaces.

Configuring the NTP Time Server and Time Services

When configuring NTP, you can specify which system on the network is the authoritative time source, or time server, and how time is synchronized between systems on the network. To do this, you configure the router to operate in one of the following modes:

- Client mode—In this mode, the local router can be synchronized to the remote system, but the remote system can never be synchronized to the local router.
- Symmetric active mode—In this mode, the local router and the remote system can synchronize each other. You use this mode in a network in which either the local router or the remote system might be a better source of time.

NOTE: Symmetric active mode can be initiated by either the local or remote system. Only one system needs to be configured to do so. This means that the local system can synchronize to any system that offers symmetric active mode without any configuration whatsoever. However, we strongly encourage you to configure authentication to ensure that the local system synchronizes only to known time servers.

- Broadcast mode—In this mode, the local router sends periodic broadcast messages to a client population at the specified broadcast or multicast *address*. Normally, you include this statement only when the local router is operating as a transmitter.
- Server mode—In this mode, the local router operates as an NTP server.

NOTE: In NTP server mode, the JUNOS software does not support authentication.

The following sections describe how to configure these modes of operation:

- Configuring the Router to Operate in Client Mode on page 421
- Configuring the Router to Operate in Symmetric Active Mode on page 422

- Configuring the Router to Operate in Broadcast Mode on page 422
- Configuring the Router to Operate in Server Mode on page 423

Configuring the Router to Operate in Client Mode

To configure the local router to operate in client mode, include the server statement and other optional statements at the [edit system ntp] hierarchy level:

```
[edit system ntp]
server address <key key-number> <version value> <prefer>;
authentication-key key-number type type value password;
boot-server address;
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the key option. The key corresponds to the key number you specify in the authentication-key statement, as described in “Configuring NTP Authentication Keys” on page 424.

By default, the router sends NTP version 3 packets to the time server. To set the NTP version level to 1 or 2, include the version option.

If you configure more than one time server, you can mark one server preferred by including the prefer option.

For information about how to configure trusted keys, see “Configuring NTP Authentication Keys” on page 424. For information about how to configure an NTP boot server, see “Configuring the NTP Boot Server” on page 419. For information about how to configure the router to operate in server mode, see “Configuring the Router to Operate in Server Mode” on page 423.

Example: Configuring Client Mode

Configure the router to operate in client mode:

```
[edit system ntp]
authentication-key 1 type md5 value "$9$EgfcrvX7VY4ZEcwgoHjkP5Q3CuREyv87";
boot-server 10.1.1.1;
server 10.1.1.1 key 1 prefer;
trusted-key 1;
```

Configuring the Router to Operate in Symmetric Active Mode

To configure the local router to operate in symmetric active mode, include the peer statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
peer address <key key-number> <version value> <prefer>;
```

Specify the address of the remote system. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the remote system, include the key option. The key corresponds to the key number you specify in the authentication-key statement, as described in "Configuring NTP Authentication Keys" on page 424.

By default, the router sends NTP version 3 packets to the remote system. To set the NTP version level to 1 or 2, include the version option.

If you configure more than one remote system, you can mark one system preferred by including the prefer option:

```
peer address <key key-number> <version value> prefer;
```

Configuring the Router to Operate in Broadcast Mode

To configure the local router to operate in broadcast mode, include the broadcast statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
broadcast address <key key-number> <version value> <ttl value>;
```

Specify the broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. Currently, the multicast address must be 224.0.1.1.

To include an authentication key in all messages sent to the remote system, include the key option. The key corresponds to the key number you specify in the authentication-key statement, as described in "Configuring NTP Authentication Keys" on page 424.

By default, the router sends NTP version 3 packets to the remote system. To set the NTP version level to 1 or 2, include the version option.

Configuring the Router to Operate in Server Mode

In server mode, the router acts as an NTP server for clients when the clients are configured appropriately. The only prerequisite for “server mode” is that the router must be receiving time from another NTP peer or server. No other configuration is necessary on the router.

To configure the local router to operate as an NTP server, include the following statements at the [edit system ntp] hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
server address <key key-number> <version value> <prefer>;
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the key option. The key corresponds to the key number you specify in the authentication-key statement, as described in “Configuring NTP Authentication Keys” on page 424.

By default, the router sends NTP version 3 packets to the time server. To set the NTP version level to 1 or 2, include the version option.

If you configure more than one time server, you can mark one server preferred by including the prefer option.

For information about how to configure trusted keys, see “Configuring NTP Authentication Keys” on page 424. For information about how to configure the router to operate in client mode, see “Configuring the Router to Operate in Client Mode” on page 421.

Example: Configuring Server Mode

Configure the router to operate in server mode:

```
[edit system ntp]
authentication-key 1 type md5 value "$9$txERuBEreWx-wtuLNdboaUjH.T3AtOESe";
server 172.17.17.27.46 prefer;
trusted-key 1;
```

Configuring NTP Authentication Keys

Time synchronization can be authenticated to ensure that the local router obtains its time services only from known sources. By default, network time synchronization is unauthenticated. The system will synchronize to whatever system appears to have the most accurate time. We strongly encourage you to configure authentication of network time services.

To authenticate other time servers, include the trusted-key statement at the [edit system ntp] hierarchy level. Only time servers transmitting network time packets that contain one of the specified key numbers and whose key matches the value configured for that key number are eligible to be synchronized to. Other systems can synchronize to the local router without being authenticated.

```
[edit system ntp]
trusted-key [ key-numbers ];
```

Each key can be any 32-bit unsigned integer except 0. Include the key option in the peer, server, or broadcast statements to transmit the specified authentication key when transmitting packets. The key is necessary if the remote system has authentication enabled so that it can synchronize to the local system.

To define the authentication keys, include the authentication-key statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
```

number is the key number, *type* is the authentication type (either Message Digest 5 [MD5] or DES), and *password* is the password for this key. The key number, type, and password must match on all systems using that particular key for authentication.

Configuring the Router to Listen for Broadcast Messages

When you are using NTP, you can configure the local router to listen for broadcast messages on the local network to discover other servers on the same subnet by including the broadcast-client statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
broadcast-client;
```

When the router hears a broadcast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *broadcast client* mode, in which it listens for, and synchronizes to, succeeding broadcast messages.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

Configuring the Router to Listen for Multicast Messages

When you are using NTP, you can configure the local router to listen for multicast messages on the local network to discover other servers on the same subnet by including the multicast-client statement at the [edit system ntp] hierarchy level:

```
[edit system ntp]
multicast-client <address>;
```

When the router hears a multicast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *multicast client* mode, in which it listens for, and synchronizes to, succeeding multicast messages.

You can specify one or more IP addresses. (You must specify an address, not a hostname.) If you do, the route joins those multicast groups. If you do not specify any addresses, the software uses 224.0.1.1.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

Chapter 24

Configuring System Log Messages

The JUNOS software generates system log messages (also called syslog messages) to record events that occur on the routing platform, 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 routing platform power-down due to excessive temperature

Each system log message identifies the JUNOS software process that generated the message and briefly describes the operation or error that occurred. This manual provides more detailed information about each system log message and, when applicable, describes possible causes of the message and action you can take to correct error conditions.

NOTE: The information in this chapter applies to JUNOS software processes and libraries, not to the services on a Physical Interface Card (PIC) such as the Adaptive Services PIC. For information about configuring system logging for PIC services, see the *JUNOS Services Interfaces Configuration Guide*.

This chapter discusses the following topics:

System Logging Configuration Statements

To configure the routing platform to log system messages, include the syslog statement at the [edit system] hierarchy level:

```
[edit system]
syslog {
    archive {
        files number;
        size size;
        (world-readable | no-world-readable);
    }
    console {
        facility severity;
    }
    file filename {
        facility severity;
        explicit-priority;
        match "regular-expression";
        archive {
            files number;
            size size;
            (world-readable | no-world-readable);
        }
    }
    host (hostname | other-routing-engine | scc-master) {
        facility severity;
        explicit-priority;
        facility-override facility;
        log-prefix string;
        match "regular-expression";
    }
    source-address source-address;
    time-format (year | millisecond | year millisecond);
    user (username | *) {
        facility severity;
        match "regular-expression";
    }
}
```

Minimum System Logging Configuration

For the JUNOS software processes to generate system log messages, you must include the syslog statement at the [edit system] hierarchy level. Specify at least one destination for system log messages, as described in Table 21. For more information about the configuration statements, see “Configuring System Logging for a Single-Chassis System” on page 429.

Table 21: Minimum Configuration Statements for System Logging

Destination	Minimum Configuration Statements

File	[edit system syslog]
	file *filename* {
	facility severity;
	}

Terminal session of one, several, or all users	[edit system syslog]
	user (*username* | *) {
	facility severity;
	}

Routing platform console	[edit system syslog]
	console {
	facility severity;
	}

Remote machine or the other Routing Engine on the routing platform	[edit system syslog]
	host (*hostname* | other-routing-engine) {
	facility severity;
	}

Configuring System Logging for a Single-Chassis System

The JUNOS system logging utility is similar to the UNIX syslogd utility. This section describes how to configure system logging for a single-chassis system that runs the JUNOS software.

System logging for the JUNOS-FIPS software is the same as for the JUNOS software. For recommendations about which facilities to log, see the *JUNOS FIPS Configuration Guide*.

For information about configuring system logging for a routing matrix, see “Configuring System Logging for a Routing Matrix” on page 446.

When you configure system logging, you can direct messages to one or more destinations by including the appropriate statement at the [edit system syslog] hierarchy level:

- To a named file in a local file system, by including the file statement. See “Directing Messages to a Log File” on page 432.
- To the terminal session of one or more specific users (or all users) when they are logged in to the routing platform, by including the user statement. See “Directing Messages to a User Terminal” on page 433.

- To the routing platform console, by including the console statement. See "Directing Messages to the Console" on page 433.
- To a remote machine that is running the syslogd utility or to the other Routing Engine on the routing platform, by including the host statement. See "Directing Messages to a Remote Machine or the Other Routing Engine" on page 433.

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). To log the messages belonging to one or more facilities to a particular destination, specify each facility name as a separate statement within the set of statements for the destination. Table 22 lists the JUNOS system logging facilities that you can specify in the configuration statements at the [edit system syslog] hierarchy level.

Table 22: JUNOS System Logging Facilities

Facility	Type of Event or Error
any	All (messages from all facilities)
authorization	Authentication and authorization attempts
change-log	Changes to the JUNOS configuration
conflict-log	Configuration that is inconsistent with routing platform hardware
daemon	Actions performed or errors encountered by various system processes
firewall	Packet filtering actions performed by a firewall filter
ftp	Actions performed or errors encountered by the FTP process
interactive-commands	Commands issued at the JUNOS command-line interface (CLI) prompt or by a JUNOScript client application
kernel	Actions performed or errors encountered by the JUNOS kernel
pfe	Actions performed or errors encountered by the Packet Forwarding Engine
user	Actions performed or errors encountered by various user-space processes

Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects routing platform functions. When you configure logging for a facility and destination, you specify a severity level for each facility; messages from the facility that are rated at that level or higher are logged to the destination.

Unlike the other severity levels, the none level disables logging of a facility instead of indicating how seriously a triggering event affects routing functions. For more information, see "Disabling Logging of a Facility" on page 444.

Table 23 lists the severity levels that you can specify in configuration statements at the [edit system syslog] hierarchy level. The levels from emergency through info are in order from highest severity (greatest effect on functioning) to lowest.

Table 23: System Log Message Severity Levels

Severity Level	Description
any	Includes all severity levels
none	Disables logging of the associated facility to a destination
emergency	System panic or other condition that causes the routing platform to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard drive errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or nonerror conditions of interest

A message's facility and severity level are together referred to as its *priority*. By default, priority information is not included in system log messages. To include priority information in messages directed to a file or a remote destination, include the explicit-priority statement. For more information, see "Including Priority in System Log Messages" on page 439.

You can modify the timestamp on system log messages to include the year, the millisecond, or both. For more information, see "Including the Year or Millisecond in Timestamps" on page 442.

When directing messages to a remote machine, you can specify the source address to use, and you can configure features that make it easier to separate JUNOS-specific messages or messages generated on particular routing platforms. For more information, see "Directing Messages to a Remote Machine or the Other Routing Engine" on page 433.

The predefined facilities group together related messages, but you can also use regular expression matching to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination. For more information, see "Using Regular Expressions to Refine the Set of Logged Messages" on page 442.

For more information about configuring system logging, see the following sections:

- Directing Messages to a Log File on page 432
- Directing Messages to a User Terminal on page 433
- Directing Messages to the Console on page 433
- Directing Messages to a Remote Machine or the Other Routing Engine on page 433

- Configuring Log File Archiving on page 438
- Including Priority in System Log Messages on page 439
- Including the Year or Millisecond in Timestamps on page 442
- Using Regular Expressions to Refine the Set of Logged Messages on page 442
- Disabling Logging of a Facility on page 444
- Examples: Configuring System Logging on page 444

Directing Messages to a Log File

To direct system log messages to a file on the local disk of the local Routing Engine, include the file statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
file filename {
    facility severity;
    explicit-priority;
    match "regular-expression";
    archive {
      files number;
      size size;
      (world-readable | no-world-readable);
    }
}
```

The default directory for log files is /var/log; to specify a different directory on the local Routing Engine's local disk, include the complete pathname. For the list of logging facilities and severity levels, see Table 22 on page 430 and Table 23 on page 431.

To prevent log files from growing too large, the JUNOS system logging utility by default writes messages to a sequence of files of a defined size. You can configure the number of files, their maximum size, and who can read them, for either all log files or a certain log file. For more information, see "Configuring Log File Archiving" on page 438.

For information about the explicit-priority statement, see "Including Priority in System Log Messages" on page 439.

For information about the match statement, see "Using Regular Expressions to Refine the Set of Logged Messages" on page 442.

Directing 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 in to the local Routing Engine, include the user statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
user (username | *) {
    facility severity;
    match "regular-expression";
}
```

Specify one or more JUNOS usernames, separating multiple values with spaces, or use the asterisk (*) to indicate all users who are logged in to the local Routing Engine.

For the list of logging facilities and severity levels, see Table 22 on page 430 and Table 23 on page 431. For information about the match statement, see “Using Regular Expressions to Refine the Set of Logged Messages” on page 442.

Directing Messages to the Console

To direct system log messages to the console of the local Routing Engine, include the console statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
console {
    facility severity;
}
```

For the list of logging facilities and severity levels, see Table 22 on page 430 and Table 23 on page 431.

Directing Messages to a Remote Machine or the Other Routing Engine

To direct system log messages to a remote machine or to the other Routing Engine on the routing platform, include the host statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
host (hostname | other-routing-engine) {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
```

To direct system log messages to a remote machine, include the host *hostname* statement to specify the remote machine’s IP address or fully qualified hostname. The remote machine must be running the standard syslogd utility. We do not recommend directing messages to another routing platform. In each system log message directed to the remote machine, the hostname of the local Routing Engine appears after the timestamp to indicate that it is the source for the message.

To direct system log messages to the other Routing Engine on a routing platform with two Routing Engines installed and operational, include the host other-routing-engine statement. Include the statement in each Routing Engine's configuration if you want them both to direct messages to the other Routing Engine. In each system log message directed to the other Routing Engine, the string re0 or re1 appears after the timestamp to indicate that the local Routing Engine is the source for the message.

For the list of logging facilities and severity levels to configure under the host statement, see Table 22 on page 430 and Table 23 on page 431.

To record facility and severity level information in each message, include the explicit-priority statement. For more information, see "Including Priority in System Log Messages" on page 439.

For information about the match statement, see "Using Regular Expressions to Refine the Set of Logged Messages" on page 442.

When directing messages to remote machines, you can use the source-address statement to specify the source address to use. In each host statement, you can also include the facility-override statement to assign an alternate facility and the log-prefix statement to add a string to each message. For more information, see the following sections:

- Specifying an Alternate Source Address on page 434
- Changing the Alternate Facility Name for Remote Messages on page 435
- Adding a String to System Log Messages on page 437

Specifying an Alternate Source Address

To specify the source address to use when directing system log messages to a remote machine, include the source-address statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
source-address source-address;
```

source-address is a valid IP address configured on one of the routing platform interfaces. The address is used only for messages directed to the remote machines specified in all host *hostname* statements at the [edit system syslog] hierarchy level, not for messages directed to the other Routing Engine.

Changing the Alternate Facility Name for Remote Messages

Some facilities assigned to messages logged on the local routing platform have JUNOS-specific names (see Table 22 on page 430). If a remote machine designated at the [edit system syslog host *hostname*] hierarchy level is not a Juniper Networks routing platform, its syslogd utility cannot interpret the JUNOS-specific names. To enable the standard syslogd utility to handle messages from these facilities, messages directed to a remote machine use a standard local*X* facility name is used instead of the JUNOS-specific facility name.

Table 24 lists the default alternate facility name used for each JUNOS-specific facility name. For facilities that are not listed, the default alternate name is the same as the local facility name.

Table 24: Default Facilities for Messages Directed to a Remote Destination

JUNOS-Specific Local Facility	Default Facility when Directed to Remote Destination
change-log	local6
conflict-log	local5
firewall	local3
interactive-commands	local7
pfe	local4

The syslogd utility on a remote machine handles all messages that belong to a facility in the same way, regardless of the source of the message (the Juniper Networks routing platform or the remote machine itself). As an example, the following statements in the configuration of the routing platform local-router direct messages from the authorization facility to the remote machine monitor.mycompany.com:

```
[edit system syslog]
host monitor.mycompany.com {
    authorization info;
}
```

The default alternate facility for the local authorization facility is also authorization. The syslogd utility on monitor is configured to write messages belonging to the authorization facility to the file /var/log/auth-attempts, which means that the file contains both the messages generated when users log in to local-router and the messages generated when users log in to monitor. Although the name of the source machine appears in each system log message, the mixing of messages from multiple machines can make it more difficult to analyze the contents of the auth-attempts file.

To change the facility used for all messages directed to a remote machine, include the facility-override statement at the [edit system syslog host *hostname*] hierarchy level:

```
[edit system syslog host hostname]
facility severity;
facility-override facility;
```

In general, it makes sense to specify an alternate facility that is not already in use on the remote machine, such as one of the local*X* facilities. On the remote machine, you must also configure the syslogd utility to handle the messages in the desired manner.

Table 25 lists the facilities that you can specify in the facility-override statement.

Table 25: Facilities for the facility-override Statement

Facility	Description
authorization	Authentication and authorization attempts
daemon	Actions performed or errors encountered by various system processes
ftp	Actions performed or errors encountered by the FTP process
kernel	Actions performed or errors encountered by the JUNOS kernel
local0	Local facility number 0
local1	Local facility number 1
local2	Local facility number 2
local3	Local facility number 3
local4	Local facility number 4
local5	Local facility number 5
local6	Local facility number 6
local7	Local facility number 7
user	Actions performed or errors encountered by various user-space processes

We do not recommend including the facility-override statement at the [edit system syslog host other-routing-engine] hierarchy level. It is not necessary to use alternate facility names when directing messages to the other Routing Engine, because its JUNOS system logging utility can interpret the JUNOS-specific names.

Examples: Assigning an Alternate Facility

Log all messages generated on the local routing platform at the error level or higher to the local0 facility on the remote machine called monitor.mycompany.com:

```
[edit system syslog]
host monitor.mycompany.com {
    any error;
    facility-override local0;
}
```

Configure routing platforms located in California and routing platforms located in New York to send messages to a single remote machine called central-logger.mycompany.com. The messages from California are aggregated into one facility (local1) and the messages from New York into another facility (local2).

- Configure California routing platforms to aggregate messages in the local1 facility:

  ```
  [edit system syslog]
  host central-logger.mycompany.com {
      change-log info;
      facility-override local1;
  }
  ```

- Configure New York routing platforms to aggregate messages in the local2 facility:

  ```
  [edit system syslog]
  host central-logger.mycompany.com {
      change-log info;
      facility-override local2;
  }
  ```

On central-logger, you can then configure the system logging utility to write messages from the local1 facility to /var/log/california-config and the messages from the local2 facility to /var/log/new-york-config.

Adding a String to System Log Messages

To add a text string to every system log message directed to a remote machine or to the other Routing Engine, include the log-prefix statement at the [edit system syslog host] hierarchy level:

```
[edit system syslog host (hostname | other-routing-engine)]
facility severity;
log-prefix string;
```

The string can contain any alphanumeric character except the space, the equal sign (=), and the colon (:). A colon and a space are appended to the string when the system log messages are written to the log. The string is inserted after the identifier for the Routing Engine that generated the message.

Example: Adding a String

Add the string M40e to all messages to indicate that the router is an M40e router, and direct the messages to the remote machine hardware-logger.mycompany.com:

```
[edit system syslog]
host hardware-logger.mycompany.com {
    any info;
    log-prefix M40e;
}
```

When these configuration statements are included on an M40e router called origin1, a message in the system logging file on hardware-logger looks like the following:

```
Mar 9 17:33:23 origin1 M40e: mgd[477]: UI_CMDLINE_READ_LINE: user 'root',
command 'run show version'
```

Configuring Log File Archiving

For each log file *logfile* that you configure, the JUNOS logging utility by default writes 128 kilobytes (KB) of messages to the file before closing it, compressing it, and naming the compressed file *logfile*.0.gz. The logging utility then opens and writes to a new file called *logfile*. When the new *logfile* reaches 128 KB in size, *logfile*.0.gz is renamed *logfile*.1.gz and the new *logfile* is closed, compressed, and renamed *logfile*.0.gz. By default, the logging utility creates up to 10 archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next 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.

You can include the archive statement in the configuration to change the maximum size of each file, how many archive files are created, and who can read log files. To configure values that apply to all log files, include the archive statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
archive {
    files number;
    size size;
    (world-readable | no-world-readable);
}
```

To configure values that apply to a particular log file, include the archive statement at the [edit system syslog file *filename*] hierarchy level:

```
[edit system syslog file filename]
facility severity;
archive {
    files number;
    size size;
    (world-readable | no-world-readable);
}
```

The number of files specified with the files statement can range from 1 through 1000. The maximum file size specified with the size statement can range from 64 KB (64k) through 1 gigabyte (1g); to represent megabytes, use the letter m after the integer. To enable all users to read log files, include the world-readable statement. To restore the default permissions, include the no-world-readable statement.

Including Priority in System Log Messages

A message's facility and severity level are together referred to as its *priority*. By default, the system logging utility does not include information about priority in system log messages.

To include the priority in messages directed to a file, include the explicit-priority statement at the [edit system syslog file *filename*] hierarchy level:

```
[edit system syslog file filename]
facility severity;
explicit-priority;
```

To include the priority in messages directed to a remote machine or the other Routing Engine, include the explicit-priority statement at the [edit system syslog host (*hostname* | other-routing-engine)] hierarchy level:

```
[edit system syslog host (hostname | other-routing-engine)]
facility severity;
explicit-priority;
```

When the explicit-priority statement is included, the message string includes the facility name and a numerical code representing the severity level. The name and number appear directly before the message code.

The message string always reports the original, local facility name. If the message belongs to a facility with a JUNOS-specific name, the JUNOS system logging utility still uses the alternate facility name for the message itself when it directs the message to a remote destination. For more information, see "Changing the Alternate Facility Name for Remote Messages" on page 435.

Table 26 on page 440 lists the facility codes that can appear in system log messages and maps them to facility names.

NOTE: If Table 26 does not provide the facility name for a code, you cannot include the facility in a statement at the [edit system syslog] hierarchy level. The JUNOS software might use these facilities—and others that are not listed—when reporting on internal operations.

Table 26: Mapping of Facility Codes to Names

Code	JUNOS Facility Name	Type of Event or Error
AUTH	authorization	Authentication and authorization attempts
AUTHPRIV		Authentication and authorization attempts that can be viewed by superusers only
CHANGE	change-log	Changes to the JUNOS configuration
CONFLICT	conflict-log	Configuration that is inconsistent with routing platform hardware
CONSOLE		Messages written to /dev/console by the kernel console output driver
CRON		Actions performed or errors encountered by the cron process
DAEMON	daemon	Actions performed or errors encountered by various system processes
FIREWALL	firewall	Packet filtering actions performed by a firewall filter
FTP	ftp	Actions performed or errors encountered by the FTP process
INTERACT	interactive-commands	Commands issued at the JUNOS CLI prompt or by a JUNOScript client application
KERN	kernel	Actions performed or errors encountered by the JUNOS kernel
NTP		Actions performed or errors encountered by the Network Time Protocol process (ntpd)
PFE	pfe	Actions performed or errors encountered by the Packet Forwarding Engine
SYSLOG		Actions performed or errors encountered by the JUNOS system logging utility
USER	user	Actions performed or errors encountered by various user-space processes

Table 27 lists the numerical severity codes that can appear in system log messages and maps them to severity levels.

Table 27: Mapping of Numerical Codes to Severity Levels

Numerical Code	Severity Level	Description
0	emergency	System panic or other condition that causes the routing platform to stop functioning
1	alert	Conditions that require immediate correction, such as a corrupted system database
2	critical	Critical conditions, such as hard drive errors
3	error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
4	warning	Conditions that warrant monitoring
5	notice	Conditions that are not errors but might warrant special handling
6	info	Events or nonerror conditions of interest
7	debug	Software debugging messages (these appear only if a technical support representative has instructed you to configure this severity level)

In the following example, the CHASSISD_PARSE_COMPLETE message belongs to the daemon facility and is assigned severity info (6):

```
Aug 21 12:36:30 router1 chassisd[522]:
%DAEMON-6-CHASSISD_PARSE_COMPLETE: Using new configuration
```

When the explicit-priority statement is not included, the priority does not appear in the message, which has the following format:

```
Aug 21 12:36:30 router1 chassisd[522]: CHASSISD_PARSE_COMPLETE: Using
new configuration
```

Including the Year or Millisecond in Timestamps

By default, the timestamp recorded in a system log message specifies the month, date, hour, minute, and second when the message was logged, as in the following example:

```
Aug 21 12:36:30
```

To include the year, the millisecond, or both in the timestamp, include the time-format statement at the [edit system syslog] hierarchy level:

```
[edit system syslog]
time-format (year | millisecond | year millisecond);
```

The modified timestamp is used for messages directed to each destination configured by a file, console, or user statement at the [edit system syslog] hierarchy level, but not to destinations configured by a host statement.

The following example illustrates the format for a timestamp that includes both the millisecond (401) and the year (2003):

```
Aug 21 12:36:30.401 2003
```

Using Regular Expressions to Refine the Set of Logged Messages

The predefined facilities group together related messages, but you can also use regular expression matching to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination.

To specify the text string that must (or must not) appear in a message for the message to be logged to a destination, include the match statement and specify the regular expression which the text string must match:

```
match "regular-expression";
```

You can include this statement at the following hierarchy levels:

- [edit system syslog file *filename*] (for a file)
- [edit system syslog user (*username* | *)] (for the terminal session of one or all users)
- [edit system syslog host (*hostname* | other-routing-engine)] (for a remote destination)

When you specify the regular expression, use the notation defined in POSIX Standard 1003.2 for extended (modern) UNIX regular expressions. Explaining regular expression syntax is beyond the scope of this document. Table 28 specifies which character or characters are matched by some of the regular expression operators that you can use in the match statement. In the descriptions, the term *term* refers to either a single alphanumeric character or a set of characters enclosed in square brackets, parentheses, or braces.

NOTE: The match statement is not case-sensitive.

Table 28: Regular Expression Operators for the match Statement

Operator	Matches
. (period)	One instance of any character except the space.
* (asterisk)	Zero or more instances of the immediately preceding term.
+ (plus sign)	One or more instances of the immediately preceding term.
? (question mark)	Zero or one instance of the immediately preceding term.
| (pipe)	One of the terms that appear on either side of the pipe operator.
! (exclamation point)	Any string except the one specified by the expression, when the exclamation point appears at the start of the expression. Use of the exclamation point is JUNOS-specific.
^ (caret)	The start of a line, when the caret appears outside square brackets. One instance of any character that does not follow it within square brackets, when the caret is the first character inside square brackets.
$ (dollar sign)	The end of a line.
[ ] (paired square brackets)	One instance of one of the enclosed alphanumeric characters. To indicate a range of characters, use a hyphen (-) to separate the beginning and ending characters of the range. For example, [a-z0-9] matches any letter or number.
( ) (paired parentheses)	One instance of the evaluated value of the enclosed term. Parentheses are used to indicate the order of evaluation in the regular expression.

Example: Using Regular Expressions

Filter messages that belong to the interactive-commands facility, directing those that include the string configure to the terminal of the root user:

```
[edit system syslog]
user root {
    interactive-commands any;
    match “.*configure.*”;
}
```

Messages like the following appear on the root user’s terminal when a user issues a configure command to enter configuration mode:

```
timestamp router-name mgd[PID]: UI_CMDLINE_READ_LINE: User 'user', command
'configure private'
```

Filter messages that belong to the daemon facility and have severity error or higher, directing them to the file /var/log/process-errors. Omit messages generated by the snmpd process, instead directing them to the file /var/log/snmpd-errors:

```
[edit system syslog]
file process-errors {
    daemon error;
    match “!(.*snmpd.*)”;
}
file snmpd-errors {
    daemon error;
    match “.*snmpd.*”;
}
```

Disabling Logging of a Facility

To disable the logging of messages from a facility, include the *facility* none statement in the configuration. This statement is useful when, for example, you want to log messages that have the same severity level and belong to all but a few facilities. Instead of including a statement for each facility you want to log, you can include the any *severity* statement and then a *facility* none statement for each facility you do not want to log. For example, the following logs all messages at the error level or higher to the console, except for messages from the daemon and kernel facilities. Messages from those facilities are logged to the file /var/log/internals instead:

```
[edit system syslog]
console {
    any error;
    daemon none;
    kernel none;
}
file internals {
    daemon info;
    kernel info;
}
```

Examples: Configuring System Logging

Log messages about all commands entered by users at the CLI prompt or by JUNOScript client applications, and all authentication or authorization attempts, both to the file cli-commands and to the terminal of any user who is logged in:

```
[edit system]
syslog {
    file cli-commands {
      interactive-commands info;
      authorization info;
    }
    user * {
      interactive-commands info;
      authorization info;
    }
}
```

Log all changes in the state of alarms to the file /var/log/alarms:

```
[edit system]
syslog {
    file alarms {
      kernel warning;
    }
}
```

Configure the handling of messages of various types, as described in the comments. Information is logged to two files, to the terminal of user alex, to a remote machine, and to the console:

```
[edit system]
syslog {
/* write all security-related messages to file /var/log/security */
      file security {
         authorization info;
         interactive-commands info;
      }
/* write messages about potential problems to file /var/log/messages: */
/* messages from “authorization” facility at level “notice” and above, */
/* messages from all other facilities at level “warning” and above */
      file messages {
         authorization notice;
         any warning;
      }
/* write all messages at level “critical” and above to terminal of user “alex” if */
/* that user is logged in */
      user alex {
         any critical;
      }
/* write all messages from the “daemon” facility at level “info” and above, and */
/* messages from all other facilities at level “warning” and above, to the */
/* machine monitor.mycompany.com */
      host monitor.mycompany.com {
         daemon info;
         any warning;
      }
/* write all messages at level “error” or above to the system console */
      console {
         any error;
      }
}
```

Configure the handling of messages generated when users issue JUNOS CLI commands, by specifying the interactive-commands facility at the following severity levels:

- info—Logs a message when users issue any command at the CLI operational or configuration mode prompt. The example writes the messages to the file /var/log/user-actions.

- notice—Logs a message when users issue the configuration mode commands rollback and commit. The example writes the messages to the terminal of user philip.

- warning—Logs a message when users issue a command that restarts a software process. The example writes the messages to the console.

```
[edit system]
syslog {
    file user-actions {
        interactive-commands info;
    }
    user philip {
        interactive-commands notice;
    }
    console {
        interactive-commands warning;
    }
}
```

Configuring System Logging for a Routing Matrix

This section explains how to configure system logging for the T640 Internet routing nodes and TX Matrix platform in a routing matrix. It assumes you are familiar with system logging for single-chassis systems, as described in "Configuring System Logging for a Single-Chassis System" on page 429. For more information about routing matrixes, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852 and the *TX Matrix Platform Hardware Guide.*

To configure system logging for all platforms in a routing matrix, include the syslog statement at the [edit system] hierarchy level on the TX Matrix platform. The syslog statement applies to every platform in the routing matrix.

```
[edit system]
syslog {
    archive {
        files number;
        size size;
        (world-readable | no-world-readable);
    }
    console {
        facility severity;
    }
    file filename {
        facility severity;
        explicit-priority;
        match "regular-expression";
        archive {
            files number;
            size size;
            (world-readable | no-world-readable);
        }
    }
    host (hostname | other-routing-engine | scc-master) {
        facility severity;
        explicit-priority;
        facility-override facility;
        log-prefix string;
        match "regular-expression";
    }
```

```
        source-address source-address;
        time-format (year | millisecond | year millisecond);
        user (username | *) {
            facility severity;
            match "regular-expression";
        }
    }
```

When included in the configuration on the TX Matrix platform, the following configuration statements have the same effect as on a single-chassis system, except that they apply to every platform in the routing matrix:

- archive—Configures the archiving of log files on each platform in the routing matrix. See "Configuring Log File Archiving" on page 438.
- console—Directs the specified messages to the console of each platform in the routing matrix. See "Directing Messages to the Console" on page 433.
- file—Directs the specified messages to a file of the same name on each platform in the routing matrix. See "Directing Messages to a Log File" on page 432.
- match—Limits the set of messages logged to a destination to those that contain (or do not contain) a text string matching a regular expression. See "Using Regular Expressions to Refine the Set of Logged Messages" on page 442.

  The separate match statement at the [edit system syslog host scc-master] hierarchy level applies to messages forwarded from the T640 routing nodes to the TX Matrix platform. See "Configuring Optional Features for Forwarded Messages" on page 451.
- source-address—Sets the source address used for messages directed to the remote machines specified in all host *hostname* statements at the [edit system syslog] hierarchy level for each platform in the routing matrix. The address is not used for messages directed to the other Routing Engine on each platform or to the TX Matrix platform from the T640 routing nodes. See "Specifying an Alternate Source Address" on page 434.
- time-format—Adds the millisecond, year, or both to the timestamp in each message. See "Including the Year or Millisecond in Timestamps" on page 442.
- user—Directs the specified messages to the terminal session of one or more specified users on each platform in the routing matrix that they are logged in to. See "Directing Messages to a User Terminal" on page 433.

The effect of the other statements differs somewhat for a routing matrix than for a single-chassis system. For more information, see the following sections:

- Configuring Message Forwarding in the Routing Matrix on page 448
- Configuring Optional Features for Forwarded Messages on page 451

- Directing Messages to a Remote Destination from the Routing Matrix on page 452
- Configuring System Logging Differently on Each Platform on page 453

Configuring Message Forwarding in the Routing Matrix

By default, the master Routing Engine on each T640 routing node forwards to the master Routing Engine on the TX Matrix platform all messages from all facilities with severity info and higher. To change the set of facilities, the severity level, or both, include the host scc-master statement at the [edit system syslog] hierarchy level on the TX Matrix platform:

```
[edit system syslog]
host scc-master {
    facility severity;
}
```

The setting applies to all T640 routing nodes in the routing matrix.

To disable message forwarding, set the facility to any and the severity level to none.

For the TX Matrix platform to record the messages forwarded by the T640 routing nodes (as well as messages generated on the TX Matrix platform itself), you must also configure system logging on the TX Matrix platform. Direct the messages to one or more destinations by including the appropriate statements at the [edit system syslog] hierarchy level on the TX Matrix platform:

- To a file, as described in "Directing Messages to a Log File" on page 432.
- To the terminal session of one or more specific users (or all users), as described in "Directing Messages to a User Terminal" on page 433.
- To the console, as described in "Directing Messages to the Console" on page 433.
- To a remote machine that is running the syslogd utility or to the other Routing Engine. For more information, see "Directing Messages to a Remote Destination from the Routing Matrix" on page 452.

As previously noted, the configuration statements included on the TX Matrix platform also configure the same destinations on each T640 routing node.

When specifying the severity level for local messages (at the [edit system syslog (file | host | console | user)] hierarchy level) and forwarded messages (at the [edit system syslog host scc-master] hierarchy level), you can set the same severity level for both, set a lower severity level for local messages, or set a higher severity level for local messages.

The following examples describe the consequence of each configuration:

- Messages Logged when Local and Forwarded Severity Level Are the Same on page 449
- Messages Logged when Local Severity Level Is Lower on page 449
- Messages Logged when Local Severity Level Is Higher on page 450

For simplicity, the examples use the any facility in every case. You can also specify different severities for different facilities, with more complex consequences.

Messages Logged when Local and Forwarded Severity Level Are the Same

When the severity level is the same for local and forwarded messages, the log on the TX Matrix platform contains all messages from the logs on the T640 routing nodes. For example, you can specify severity info for the /var/log/messages file, which is the default severity level for messages forwarded by T640 routing nodes:

```
[edit system syslog]
file messages {
    any info;
}
```

Table 29 specifies which messages are included in the logs on the T640 routing nodes and the TX Matrix platform.

Table 29: Example: Local and Forwarded Severity Level Are Both info

Log Location	Source of Messages	Lowest Severity Included
T640 routing node	Local	info
TX Matrix platform	Local	info
	Forwarded from T640 routing nodes	info

Messages Logged when Local Severity Level Is Lower

When the severity level is lower for local messages than for forwarded messages, the log on the TX Matrix platform includes fewer forwarded messages than when the severities are the same. Locally generated messages are still logged at the lower severity level, so their number in each log is the same as when the severities are the same.

For example, you can specify severity notice for the /var/log/messages file and severity critical for forwarded messages:

```
[edit system syslog]
file messages {
    any notice;
}
host scc-master {
    any critical;
}
```

Table 30 specifies which messages are included in the logs on the T640 routing nodes and the TX Matrix platform. The T640 routing nodes forward only those messages with severity critical and higher, so the log on the TX Matrix platform does not include the messages with severity error, warning, or notice that the T640 routing nodes log locally.

Table 30: Example: Local Severity Is notice, Forwarded Severity Is critical

Log Location	Source of Messages	Lowest Severity Included
T640 routing node	Local	notice
TX Matrix platform	Local	notice
	Forwarded from T640 routing nodes	critical

Messages Logged when Local Severity Level Is Higher

When the severity level is higher for local messages than for forwarded messages, the log on the TX Matrix platform includes fewer forwarded messages than when the severities are the same, and all local logs contain fewer messages overall.

For example, you can specify severity critical for the /var/log/messages file and severity notice for forwarded messages:

```
[edit system syslog]
file messages {
    any critical;
}
host scc-master {
    any notice;
}
```

Table 31 specifies which messages are included in the logs on the T640 routing nodes and the TX Matrix platform. Although the T640 routing nodes forward messages with severity notice and higher, the TX Matrix platform discards any of those messages with severity critical or lower (does not log forwarded messages with severity error, warning, or notice). None of the logs include messages with severity error or lower.

Table 31: Example: Local Severity critical, Forwarded Severity is notice

Log Location	Source of Messages	Lowest Severity Included
T640 routing node	Local	critical
TX Matrix platform	Local	critical
	Forwarded from T640 routing nodes	critical

Configuring Optional Features for Forwarded Messages

You can configure additional optional features when specifying how the T640 routing nodes forward messages to the TX Matrix platform, by including statements at the [edit system syslog host scc-master]. To include priority information (facility and severity level) in each forwarded message, include the explicit-priority statement. To insert a string in each forwarded message, include the log-prefix statement. To use regular expression matching to specify more exactly which messages from a facility are forwarded, include the match statement.

```
[edit system syslog host scc-master]
    facility severity;
    explicit-priority;
    log-prefix string;
    match "regular-expression";

}
```

NOTE: You can also include the facility-override statement at the [edit system syslog host scc-master] hierarchy level, but we do not recommend doing so. It is not necessary to use alternate facilities for messages forwarded to the TX Matrix platform, because it runs the JUNOS system logging utility and can interpret the JUNOS-specific facilities. For more information about alternate facilities, see “Changing the Alternate Facility Name for Remote Messages” on page 435.

When you include the log-prefix statement, the string that you define appears in every message forwarded to the TX Matrix platform. For more information, see “Adding a String to System Log Messages” on page 437.

When you include the explicit-priority statement, the messages forwarded to the TX Matrix platform include priority information. For the information to appear in a log file on the TX Matrix platform, you must also include the explicit-priority statement at the [edit system syslog file *filename*] hierarchy level. The log file on each platform in the routing matrix also includes priority information for locally generated messages as a consequence.

To include priority information in messages directed to a remote machine from the routing matrix, also include the explicit-priority statement at the [edit system syslog host *hostname*] hierarchy level. For more information, see “Directing Messages to a Remote Destination from the Routing Matrix” on page 452.

In the following example, the /var/log/messages file on all platforms includes priority information for messages with severity notice and higher from all facilities. The log on the TX Matrix platform also includes messages with those characteristics forwarded from the T640 routing nodes.

```
[edit system syslog]
host scc-master {
    any notice;
    explicit-priority;

}
file messages {
    any notice;
    explicit-priority;
}
```

When you include the match statement, the regular expression that you specify controls which messages from the T640 routing nodes are forwarded to the TX Matrix platform. The regular expression is not applied to messages from the T640 routing notes that are directed to destinations other than the TX Matrix platform. For more information about regular expression matching, see "Using Regular Expressions to Refine the Set of Logged Messages" on page 442.

Directing Messages to a Remote Destination from the Routing Matrix

You can configure a routing matrix to direct system logging messages to a remote machine or the other Routing Engine on each routing platform, just as on a single-chassis system. Include the host statement at the [edit system syslog] hierarchy level on the TX Matrix platform:

```
[edit system syslog]
host (hostname | other-routing-engine) {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
```

The TX Matrix platform directs messages to a remote machine or the other Routing Engine in the same way as a single-chassis system, and the optional statements (explicit-priority, facility-override, log-prefix, match, and source-address) also have the same effect as on a single-chassis system. For more information, see "Directing Messages to a Remote Machine or the Other Routing Engine" on page 433.

For the TX Matrix platform to include priority information when it directs messages that originated on a T640 routing node to the remote destination, you must also include the explicit-priority statement at the [edit system syslog host scc-master] hierarchy level.

The other-routing-engine statement does not interact with message forwarding from the T640 routing nodes to the TX Matrix platform. For example, if you include the statement in the configuration for the Routing Engine in slot 0 (re0), the re0 Routing Engine on each T640 routing node sends messages to the re1 Routing Engine on its platform only. It does not also send messages directly to the re1 Routing Engine on the TX Matrix platform.

Because the configuration on the TX Matrix platform applies to the T640 routing nodes, any T640 routing node that has interfaces for direct access to the Internet also directs messages to the remote machine. The consequences include the following:

- If the T640 routing nodes are configured to forward messages to the TX Matrix platform (as in the default configuration), the remote machine receives two copies of some messages: one directly from the T640 routing node and the other from the TX Matrix platform. Which messages are duplicated depends on whether the severities are the same for local logging and for forwarded messages. For more information, see "Configuring Message Forwarding in the Routing Matrix" on page 448.

- If the source-address statement is configured at the [edit system syslog] hierarchy level, all platforms in the routing matrix use the same source address for messages directed to the remote machine. This is appropriate, as the routing matrix is intended to function as a single routing platform.

- If the log-prefix statement is included, the messages from all platforms in the routing matrix include the same string. You cannot use the string to distinguish between the platforms in the routing matrix.

Configuring System Logging Differently on Each Platform

We recommend that all platforms in a routing matrix use the same configuration, which implies that you include system logging configuration statements on the TX Matrix platform only. In rare circumstances, however, you might choose to log different messages on different platforms. For example, if one platform in the routing matrix is experiencing problems, a Juniper Networks support representative might instruct you to log messages with severity debug on that one platform.

To configure platforms separately, include system logging configuration statements in the appropriate groups at the [edit groups] hierarchy level on the TX Matrix platform:

- To configure settings that apply to the TX Matrix platform but not the T640 routing nodes, include them in the re0 and re1 configuration groups.

- To configure settings that apply to particular T640 routing nodes, include them in the lcc*n*-re0 and lcc*n*-re1 configuration groups, where *n* is the line-card chassis (LCC) index number of the routing node.

When you use configuration groups, do not issue CLI configuration mode commands on the TX Matrix platform that affect the [edit system syslog] hierarchy level. The resulting statements overwrite the statements defined in configuration groups and apply to the T640 routing nodes also. (We further recommend that you do not issue CLI configuration mode commands on the T640 routing nodes at any time.)

For more information about the configuration groups for a routing matrix, see "Creating a Configuration Group" on page 618.

The following sample statements configure the /var/log/messages files on three platforms to include different sets of messages:

- On the TX Matrix platform, local messages with severity info and higher from all facilities. The file does not include messages from the T640 routing nodes, because the host scc-master statement disables message forwarding.
- On the T640 routing node designated LCC0, messages with severity debug from all facilities.
- On the T640 routing node designated LCC1, messages with severity notice from all facilities.

```
[edit groups]
re0 {
    system {
        syslog {
            file messages {
                any info;
            }
            host scc-master {
                any none;
            }
        }
    }
}
re1 {
    ... same statements as for re0...
}
lcc0-re0 {
    system {
        syslog {
            file messages {
                any debug;
            }
        }
    }
}
lcc0-re1 {
    ... same statements as for lcc0-re0...
}
lcc1-re0 {
    system {
        syslog {
            file messages {
                any notice;
            }
        }
    }
}
lcc0-re1 {
    ... same statements as for lcc1-re0...
}
```

Chapter 25

Configuring Miscellaneous System Management Features

This chapter discusses the following topics:

- Configuring the Range of Port Addresses on page 490
- Configuring ARP Learning and Aging on page 490
- Configuring System Alarms to Show Automatically on page 491

Configuring Console and Auxiliary Port Properties

The router's craft interface has two ports—a console port and an auxiliary port—for connecting terminals to the router. The console port is enabled by default, and its speed is 9600 baud. The auxiliary port is disabled by default.

To configure the properties for the console and auxiliary ports, include the ports statement at the [edit system] hierarchy level:

```
[edit system]
ports {
    auxiliary {
      type terminal-type;
    }
    console {
      insecure;
      log-out-on-disconnect;
      type terminal-type;
    }
}
```

By default, the terminal type is unknown, and the terminal speed is 9600 baud for both the console and auxiliary ports. To change the terminal type, include the type statement, specifying a *terminal-type* of ansi, vt100, small-xterm, or xterm. The first three terminal types set a screen size of 80 columns by 24 lines. The last type, xterm, sets the size to 80 columns by 65 rows.

By default, the console session is not logged out when the data carrier is lost on the console modem control lines. To log out the session when the data carrier on the console port is lost, include the log-out-on-disconnect statement.

By default, terminal connections to the console and auxiliary ports are secure. When you configure the console as insecure, root logins are not allowed to establish terminal connections. To disable root login connections to the console and auxiliary ports, include the insecure statement.

Disabling the Sending of Redirect Messages on the Router

By default, the router sends protocol redirect messages. To disable the sending of redirect messages by the router, include the no-redirects statement at the [edit system] hierarchy level:

```
[edit system]
no-redirects;
```

To re-enable the sending of redirect messages on the router, delete the no-redirects statement from the configuration.

To disable the sending of redirect messages on a per-interface basis, include the no-redirects statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family*] hierarchy level, as described in the *JUNOS Network Interfaces Configuration Guide*.

Configuring the Source Address for Locally Generated TCP/IP Packets

By default, the source address included in locally generated Transmission Control Protocol/IP (TCP/IP) packets, such as FTP traffic, and in User Datagram Protocol (UDP) and IP packets, such as Network Time Protocol (NTP) requests, is chosen as the local address for the interface on which the traffic is transmitted. This means that the local address chosen for packets to a particular destination might change from connection to connection based on the interface that the routing protocol has chosen to reach the destination when the connection is established. If multiple equal-cost next hops are present for a destination, locally generated packets use the lo0 address as a source.

To configure the software to select a fixed address to use as the source for locally generated IP packets, include the default-address-selection statement at the [edit system] hierarchy level:

```
[edit system]
default-address-selection;
```

If you include the default-address-selection statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. The default address is usually an address configured on the lo0 loopback interface. For example, if you specified that SSH and telnet use a particular address, but you also have default-address selection configured, the system default address is used. For more information about how the default address is chosen, see the *JUNOS Network Interfaces Configuration Guide*.

For IP packets sent by IP routing protocols—including Open Shortest Path First (OSPF), Routing Information Protocol (RIP), Resource Reservation Protocol (RSVP), and the multicast protocols, but not including Intermediate System-to-Intermediate System (IS-IS)—the local address selection is often constrained by the protocol specification so that the protocol operates correctly. When this constraint exists in the routing protocol, the packet's source address is unaffected by the presence of the default-address-selection statement in the configuration. For protocols in which the local address is unconstrained by the protocol specification, for example, internal Border Gateway Protocol (IBGP) and multihop external BGP (EBGP), if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same method as other locally generated IP packets.

Configuring the Router or Interface to Act as a DHCP/BOOTP Relay Agent

To configure a router or interface to act as a bootstrap protocol (DHCP/BOOTP) relay agent, you include statements at the [edit forwarding-options helpers] hierarchy level. For more information, see the *JUNOS Policy Framework Configuration Guide*.

For J-series Services Routers, you can configure a router or interface as a DHCP server by including statements at the [edit system services] hierarchy level. For more information, see "Configuring a DHCP Server" on page 459.

NOTE: You cannot configure a router or interface as a DHCP server and a BOOTP relay agent at the same time.

Configuring System Services

For security reasons, remote access to the router is disabled by default. You must configure the router explicitly so that users on remote systems can access it. The router can be accessed from a remote system by means of the DHCP, finger, FTP, JUNOScript clear-text, JUNOScript Secure Sockets Layer (SSL), rlogin, SSH, and telnet services.

NOTE: To protect the system resources, you can configure connection limits for each service. In addition, JUNOS software limits the number of processes a single user may have. For a service connection to succeed, the cumulative number of connections must be within the configured limits and the user must not have exceeded the process quota.

This section discusses the following topics:

- Configuring a DHCP Server on page 459
- Configuring Finger Service on page 475
- Configuring FTP Service on page 476
- Configuring JUNOScript Clear-Text Service on page 476
- Configuring JUNOScript SSL Service on page 477
- Configuring SSH Service on page 477
- Configuring telnet Service on page 479

JUNOS-FIPS disables many of the usual JUNOS protocols and services. Services that cannot be configured include DHCP, finger, FTP, JUNOScript clear text, rlogin, rsh, and telnet.

Configuring a DHCP Server

The Dynamic Host Configuration Protocol (DHCP) server provides a framework for passing configuration information to client hosts (such as PCs) on a TCP/IP network. A router or interface that acts as a DHCP server can allocate network IP addresses and deliver configuration settings to client hosts without user intervention. DHCP access service minimizes the overhead required to add clients to the network by providing a centralized, server-based setup. You do not have to manually create and maintain IP address assignments for clients. DHCP is defined in RFC 2131, *Dynamic Host Configuration Protocol*.

NOTE: The DHCP server is supported in J-series Services Routers and is compatible with the autoinstallation feature.

To configure a J-series Services Router to accept DHCP as an access service, include the dhcp statement at the [edit system services] hierarchy level:

```
[edit system services]
dhcp {
    boot-file filename;
    boot-server (address | hostname);
    domain-name domain-name;
    domain-search [domain-list];
    default-lease-time seconds;
    maximum-lease-time seconds;
    name-server {
      address;
    }
    option {
      [ (id-number option-type option-value) | (id-number array option-type
        option-values) ] ;
    }
    pool (address /prefix-length) {
      address-range {
        low address;
        high address;
      }
      exclude-address {
        address;
      }
    }
    router {
      address;
    }
    static-binding MAC-address {
      fixed-address {
        address;
      }
      host hostname;
      client-identifier (ascii client-id | hexadecimal client-id);
    }
    server-identifier address;
    wins-server {
      address;
    }
}
```

To configure DHCP properties, go to one or more of the following sections:

- DHCP Overview on page 460
- Configuring Address Pools on page 466
- Configuring Manual (Static) Bindings on page 467
- Specifying DHCP Lease Times on page 468
- Configuring a Boot File and Boot Server on page 468
- Configuring a DHCP Server Identifier on page 469
- Configuring a Domain Name and Domain Search List on page 470
- Configuring Routers Available to the Client on page 471
- Creating User-Defined DHCP Options on page 471

DHCP Overview

DHCP access service consists of two components: a protocol for delivering host-specific configuration information from a server to a client host and a method for allocating network addresses to a client host. The client sends a message to request configuration information. A DHCP server sends the configuration information back to the client.

With DHCP, clients can be assigned a network address for a fixed *lease*, enabling serial reassignment of network addresses to different clients. A DHCP server leases IP addresses for specific times to various clients. If a client does not use its assigned address for some period of time, the DHCP server can assign that IP address to another host. When assignments are made or changed, the DHCP server updates information in the DNS server. The DHCP server provides clients with their previous lease assignments whenever possible.

A DHCP server provides persistent storage of network parameters for clients. Because DHCP is an extension of BOOTP, DHCP servers can handle BOOTP requests.

The DHCP server includes IPv4 address assignment and commonly used DHCP options. The server is compatible with DHCP servers from other vendors on the network. The server does not support IPv6 address assignment, user class-specific configuration, DHCP failover protocol, dynamic DNS updates, or VPN connections. The DHCP server is not supported in JUNOS-FIPS.

NOTE: You cannot configure a router as a DHCP server and a BOOTP relay agent at the same time.

Network Address Assignments (Allocating a New Address)

To receive configuration information and a network address assignment, a DHCP client negotiates with DHCP servers in a series of messages. The following steps show the messages exchanged between a DHCP client and servers to allocate a new network address. When allocating a new network address, the DHCP process can involve more than one server, but only one server is selected by the client.

1. When a client computer is started, it broadcasts a DHCPDISCOVER message on the local subnet requesting a DHCP server. This request includes the hardware address of the requesting client.

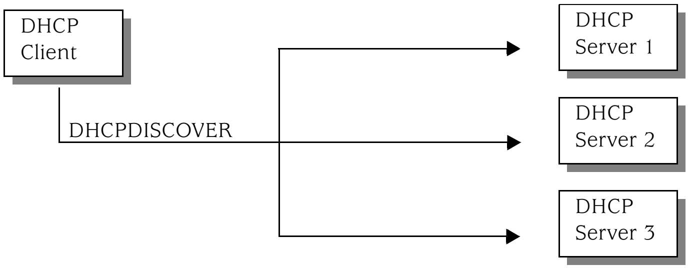

2. Each DHCP server receiving the broadcast sends a DHCPOFFER message to the client offering an IP address for a set period of time, known as the *lease period*.

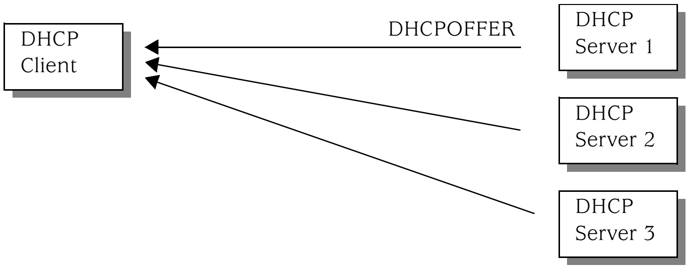

3. The client receives one or more DHCPOFFER messages from one or more servers and selects one of the offers received. Normally, a client looks for the longest lease period.

4. The client broadcasts a DHCPREQUEST message indicating the client has selected an offered leased IP address and identifies the selected server.

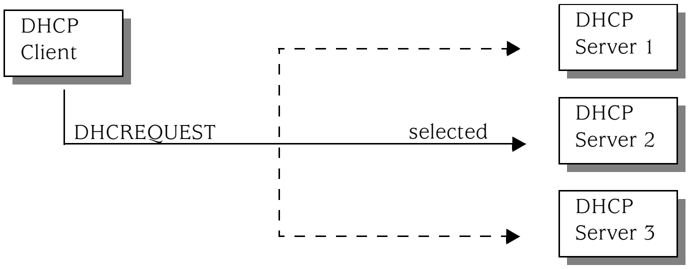

5. Those servers not selected by the DHCPREQUEST message return the unselected IP addresses to the pool of available addresses.

6. The selected DHCP server sends a DHCPACK acknowledgement that includes configuration information such as the IP address, subnet mask, default gateway, and the lease period.

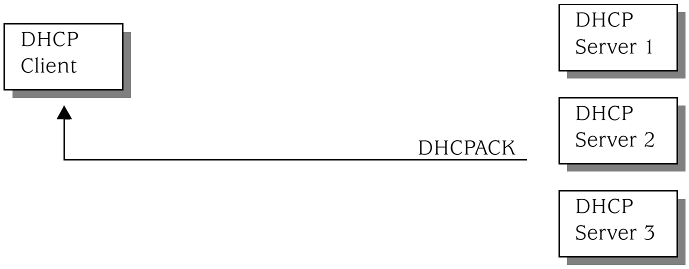

The information offered by the server is configurable. See "Configuring a DHCP Server" on page 459 for more information.

7. The client receives the DHCPACK message with configuration information. The process is complete. The client is configured and has access to the network.

   - If the client receives a DHCPNAK message (for example, if the client has moved to a new subnet), the client restarts the negotiation process.

   - The client can relinquish its lease on a network address by sending a DHCPRELEASE message to the server (for example, when the client is restarted). When the server receives the DHCPRELEASE message, it marks the lease as free and the IP address becomes available again.

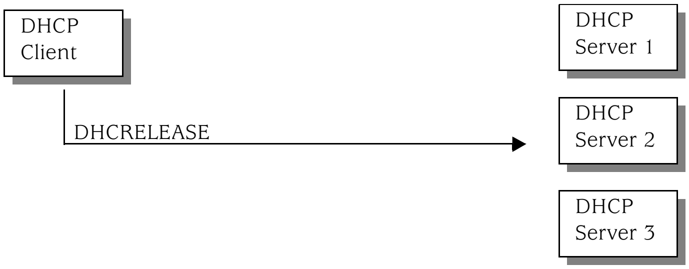

Network Address Assignments (Reusing a Previously Assigned Address)

To reuse a previously allocated network address:

1. A client that previously had a lease broadcasts a DHCPREQUEST message on the local subnet.

2. The server with knowledge of the client's configuration responds with a DHCPACK message.

3. The client verifies the DHCP configuration information sent by the server and uses this information to reestablish the lease.

Static and Dynamic Bindings

DHCP supports both dynamic and static bindings. For dynamic bindings, IP addresses are assigned to clients from a pool of addresses. Static bindings provide configuration information for a specific client and can include one or more fixed IP addresses for the client. You can configure a DHCP server to include both address pools and static bindings. For any individual client, static bindings take priority over address pools.

Compatibility with Autoinstallation

The DHCP server is compatible with the autoinstallation feature on J-series Services Routers. The server automatically checks autoinstallation settings for conflicts and gives autoinstallation settings priority over corresponding DHCP settings. For example, an IP address set by autoinstallation takes priority over an IP address set by the DHCP server.

NOTE: The autoinstallation feature includes a fixed address pool and a fixed lease time. With DHCP, you can create address pools and modify lease times.

Conflict Detection and Resolution

When a client receives an IP address from the DHCP server, the client performs a series of ARP tests to verify that the IP address is available and no conflicts exist. If the client detects an address conflict, the client notifies the DHCP server about the conflict and may request another IP address from the DHCP server.

The DHCP server keeps a log of all conflicts and removes addresses with conflicts from the pool. These addresses remain excluded until you manually clear the conflicts list with the clear system services dhcp conflict command. For more information on this command, see the *JUNOS System Basics and Services Command Reference*.

DHCP Statement Hierarchy and Inheritance

DHCP configuration statements are organized hierarchically. Statements at the top of the hierarchy apply to the DHCP server and network, branches contain statements that apply to address pools in a subnetwork, and leaves contain statements that apply to static bindings for individual clients. See Table 32.

The pool and static-binding statements appear at the [edit system services dhcp] hierarchy level. You can include the remaining statements at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

Table 32: Pool and Binding Statements

Statement	Description	Hierarchy Level
pool	Configure a pool of IP addresses for DHCP clients on a subnet. When a client joins the network, the DHCP server dynamically allocates an IP address from this pool.	[edit system services dhcp]
static-binding	Sets static bindings for DHCP clients. A static binding is a mapping between a fixed IP address and the client’s MAC address.	

To minimize configuration changes, include common configuration statements shown in Table 33 (for example, the domain-name statement) at the highest applicable level of the hierarchy (network or subnetwork). Configuration statements at lower levels of the hierarchy override statements inherited from a higher level. For example, if a statement appears at both the [edit system services dhcp] and [edit system services dhcp pool] hierarchy levels, the value assigned to the statement at the [edit system services dhcp pool] level takes priority.

Table 33: Common Configuration Statements

Statement	Description	Hierarchy Level
boot-file	Set the boot filename advertised to clients. The client uses the boot image stored in the boot file to complete configuration.	[edit system services dhcp] [edit system services dhcp pool] [edit system services dhcp static-binding]
boot-server	Set the server that contains the boot file.	
default-lease-time	Set the default lease time assigned to any client that does not request a specific lease time.	
domain-name	Configure the name of the domain in which clients will search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified.	
domain-search	Defines a domain search list.	
maximum-lease-time	Sets the maximum lease time allowed by the server.	
name-server	Specifies the DNS server that maintains the database of client name to IP address mappings.	
option	Configures user-defined DHCP options.	
router	Specifies IP address for routers on the client’s subnetwork. Routers are listed in order of preference.	
server-identifier	Sets the IP address of the DHCP server.	

Configuring Address Pools

For dynamic bindings, set aside a pool of IP addresses that can be assigned to clients. Addresses in a pool must be available to clients on the same subnet.

To configure an address pool, include the following statements at the [edit system services dhcp] hierarchy level:

```
[edit system services dhcp]
pool (address /prefix-length) {
    address-range {
      low address;
      high address;
      }
    exclude-address {
      address;
    }
}
```

The pool definition must include the client subnet number and prefix length (in bits). Optionally, the definition can include an address range and a list of excluded addresses.

The address-range statement defines the lowest and highest IP addresses in the pool that are available for dynamic address assignment. This statement is optional. If no range is specified, the pool will use all available addresses within the subnet specified. (Broadcast addresses, interface addresses, and excluded addresses are not available.)

The exclude-address statement specifies addresses within the range that are not used for dynamic address assignment. You can exclude one or more addresses within the range. This statement is optional.

The following is an example of a pool configuration.

```
[edit system services dhcp]
pool 3.3.3.0/24 {
    address-range low 3.3.3.2 high 3.3.3.254;
    exclude-address {
        3.3.3.33;
    }
}
```

For dynamic address assignment, configure an address pool for each client subnet the DHCP server supports. You can configure multiple address pools for a DHCP server, but only one address range per pool is supported.

DHCP maintains the state information for all pools configured. Clients are assigned addresses from pools with subnets that match the interface on which the DHCPDISCOVER packet is received. When more than one pool exists on the same interface, addresses are assigned on a rotating basis from all available pools.

Configuring Manual (Static) Bindings

Static bindings provide configuration information for specific clients. This information can include one or more fixed Internet addresses, the client hostname, and a client identifier.

To configure static bindings, include the following statements at the [edit system services dhcp] hierarchy level:

```
[edit system services dhcp]
static-binding MAC-address {
      fixed-address {
         address;
      }
      host client-hostname;
      client-identifier (ascii client-id | hexadecimal client-id);client-id;
}
```

A static binding defines a mapping between a fixed IP address and the client's MAC address.

The MAC-address statement specifies the MAC address of the client. This is a hardware address that uniquely identifies each client on the network.

The fixed-address statement specifies the fixed IP address assigned to the client. Typically a client has one address assigned, but you can assign more.

The host statement specifies the hostname of the client requesting the DHCP server. The name can include the local domain name. Otherwise, the name is resolved based on the domain-name statement.

The client-identifier statement is used by the DHCP server to index the database of address bindings. The client identifier is either an ASCII string or hexadecimal digits. It can include a type-value pair as specified in RFC 1700, *Assigned Numbers*. Either a client identifier or the client's MAC address must be configured to uniquely identify the client on the network.

The following is an example of a static binding configuration:

```
[edit system services dhcp]
static-binding 00:0d:56:f4:01:ab {
      fixed-address {
         5.5.5.5;
         6.6.6.6;
      }
      host-name "another-host.domain.tld";
      client-identifier hexadecimal 01001122aabbcc;
}
```

Specifying DHCP Lease Times

For clients that do not request a specific lease time, the default lease time is one day. You can configure a maximum lease time for IP address assignments or change the default lease time.

To configure lease times, include the maximum-lease-time and default-lease-time statements:

```
maximum-lease-time seconds;
default-lease-time seconds;
```

You can include these statements at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

Lease times defined for static bindings and address pools take priority over lease times defined at the [edit system services dhcp] hierarchy level.

The maximum-lease-time statement configures the maximum length of time in seconds for which a client can request and hold a lease. If a client requests a lease longer than the maximum specified, the lease is granted only for the maximum time configured on the server. After a lease expires, the client must request a new lease.

NOTE: Maximum lease times do not apply to dynamic BOOTP leases. These leases are not specified by the client and can exceed the maximum lease time configured.

The following example shows a configuration for maximum and default lease times:

```
[edit system services dhcp]
maximum-lease-time 7200;
default-lease-time 3600;
```

Configuring a Boot File and Boot Server

When a client starts, it contacts a boot server to download the boot file.

To configure a boot file and boot server, include the boot-file and boot-server statements:

```
boot-file filename;
boot-server address;
```

You can include these statements at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

After a client receives a DHCPOFFER response from a DHCP server, the client can communicate directly with the boot server (instead of the DHCP server) to download the boot file. This minimizes network traffic and enables you to specify separate boot server/file pairs for each client pool or subnetwork.

The boot-file statement configures the name and location of the initial boot file that the DHCP client loads and executes. This file stores the boot image for the client. In most cases, the boot image is the operating system the client uses to load.

The boot-server statement configures the IP address of the TFTP server that contains the client's initial boot file. You must configure an IP address (not a hostname) for the server.

You must configure at least one boot file and boot server. Optionally, you can configure multiple boot files and boot servers. For example, you might configure two separate boot servers and files: one for static binding and one for address pools. Boot file configurations for pools or static bindings take precedence over boot file configurations at the [edit system services dhcp] hierarchy level.

The following example specifies a boot file and server for an address pool:

```
[edit system services dhcp]
pool 4.4.4.0/24 {
    boot-file "boot.client";
    boot-server 4.4.4.1;
}
```

Configuring a DHCP Server Identifier

The host running the DHCP server itself must use a manually assigned, static IP address. It cannot send a request and receive an IP address from itself or another DHCP server.

To configure a server identifier, include the server-identifier statement:

```
server-identifier address;
```

You can include this statement at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

The server-identifier statement specifies the IP address of the DHCP server. The host must be a TFTP server that is accessible by all clients served within a range of IP addresses (based on either an address pool or static binding).

The following example shows a DHCP server identifier configured for an address pool:

```
[edit system services dhcp]
pool 3.3.3.0/24 {
        address-range low 3.3.3.2 high 3.3.3.254;
        exclude-address {
           3.3.3.33;
        }
        router {
           3.3.3.1;
        }
        server-identifier 3.3.3.1;
```

Configuring a Domain Name and Domain Search List

To configure the name of the domain in which clients will search for a DHCP server host, include the domain-name statement:

```
domain-name domain;
```

You can include this statement at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

The domain-name statement sets the domain name that is appended to hostnames that are not fully qualified. This statement is optional. If you do not configure a domain name, the default is the client's current domain.

To configure a domain search list, include the domain-search statement:

```
domain-search [domain-list];
```

You can include this statement at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

The domain-search statement sets the order in which clients append domain names when searching for the IP address of a host. You can include one or more domain names in the list. For more information, see RFC 3397, *Dynamic Host Configuration Protocol (DHCP) Domain Search Option*.

The domain-search statement is optional, if you do not configure a domain search list, the default is the client's current domain.

Configuring Routers Available to the Client

After a DHCP client loads the boot image and has booted, the client sends packets to a router.

To configure routers available to the client, include the router statement:

```
router {
  address1;
  address2;
}
```

You can include this statement at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

The router statement specifies a list of IP addresses for routers on the client's subnet. List routers in order of preference. You must configure at least one router for each client subnet.

The following example shows routers configured at the [edit system services dhcp] hierarchy level:

```
[edit system services]
dhcp {
    router {
      6.6.6.1;
      7.7.7.1;
    }
}
```

Creating User-Defined DHCP Options

You can configure one or more user-defined options that are not included in the JUNOS default implementation of the DHCP server. For example, if a client requests a DHCP option that is not included in the DHCP server, you can create a user-defined option that enables the server to respond to the client's request.

To configure a user-defined DHCP option, include the option statement:

```
option {
    [ (id-number option-type option-value) | (id-number array option-type
        option-values) ] ;
}
```

The option statement includes:

- *id-number*—Any whole number. The ID number is used to index the option and must be unique across a DHCP server.
- *option-type*—Any of the following types: flag, byte, string, short, unsigned-short, integer, unsigned-integer, ip-address.
- array—An option can include an array of values.

- *option-value*—Value associated with an option. The option value must be compatible with the option type (for example, an On or Off value for a flag type).

You can include this statement at the following hierarchy levels:

```
[edit system services dhcp]
[edit system services dhcp pool]
[edit system services dhcp static-binding]
```

The following example shows user-defined DHCP options:

```
[edit system services dhcp]
option 19 flag off;                   # 19: "IP Forwarding" option
option 40 string "domain.tld";        # 40: "NIS Domain" option
option 16 ip-address 3.3.3.33;       # 16: "Swap Server" option
```

User-defined options that conflict with DHCP configuration statements are ignored by the server. For example, in the following configuration, the DHCP server ignores the user-defined option 3 router statement and uses the edit system services dhcp router statement instead:

```
[edit system services]
dhcp {
    option 3 router 7.7.7.2;      # 3: "Default Router" option
    router {
      7.7.7.1;
    }
}
```

Example: Complete DHCP Server Configuration

This section shows a complete DHCP server configuration with address pools, static bindings, and user-defined options.

The following example shows statements at the [edit interfaces] hierarchy level. The interface's primary address (3.3.3.1/24) has a corresponding address pool (3.3.3.0/24) defined at the [edit system services] hierarchy level.

```
[edit interfaces]
fe-0/0/1 {
   unit 0 {
      family inet {
         address 3.3.3.1/24;
      }
   }
}
```

NOTE: You can configure a DHCP server only on an interface's primary IP address.

Statements at the [edit system services] hierarchy level include:

```
[edit system services]
dhcp {
    domain-name "domain.tld";
    maximum-lease-time 7200;
    default-lease-time 3600;
    name-server {
        6.6.6.6;
        6.6.6.7;
    }
    domain-search [ subnet1.domain.tld subnet2.domain.tld ];
    wins-server {
        7.7.7.7;
        7.7.7.9;
    }
    router {
        6.6.6.1;
        7.7.7.1;
    }
    option 19 flag off;                # 19: "IP Forwarding" option
    option 40 string "domain.tld";     # 40: "NIS Domain" option
    option 16 ip-address 3.3.3.33;     # 16: "Swap Server" option
    pool 3.3.3.0/24 {
        address-range low 3.3.3.2 high 3.3.3.254;
        exclude-address {
            3.3.3.33;
        }
        router {
            3.3.3.1;
        }
        server-identifier 3.3.3.1;
    }
    pool 4.4.4.0/24 {
        boot-file "boot.client";
        boot-server 4.4.4.1;
    }
    static-binding 00:0d:56:f4:20:01 {
        fixed-address 4.4.4.4;
        host-name "host.domain.tld";
    }
    static-binding 00:0d:56:f4:01:ab {
        fixed-address {
            5.5.5.5;
            6.6.6.6;
        }
        host-name "another-host.domain.tld";
        client-identifier "01aa.001a.bc65.3e";
    }
}
```

Example: Viewing DHCP Bindings

Use the CLI command show system services dhcp binding to view information about DHCP address bindings, lease times, and address conflicts.

The following example shows the binding type and lease expiration times for IP addresses configured on a router that supports a DHCP server:

```
user@host> show system services dhcp binding

IP Address      Hardware Address    Type            Lease expires at
192.168.1.2     00:a0:12:00:12:ab   static          never
192.168.1.3     00:a0:12:00:13:02   dynamic         2004-05-03 13:01:42 PDT
```

Enter an IP address to show binding for a specific IP address:

```
user@host> show system services dhcp binding 192.168.1.3

DHCP binding information:
    IP address          192.168.1.3
    Hardware address    00:a0:12:00:12:ab
    Client identifier
        61 63 65 64 2d 30 30 3a 61 30 3a 31 32 3a 30 30 aced-00:a0:12:00
        3a 31 33 3a 30 32

Lease information:
    Type                dynamic
    Obtained at         2004-05-02 13:01:42 PDT
    Expires at          2004-05-03 13:01:42 PDT
```

Use the detail option to show detailed binding information:

```
user@host> show system services dhcp binding 192.168.1.3 detail

DHCP binding information:
    IP address              192.168.1.3
    Hardware address        00:a0:12:00:12:ab
    Pool                    192.168.1.0/24
    Interface               fe-0/0/0, relayed by 192.168.4.254

Lease information:
    Type                    dynamic
    Obtained at             2004-05-02 13:01:42 PDT
    Expires at              2004-05-03 13:01:42 PDT

DHCP options:
    name-server foo.mydomain.tld
    domain-name mydomain.tld
    option 19 flag off
```

Example: Viewing DHCP Address Pools

Use the CLI command show system services dhcp pools to view information about DHCP address pools.

The following example show address pools configured on a DHCP server:

```
user@ host > show system services dhcp pools
Pool name       Low address     High address    Excluded addresses
10.40.1.0/24    10.40.1.1       10.40.1.254     10.40.1.254
```

Example: Viewing and Clearing DHCP Conflicts

When the DHCP server provides an IP address, the client performs an ARP check to make sure the address is not being used by another client and reports any conflicts back to the server. The server keeps track of addresses with conflicts and removes them from the address pool. Use the CLI command show system services dhcp conflict to show conflicts.

```
user@host> show system services dhcp conflict

Detection time              Detection method    Address
2004-08-03 19:04:00 PDT     client              192.168.1.5
2004-08-04 04:23:12 PDT     ping                192.168.1.8
```

Use the clear system services dhcp conflicts command to clear the conflicts list and return IP addresses to the pool. The following command shows how to clear an address on the server that has a conflict:

```
user@host> clear system services dhcp conflict 192.168.1.5
```

NOTE: For more information about CLI commands you can use with the DHCP server, see the *JUNOS System Basics and Services Command Reference*.

Configuring Finger Service

To configure the router to accept finger as an access service, include the finger statement at the [edit system services] hierarchy level:

```
[edit system services]
finger {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

You can optionally specify the maximum number of concurrently established connections and the maximum number of connections attempted per minute.

- connection-limit *limit*—The maximum number of established connections (1 through 250) per minute. By default, the maximum number of establish connections is set to 75 connections per minute.

- rate-limit *limit*—The maximum number of connection attempts allowed per minute (1 through 250). By default, the maximum number of connection limits is set to 150 per minute.

JUNOS-FIPS does not support finger as an access service.

Configuring FTP Service

To configure the router to accept the FTP as an access service, include the ftp statement at the [edit system services] hierarchy level:

```
[edit system services]
ftp {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

You can optionally specify the maximum number of concurrently established connections and the maximum number of connections attempted per minute.

- connection-limit *limit*—The maximum number of established connections (1 through 250) per minute. By default, the maximum number of establish connections is set to 75 connections per minute.

- rate-limit *limit*—The maximum number of connection attempts allowed per minute (1 through 250). By default, the maximum number of connection limits is set to 150 per minute.

JUNOS-FIPS does not support FTP as an access service.

Configuring JUNOScript Clear-Text Service

To configure the router to accept JUNOScript clear text as an access service, include the xnm-clear-text statement at the [edit system services] hierarchy level:

```
[edit system services]
xnm-clear-text {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

You can optionally specify the maximum number of concurrently established connections and the maximum number of connections attempted per minute.

- connection-limit *limit*—The maximum number of established connections (1 through 250) per minute. By default, the maximum number of establish connections is set to 75 connections per minute.

- rate-limit *limit*—The maximum number of connection attempts allowed per minute (1 through 250). By default, the maximum number of connection limits is set to 150 per minute.

JUNOS-FIPS does not support JUNOScript clear text as an access service.

Configuring JUNOScript SSL Service

To configure the router to accept JUNOScript SSL as an access service, include the xnm-ssl statement at the [edit system services] hierarchy level:

```
[edit system services]
xnm-ssl {
    <connection-limit limit>;
    <rate-limit limit>;
    <local-certificate name>;
}
```

You can optionally specify the maximum number of concurrently established connections and the maximum number of connections attempted per minute.

- connection-limit *limit*—The maximum number of established connections (1 through 250) per minute. By default, the maximum number of establish connections is set to 75 connections per minute.
- local-certificate *name*—Name of the local X.509 certificate to use. You must import the certificate before you reference it. For more information about how to import an SSL certificate, see “Using JUNOScript SSL Service” on page 757.
- rate-limit *limit*—The maximum number of connection attempts allowed per minute (1 through 250). By default, the maximum number of connection limits is set to 150 per minute.

Configuring SSH Service

To configure the router to accept SSH as an access service, include the SSH statement at the [edit system services] hierarchy level:

```
[edit system services ]
ssh {
    root-login (allow | deny | deny-password);
    protocol-version [v1 v2];
    <connection-limit limit>;
    <rate-limit limit>;
}
```

You can optionally specify the maximum number of concurrently established connections and the maximum number of connections attempted per minute.

- connection-limit *limit*—The maximum number of established connections (1 through 250) per minute. By default, the maximum number of establish connections is set to 75 connections per minute.
- rate-limit *limit*—The maximum number of connection attempts allowed per minute (1 through 250). By default, the maximum number of connection limits is set to 150 per minute.

This section includes the following topics:

- Configuring the Root Login on page 478
- Configuring the SSH Protocol Version on page 478

Configuring the Root Login

By default, users are allowed to log in to the router as root through SSH. To control user access through SSH, include the root-login statement at the [edit systems services ssh] hierarchy level:

```
[edit system services ssh]
root-login (allow | deny | deny-password);
```

allow—Allows users to log in to the router as root through SSH. The default is allow.

deny—Disables users from logging in to the router as root through SSH.

deny-password—Allows users to log in to the router as root through SSH when the authentication method (for example, RSA) does not require a password.

NOTE: The root-login and protocol-version statements are supported in JUNOS Release 5.0 and later. If you downgrade to a release prior to release 5.0, the root-login and protocol-version statements are ignored if they are present in the configuration file.

Configuring the SSH Protocol Version

By default, version 2 of the SSH protocol is enabled. To configure the router to use only version 1of the SSH protocol, include the protocol-version statement and specify v1 at the [edit system services ssh] hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 ];
```

To configure the router to use version 1 and 2 of the SSH protocol, include the protocol-version statement and specify v1 and v2 at the [edit system services ssh] hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 v2 ];
```

You can specify v1, v2, or both versions [v1 v2] of the SSH protocol. The default is v2.

NOTE: The root-login and protocol-version statements are supported in JUNOS Release 5.0 and later. If you downgrade to a release prior to release 5.0, the root-login and protocol-version statements are ignored if they are present in the configuration file.

Configuring telnet Service

To configure the router to accept telnet as an access service, include the telnet statement at the [edit system services] hierarchy level:

```
[edit system services]
telnet {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

You can optionally specify the maximum number of concurrently established connections and the maximum number of connections attempted per minute.

- connection-limit *limit*—The maximum number of established connections (1 through 250) per minute. By default, the maximum number of establish connections is set to 75 connections per minute.

- rate-limit *limit*—The maximum number of connection attempts allowed per minute (1 through 250). By default, the maximum number of connection limits is set to 150 per minute.

JUNOS-FIPS does not support telnet as an access service.

Configuring Console Access to PICs

By default, there is no password setting for console access. To configure console access to the Physical Interface Cards (PICs), include the pic-console-authentication statement at the [edit system] hierarchy level:

```
[edit system]
pic-console-authentication {
    (encrypted-password "password" | plain-text-password);
}
```

encrypted-password "*password*"—Use Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.

plain-text-password—Use a plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password. For information about how to create plain-text passwords, see “Specifying Plain-Text Passwords” on page 17.

Configuring a System Login Message

By default, no login message is displayed. To configure a system login message, include the message statement at the [edit system login] hierarchy level:

```
[edit system login]
message text;
```

If the message text contains any spaces, enclose it in quotation marks.

A system login message appears before the user logs in. A system login announcement appears after the user logs in. See "Configuring a System Login Announcement" on page 480.

Configuring a System Login Announcement

By default, no login announcement is displayed. To configure a system login announcement, include the message statement at the [edit system login] hierarchy level:

```
[edit system login]
announcement text;
```

If the announcement text contains any spaces, enclose it in quotation marks.

A system login announcement appears after the user logs in. A system login message appears before the user logs in. See "Configuring a System Login Message" on page 480.

Configuring JUNOS Software Processes

By default, all JUNOS software processes are enabled on the router. To control the software processes on the router, you can do the following:

- Disabling JUNOS Software Processes on page 481
- Configuring Failover to Backup Media if a Software Process Fails on page 481

Disabling JUNOS Software Processes

CAUTION: Never disable any of the software processes unless instructed to do so by a Customer Support engineer.

To disable a software process, specify the appropriate option in the processes statement at the [edit system] hierarchy level:

```
[edit system]
processes {
    adaptive-services (enable | disable);
    alarm-control (enable | disable);
    chassis-control (enable | disable);
    class-of-service (enable | disable);
    craft-control (enable | disable);
    dhcp (enable | disable);
    disk-monitoring (enable | disable);
    ecc-error-logging (enable | disable);
    firewall (enable | disable);
    inet-process (enable | disable);
    interface-control (enable | disable);
    kernel-replication (enable | disable);
    l2tp-service (enable | disable);
    link-management (enable | disable);
    mib-process (enable | disable);
    network-access (enable | disable);
    ntp (enable | disable);
    pgm (enable | disable);
    pic-services-logging (enable | disable);
    pppoe (enable | disable);
    redundancy-device (enable | disable);
    remote-operations (enable | disable);
    routing (enable | disable);
    sampling (enable | disable);
    service-deployment (enable | disable);
    snmp (enable | disable);
    usb-control (enable | disable);
    watchdog (enable | disable);
    web-management (enable | disable);
    timeout seconds;
}
```

Configuring Failover to Backup Media if a Software Process Fails

For routers with redundant Routing Engines, in the event that a software process fails repeatedly, you can configure the router to switch to backup media containing an alternate version of the system, either the alternate media or the other Routing Engine. To configure the switch to the backup media, include the failover statement at the [edit system processes *process-name*] hierarchy level:

```
[edit system processes]
process-name failover (alternate-media | other-routing-engine);
```

process-name is one of the valid process names. If this statement is configured for a process, and that process fails four times within thirty seconds, the router reboots from either the alternative media or the other Routing Engine.

Configuring the Password on the Diagnostics Port

If you have been asked by Customer Support personnel to connect a physical console to a control board or forwarding component on the router (such as the System Control Board [SCB], System and Switch Board [SSB], or Switching and Forwarding Module [SFM]) to perform diagnostics, you can configure a password on the diagnostics port. This password provides an extra level of security.

To configure a password on the diagnostics port, include the diag-port-authentication statement at the [edit system] hierarchy level:

```
[edit system]
diag-port-authentication (encrypted-password "password" | plain-text-password);
```

You can use an MD5 password, or you can enter a plain-text password that the JUNOS software encrypts (using MD5-style encryption) before it places it into the password database. For an MD5 password, specify the password in the configuration.

If you configure the plain-text-password option, the CLI prompts you for the password. For information about how to create a plain-text passwords, see "Specifying Plain-Text Passwords" on page 17.

For routers that have more than one SSB, the same password is used for both SSBs.

Saving Core Files from JUNOS Processes

By default, when an internal JUNOS process generates a core file, the file and associated context information are saved for debugging purposes in a compressed tar file named /var/tmp/*process-name*.core.*core-number*.tgz. The contextual information includes the configuration and log messages files.

To disable the saving of core files and associated context information, include the no-saved-core-context statement at the [edit system] hierarchy level:

```
[edit system]
no-saved-core-context;
```

To save the core files only, include the saved-core-files statement, specifying the number of files to save at the [edit system] hierarchy level:

```
[edit system]
saved-core-files saved-core-files;
```

saved-core-files is the number of core files to save and can be a value from 1 through 64.

To save the core files along with the contextual information, include the saved-core-context statement at the [edit system] hierarchy level:

```
[edit system]
saved-core-context;
```

Configuring a Router to Transfer its Configuration to an Archive Site

If you want to back up your router's current configuration to an archive site, you can configure the router to transfer its currently active configuration by FTP periodically or after each commit.

To configure the router to transfer its currently active configuration to an archive site, include statements at the [edit system archival configuration] hierarchy level:

```
[edit system archival configuration]
transfer-interval interval;
transfer-on-commit;
archive-sites {
        ftp://<username>:<password>@<host>:<port>/<url-path>;
}
```

This section includes the following topics:

- Configuring the Transfer Interval on page 483
- Configuring Transfer on Commit on page 483
- Configuring Archive Sites on page 484

Configuring the Transfer Interval

To configure the router to periodically transfer its currently active configuration to an archive site, include the transfer-interval statement at the [edit system archival configuration] hierarchy level:

```
[edit system archival configuration]
transfer-interval interval;
```

The *interval* is a period of time ranging from 15 through 2880 minutes.

Configuring Transfer on Commit

To configure the router to transfer its currently active configuration to an archive site each time you commit a candidate configuration, include the transfer-on-commit statement at the [edit system archival configuration] hierarchy level:

```
[edit system archival configuration]
transfer-on-commit;
```

Configuring Archive Sites

When you configure the router to transfer its configuration files, you specify an archive site to which the files are transferred. If you specify more than one archive site, the router attempts to transfer to the first archive site in the list, moving to the next site only if the transfer fails.

To configure the archive site, include the archive-sites statement at the [edit system archival configuration] hierarchy level:

```
[edit system archival configuration]
archive-sites {
    ftp://username@host:<port>url-path password password;
}
```

When you specify the archive site, do not add a forward slash (/) to the end of the URL. The format for the destination filename is
< router-name > _juniper.conf[.gz]_*YYYYMMDD_HHMMSS*

Specifying the Number of Configurations Stored on the Flash Drive

By default, the JUNOS software saves the current configuration and three previous versions of the committed configuration on the flash drive. The currently operational JUNOS software configuration is stored in the file juniper.conf.gz, 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 router's flash drive in the directory /config.

In addition to saving the current configuration and the current operational version, you can also specify how many previous versions of the committed configurations you want stored on the flash drive in the directory /config. The remaining previous versions of committed configurations are stored in the directory /var/db/config on the hard disk. This is useful when you have very large configurations that might not fit on the flash drive.

To specify how many previous versions of the committed configurations you want stored on the flash drive, include the max-configurations-on-flash statement at the [edit system] hierarchy level:

```
[edit system]
max-configurations-on-flash number;
```

number can be from 0 to 49.

For more information about how the configuration is stored, see "How the Configuration Is Stored" on page 221.

Configuring RADIUS System Accounting

With RADIUS accounting enabled, Juniper Network routers, acting as RADIUS clients, can notify the RADIUS server about user activities such as software logins, configuration changes, and interactive commands. The framework for RADIUS accounting is described in RFC 2866.

To audit user events, include the following statements at the [edit system accounting] hierarchy level:

```
[edit system accounting]
events [ events ];
destination {
    radius {
        server {
            server-address {
                accounting-port port-number;
                secret password;
                source-address address;
                retry number;
                timeout seconds;
            }
        }
    }
}
```

This section includes the following topics:

- Specifying Events on page 488
- Configuring RADIUS Accounting on page 486
- Example: Configuring RADIUS Accounting on page 487

Specifying Events

To specify the events you want to audit, include the events statement at the [edit system accounting] hierarchy level:

```
[edit system accounting]
events [ events ];
```

events is one or more of the following:

- login—Audit logins
- change-log—Audit configuration changes
- interactive-commands—Audit interactive commands (any command-line input)

Configuring RADIUS Accounting

To configure RADIUS server accounting, include the server statement at the [edit system accounting destination radius] hierarchy level:

```
server {
    server-address {
        accounting-port port-number;
        secret password;
        source-address address;
        retry number;
        timeout seconds;
    }
}
```

server-address specifies the address of the RADIUS server. To configure multiple RADIUS servers, include multiple server statements.

NOTE: If no RADIUS servers are configured at the [edit system accounting destination tacplus] statement hierarchy level, the JUNOS software uses the RADIUS servers configured at the [edit system radius-server] hierarchy level.

accounting-port *port-number* specifies the RADIUS server accounting port number.

You must specify a secret (password) that the local router passes to the RADIUS client by including the secret statement. Secrets can contain spaces. The secret used by the local router must match that used by the server.

In the source-address statement, specify a source address for the RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 address configured on one of the router interfaces.

You can optionally specify the number of times that the router attempts to contact a RADIUS authentication server by including the retry statement. By default, the router retries 3 times. You can configure the router to retry from 1 through 10 times.

You can optionally specify the length of time that the local router waits to receive a response from a RADIUS server by including the timeout statement. By default, the router waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

Example: Configuring RADIUS Accounting

The following example shows three servers (5.5.5.5, 6.6.6.6, and 7.7.7.7) configured for RADIUS accounting.

```
system {
    accounting {
      events [ login change-log interactive-commands ];
      destination {
        radius {
          server {
            5.5.5.5 {
              accounting-port 3333;
              secret $9$dkafeqwrew;
              source-address 1.1.1.1;
              retry 3;
              timeout 3;
            }
            6.6.6.6 secret $9$fe3erqwrez;
            7.7.7.7 secret $9$f34929ftby;
          }
        }
      }
    }
}
```

Configuring TACACS+ System Accounting

You can use TACACS+ to track and log software logins, configuration changes, and interactive commands. To audit these events, include the following statements at the [edit system accounting] hierarchy level:

```
    [edit system accounting]
    events [ events ];
    destination {
        tacplus {
          server {
            server-address {
              port port-number;
              secret password;
              single-connection;
              timeout seconds;
            }
          }
        }
    }
```

This section includes the following topics:

- Specifying Events on page 488
- Configuring TACACS+ Accounting on page 488

Specifying Events

To specify the events you want to audit, include the events statement at the [edit system accounting] hierarchy level:

```
[edit system accounting]
events [ events ];
```

events is one or more of the following:

- login—Audit logins
- change-log—Audit configuration changes
- interactive-commands—Audit interactive commands (any command-line input)

Configuring TACACS+ Accounting

To configure TACACS+ server accounting, include the server statement at the [edit system accounting destination tacplus] hierarchy level:

```
[edit system accounting destination tacplus]
server {
    server-address {
      port port-number;
      secret password;
      single-connection;
      timeout seconds;
    }
}
```

server-address specifies the address of the TACACS+ server. To configure multiple TACACS+ servers, include multiple server statements.

NOTE: If no TACACS+ servers are configured at the [edit system accounting destination tacplus] statement hierarchy level, the JUNOS software uses the TACACS+ servers configured at the [edit system tacplus-server] hierarchy level.

port-number specifies the TACACS+ server port number.

You must specify a secret (password) that the local router passes to the TACACS+ client by including the secret statement. Secrets can contain spaces. The secret used by the local router must match that used by the server.

You can optionally specify the length of time that the local router waits to receive a response from a TACACS+ server by including the timeout statement. By default, the router waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds.

You can optionally maintain one open TCP connection to the server for multiple requests, rather than opening a connection for each connection attempt, by including the single-connection statement.

Enabling the SDX Software

You can enable JUNOS software to work with the Service Deployment System (SDX) software. The SDX software supports dynamic service activation engine (SAE) functionality on JUNOS routers. To do this, include the following statements at the [edit system services service-deployment] hierarchy level:

```
[edit system services service-deployment]
servers server-address {
    port port-number;
}
source-address source-address;
```

server-address is the IPv4 address of the SDX server.

By default, *port-number* is set to 3333 and is a TCP port number.

source-address is optional, and is the local IP version 4 (IPv4) address to be used as the source address for traffic to the SDX server.

For more information about SDX software, see the SDX documentation set.

Configuring the Path MTU Discovery

By default, path maximum transmission unit (MTU) discovery on outgoing TCP connections is disabled. To enable path MTU discovery, include the path-mtu-discovery statement at the [edit system internet-options] hierarchy level:

```
[edit system internet-options]
path-mtu-discovery;
```

Configuring Source Quench

By default, Internet Control Message Protocol (ICMP) source quench is disabled. You enable source quench when you want the JUNOS software to ignore ICMP source quench messages. To do this, include the source-quench statement at the [edit system internet-options] hierarchy level:

```
[edit system internet-options]
source-quench;
```

Configuring the Range of Port Addresses

By default, the upper range of a port address is 5000. You can increase the range from which the port number can be selected to decrease the probability that someone can determine your port number. To do so, include the source-port statement at the [edit system internet-options] hierarchy level:

```
[edit system internet-options]
source-port upper-limit <upper-limit>;
```

upper-limit *<upper-limit>*—Is the upper limit of a source port address and can be a value from 5000 through 65,355.

Configuring ARP Learning and Aging

The Address Resolution Protocol (ARP) is a protocol used by IPv4 to map IP network addresses to MAC addresses. This section describes how to set passive ARP learning and ARP aging options for routers. For more information about configuring ATM on Juniper Networks routing platforms, see the *JUNOS Network Interfaces Configuration Guide*.

For more information, see the following sections:

- Configuring Passive ARP Learning for Backup VRRP Routers on page 490
- Adjusting the ARP Aging Timer on page 491

Configuring Passive ARP Learning for Backup VRRP Routers

By default, the backup VRRP router drops ARP requests for the VRRP-IP to VRRP-MAC address translation. The backup router does not learn the ARP (IP-to-MAC address) mappings for the hosts sending the requests. When it detects a failure of the master router and becomes the new master, the backup router must learn all the entries that were present in the ARP cache of the master router. In environments with many directly attached hosts, such as metro Ethernet environments, the number of ARP entries to learn can be high. This can cause a significant transition delay, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup router to hold approximately the same contents as the ARP cache in the master router, thus preventing the problem of learning ARP entries in a burst. To enable passive ARP learning, include the passive-learning statement at the [edit system arp] hierarchy level:

```
[edit system arp]
passive-learning;
```

We recommend setting passive learning on both the backup and master VRRP routers. This prevents the need to intervene manually when the master router becomes the backup router. While a router is operating as the master, the passive learning configuration has no operational impact. The configuration takes effect only when the router is operating as a backup router.

Adjusting the ARP Aging Timer

By default, the ARP aging timer is set at 20 minutes. In environments with many directly attached hosts, such as metro Ethernet environments, increasing the amount of time between ARP updates by configuring the ARP aging timer can improve performance.

To configure the ARP aging timer, include the aging-timer statement at the [edit system arp] hierarchy level:

```
[edit system arp]
aging-timer minutes;
```

The aging timer range is from 20 through 240 minutes. The timer value you configure takes effect as ARP entries expire. Each refreshed ARP entry receives the new timer value. The new timer value does not apply to ARP entries that exist at the time you commit the configuration.

Configuring System Alarms to Show Automatically

You can configure J-series Services Routers to execute a show system alarms command whenever a user with the login class admin logs on to the router. To do so, include the login-alarms statement at the [edit system login class admin] hierarchy level:

```
[edit system login class admin]
    login-alarms;
```

Table 34 describes system alarms that may occur. These alarms are preset and cannot be modified.

Table 34: System Alarms

Alarm Type	Alarm summary	Remedy
Configuration	This alarm appears if you have not created a rescue configuration for the router. If you inadvertently commit a configuration that denies management access to the router, you must either connect a console to the router or invoke a rescue configuration. Using a rescue configuration is the recommended method. A rescue configuration is one that you know allows management access to the router.	Set the rescue configuration. For more information on setting the rescue configuration, see “Creating and Returning to a Rescue Configuration” on page 268.

Table 34: System Alarms

Alarm Type	Alarm summary	Remedy
License	This alarm appears if you have configured at least one software feature that requires a feature license, but no valid license for the feature is currently installed.	Install a valid license key.

For more information on system alarms for J-series Services Routers, see the *J-series Services Router Administration Guide*. For more information on the show system alarms command, see the *JUNOS System Basics and Services Command Reference.*

Chapter 26

Security Configuration Example

This chapter provides an example of a configuration that applies sound security policies so that the router can operate securely. This configuration is specific to IP version 4 (IPv4).

The following sections explain how to configure a router securely:

- Configuring System Information on page 494
- Configuring Interfaces on page 499
- Configuring SNMP on page 501
- Configuring Protocol-Independent Routing Properties on page 504
- Configuring Routing Protocols on page 505
- Configuring Firewalls on page 508

The final section in this example, "Example: Consolidated Security Configuration" on page 512, shows the complete configuration example.

NOTE: For advanced network security, a special version of JUNOS, called JUNOS FIPS is available. For information on how to configure a network of Juniper Networks routers in a FIPS environment, see the *JUNOS-FIPS Configuration Guide*

Configuring System Information

Configure the router name and domain name:

```
[edit]
system {
    host-name Secure-Router;
    domain-name company.com;
    default-address-selection;
}
```

This section includes the following topics:

- Configuring RADIUS on page 494
- Creating Login Classes on page 495
- Defining User Login Accounts on page 496
- Defining RADIUS Template Accounts on page 496
- Enabling Connection Services on page 497
- Configuring System Logging on page 497
- Configuring the Time Source on page 498

Configuring RADIUS

The JUNOS software supports two protocols for central authentication of users on multiple routers: Remote Authentication Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+). We recommend RADIUS because it is a multivendor IETF standard, and its features are more widely accepted than those of TACACS+ or other proprietary systems. In addition, we recommend using a one-time-password system for increased security, and all vendors of these systems support RADIUS.

In the JUNOS model for centralized RADIUS authentication, you create one or more template accounts on the router, and the users' access to the router is configured to use the template account. In this configuration, if the RADIUS server is not reachable, the fallback authentication mechanism is through the local account set up on the router.

```
[edit]
system {
    authentication-order [ radius password ];
    root-authentication {
      encrypted-password "$9$aH1j8gqQ1gjyjgjhgjgiiiii"; # SECRET-DATA
    }
    name-server {
      10.1.1.1;
      10.1.1.2;
    }
}
```

Enable RADIUS authentication and define the shared secret between the client and the server so each know that they are talking to the trusted peer. Define a timeout value for each server so if there is no response within the specified number of seconds, the router can try either the next server or the next authentication mechanism.

```
[edit]
system {
    radius-server {
        10.1.2.1 {
            secret "$9$aH1j8gqQ1sdjerrrhser"; # SECRET-DATA
            timeout 5;
        }
        10.1.2.2 {
            secret "$9$aH1j8gqQ1csdoiuardwefoiud"; # SECRET-DATA
            timeout 5;
        }
    }
}
```

Creating Login Classes

Create several user classes, each with specific privileges. In this example, you configure timeouts to disconnect the class members after a period of inactivity. Users' privilege levels, and therefore the classes of which they are members, should be dependent on their responsibilities within the organization, and the permissions shown here are only examples.

The first class of users (called “observation”) can only view statistics and configuration. They are not allowed to modify any configuration. The second class of users (called “operation”) can view and modify the configuration. The third class of users (called “engineering”) has unlimited access and control.

```
[edit]
system {
    login {
        class observation {
            idle-timeout 5;
            permissions [ view ];
        }
        class operation {
            idle-timeout 5;
            permissions [ admin clear configure interface interface-control network
            reset routing routing-control snmp snmp-control trace-control
            firewall-control rollback ];
        }
        class engineering {
            idle-timeout 5;
            permissions all;
        }
    }
}
```

Defining User Login Accounts

Define the local superuser account. If RADIUS fails or become unreachable, we revert to the local accounts on the router.

```
[edit]
system {
    login {
        user admin {
            uid 1000;
            class engineering;
            authentication {
                encrypted-password "<PASSWORD>"; # SECRET-DATA
            }
        }
    }
}
```

Defining RADIUS Template Accounts

Define RADIUS template accounts for different users or groups of users:

```
[edit]
system {
    login {
        user observation {
            uid 1001;
            class observation;
        }
        user operation {
            uid 1002;
             class operation;
        }
        user engineering {
            uid 1003;
            class engineering;
        }
    }
}
```

Enabling Connection Services

Enable connection services on the router. SSH provides secure encrypted communications over an insecure network and is therefore useful for inband router management. Like all other types of network-based access, however, SSH access to the router is disabled by default in the JUNOS software. The following configuration enables SSH access and sets optional parameters that can be used to control the number of concurrent SSH sessions and the maximum number of SSH sessions that can be established in one minute. The rate-limit option can be useful in protecting against SYN flood denial-of-service (DoS) attacks on the SSH port.

```
  [edit]
  system {
      login {
      services {
         ssh connection-limit 10 rate-limit 4;
      }
  }
}
```

Configuring System Logging

A file that records when authentication and authorization is granted and rejected, as well as all user commands, provides an excellent way to track all management activity on the router. Checking these files for failed authentication events can help identify attempts to hack into the router. These files can also provide logs of all the command executed on the router and who has performed them. You can review logs of the commands executed on the router and correlate any event in the network with changes made at a particular time. These files are stored locally on the router. Place the firewall logs in a separate system log file.

```
[edit]
system {
    syslog {
       file messages {
          any notice;
          authorization info;
          daemon any;
          kernel any;
          archive size 10m files 5 no-world-readable;
       }
       file authorization-commands {
          authorization any;
          interactive-commands any;
       }
       file firewall-logs {
          firewall any;
       }
    }
}
```

Configuring the Time Source

Debugging and troubleshooting are much easier when the timestamps in the log files of all routers are synchronized, because events that span the network can be correlated with synchronous entries in multiple logs. We strongly recommend the using the Network Time Protocol (NTP) to synchronize the system clocks of routers and other network equipment.

By default, NTP operates in an entirely unauthenticated manner. If a malicious attempt to influence the accuracy of a router's clock succeeds, it could have negative effects on system logging, make troubleshooting and intrusion detection more difficult, and impede other management functions.

The following configuration synchronizes all the routes in the network to a single time source. We recommend using authentication to make sure that the NTP peer is trusted. The boot-server statement identifies the server from which the initial time of day and date is obtained when the router boots. The server statement identifies the NTP server used for periodic time synchronization. The authentication-key statement specifies that an HMAC-Message Digest 5 (MD5) scheme is used to hash the key value for authentication, which will prevent the router from synchronizing with a attacker's host posing as the time server.

```
[edit]
system {
    ntp {
        authentication-key 2 type md5 value "$9$aH1j8gqQ1gjyjgjhgjgiiiii"; #
SECRET-DATA
        boot-server 10.1.4.1;
        server 10.1.4.2;
    }
}
```

Configuring Interfaces

Configure the interfaces on your router. This example shows configurations for Asynchronous Transfer Mode (ATM), SONET, loopback, and out-of-band management interfaces. For more information about configuring interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Configure an ATM interface:

```
[edit]
interfaces {
    at-4/0/0 {
      description core-router;
      atm-options {
        vpi 0 maximum-vcs 1024;
        ilmi;
      }
      unit 131 {
        description to-other-core-router;
        encapsulation atm-snap;
        point-to-point;
        vci 0.131;
        family inet {
          address 12.1.1.1/30;
        }
        family iso;
      }
    }
}
```

The fxp0 interface can be used for out-of-band management. However, because most service providers use inband communication for management (because of lower operating costs), you can disable this interface to make the router more secure.

```
[edit]
interfaces {
    fxp0 {
      disable;
    }
}
```

Configure the loopback interface. To protect the Routing Engine, apply a firewall filter to the router's loopback interface. This filter, which you define at the [edit firewall] hierarchy level, checks all traffic destined for the Routing Engine that enters the router from the customer interfaces. Adding or modifying filters for every interface on the router is not necessary.

```
[edit]
interfaces {
    lo0 {
        unit 0 {
            family inet {
                filter {
                    input protect-routing-engine;
                }
                address 10.10.5.1/32;
            }
            family iso {
                address 48.0005.80dd.f900.0000.0001.0001.0000.0000.011.00;
            }
        }
    }
}
```

Configure a SONET interface:

```
[edit]
interfaces {
    so-2/0/0 {
        description To-other-router;
        clocking external;
        sonet-options {
            fcs 32;
            payload-scrambler;
        }
        unit 0 {
            family inet {
                address 10.1.5.1/30;
            }
            family iso;
        }
    }
}
```

Configuring SNMP

Configure Simple Network Management Protocol version 3 (SNMPv3):

```
[edit snmp]
engine-id {
      use-fxp0-mac-address;
}
view jnxAlarms; {
      oid 1.3.6.1.4.1.2636.3.4 include;
)
view interfaces {
      oid 1.3.6.1.2.1.2 include;
)
view ping-mib; {
      oid 1.3.6.1.2.1.80 include;
)

[edit snmp v3]
notify n1 {
      tag router1;                # Identifies a set of target addresses
      type trap;                  # Defines type of notification
}
notify n2 {
      tag host1;
      type trap;
}
notify-filter nf1 {
      oid .1 include;     # Defines which traps (or which objects for which traps)
}                          # that will be sent. In this case, include all traps.
notify-filter nf2 {
      oid 1.3.6.1.4.1 include;     # Send enterprise-specific traps only
}
notify-filter nf3 {
       oid 1.3.6.1.2.1.1.5 include;  # Send BGP traps only
}
snmp-community index1 {
      community-name "$9$JOZi.QF/AtOz3"; # SECRET-DATA
      security-name john; # Matches the security name at the target-parameters
      tag host1;                 # Finds the addresses that are allow to be used
}                                # with this community string
target-address ta1 {            # Associates the target address with the group
san-francisco
      address 10.1.1.1
      address-mask 255.255.255.0;# Defines the range of addresses
      port 162;
      tag-list router1;
      target-parameters tp1;        # Apply configured target parameters
}
target-address ta2 {
      address 10.1.1.2
      address-mask 255.255.255.0;
      port 162;
      tag-list host1;
      target-parameters tp2;
}
```

```
target-address ta3 {
    address 10.1.1.3;
    address-mask 255.255.255.0;
    port 162;
    tag-list [router1 host1];
    target-parameters tp3;
}
target-parameters tp1 {         # Define the target parameters
    notify-filter nf1;          # Specify which notify filter to apply
    parameters {
      message-processing-model v1;
      security-model v1‘;
      security-level none;
      security-name john; # Matches the security name configured at the [edit
    }                   # snmp v3  snmp-community community-index] hierarchy
}                       #level
target-parameters tp2 {
    notify-filter nf2;
    parameters {
      message-processing-model v1;
      security-model v1‘;
      security-level none;
      security-name john;
    }
}
target-parameters tp3 {
    notify-filter nf3;
    parameters {
      message-processing-model v1;
      security-model v1‘;
      security-level none;
      security-name john;
    }
}
usm {
    local-engine {            #Define authentication and encryption for SNMP3
      user user1 {            #users.
        authentication-md5 {
          authentication-password authentication-password;
        }
        privacy-des {
          privacy-password password;
        }
      }
      user user2 {
        authentication-sha {
          authentication-password authentication-password;
        }
        privacy-none;
      }
      user user3 {
        authentication-none;
        privacy-none;
      }
```

```
         user user4 {
            authentication-md5 {
               authentication-password authentication-password;
            }
            privacy-3des {
               privacy-password password;
            }
         }
         user user5 {
            authentication-sha {
               authentication-password authentication-password;
            }
            privacy-aes128 {
               privacy-password password;
            }
         }
      }
}
vacm {
       access {
       group san-francisco {             # Defines the access privileges for the group
            default-context-prefix { #san-francisco
               security-model v1 {
                  security-level none {
                     notify-view ping-mib;
                     read-view interfaces;
                     write-view jnxAlarms;
                  }
               }
            }
         }
      }
      security-to-group {
        security-model v1 {
           security-name john {             # Assigns john to the security group san-francisco
               group san-francisco;
           }
           security-name bob {
               group new-york;
           }
           security-name elizabeth {
               group chicago;
           }
         }
      }
}
```

For more information about configuring SNMP, see the *JUNOS Network Management Configuration Guide*.

Configuring Protocol-Independent Routing Properties

Configure a router ID and autonomous system (AS) number for Border Gateway Protocol (BGP):

```
[edit]
routing-options {
    router-id 10.1.7.1;
    autonomous-system 222;
}
```

Configure martian addresses, which are reserved host or network addresses about which all routing information should be ignored. By default, the JUNOS software blocks the following martian addresses: 0.0.0.0/8, 127.0.0.0/8, 128.0.0.0/16, 191.255.0.0/16, 192.0.0.0/24, 223.255.55.0/24, and 240.0.0.0/4. It is also a good idea to block private address space (addresses defined in RFC 1918). You can add these addresses and other martian addresses to the default martian addresses.

```
[edit]
routing-options {
    martians {
      1.0.0.0/8 exact;
      10.0.0.0/8 exact;
      19.255.0.0/16 exact;
      59.0.0.0/8 exact;
      129.156.0.0/16 exact;
      172.16.0.0/12 exact;
      192.0.2.0/24 exact;
      192.5.0.0/24 exact;
      192.9.200.0/24 exact;
      192.9.99.0/24 exact;
      192.168.0.0/16 exact;
      224.0.0.0/3 exact;
    }
}
```

For more information about configuring protocol-independent routing properties, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring Routing Protocols

The main task of a router is to use its routing and forwarding tables to forward user traffic to its intended destination. Attackers can send forged routing protocol packets to a router with the intent of changing or corrupting the contents of its routing table or other databases, which in turn can degrade the functionality of the router and the network. To prevent such attacks, routers must ensure that they form routing protocol relationships (peering or neighboring relationships) to trusted peers. One way of doing this is by authenticating routing protocol messages. We strongly recommend using authentication when configuring routing protocols. The JUNOS software supports HMAC-MD5 authentication for BGP, Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and Resource Reservation Protocol (RSVP). HMAC-MD5 uses a secret key that is combined with the data being transmitted to compute a hash. The computed hash is transmitted along with the data. The receiver uses the matching key to recompute and validate the message hash. If an attacker has forged or modified the message, the hash will not match and the data will be discarded.

In this example, we configure BGP and, as the interior gateway protocol (IGP), IS-IS. If you use OSPF, configure it similarly to the IS-IS configuration shown.

This section includes the following topics:

- Configuring BGP on page 505
- Configuring IS-IS on page 507

For more information about configuring BGP and IS-IS, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring BGP

The following example shows the configuration of a single authentication key for the BGP peer group internal peers. You can also configure BGP authentication at the neighbor or routing instance levels, or for all BGP sessions. As with any security configuration, there is a tradeoff between the degree of granularity (and to some extent the degree of security) and the amount of management necessary to maintain the system. This example also configures a number of tracing options for routing protocol events and errors, which can be good indicators of attacks against routing protocols. These events include protocol authentication failures, which might point to an attacker that is sending spoofed or otherwise malformed routing packets to the router in an attempt to elicit a particular behavior.

```
[edit]
protocols {
    bgp {
        group ibgp {
            type internal;
            traceoptions {
                file bgp-trace size 1m files 10;
                flag state;
                flag general;
            }
            local-address 10.10.5.1;
            log-updown;
            neighbor 10.2.1.1;
            authentication-key "$9$aH1j8gqQ1gjyjgjhgjgiiiii";
        }
        group ebgp {
            type external;
            traceoptions {
                file ebgp-trace size 10m files 10;
                flag state;
                flag general;
            }
            local-address 10.10.5.1;
            log-updown;
            peer-as 2;
            neighbor 10.2.1.2;
            authentication-key "$9$aH1j8gqQ1gjyjgjhgjgiiiii";
        }
    }
}
```

Configuring IS-IS

Although all JUNOS IGPs support authentication, some are inherently more secure than others. Most service providers use OSPF or IS-IS to allow fast internal convergence and scalability and to use traffic engineering capabilities with Multiprotocol Label Switching (MPLS). Because IS-IS does not operate at the network layer, it is more difficult to spoof than OSPF, which is encapsulated in IP and is therefore subject to remote spoofing and DoS attacks. This example also configures a number of tracing options for routing protocol events and errors, which can be good indicators of attacks against routing protocols. These events include protocol authentication failures, which might point to an attacker that is sending spoofed or otherwise malformed routing packets to the router in an attempt to elicit a particular behavior.

```
[edit]
protocols {
    isis {
      authentication-key "$9$aH1j8gqQ1gjyjgjhgjgiiiii"; # SECRET-DATA
      authentication-type md5;
      traceoptions {
        file isis-trace size 10m files 10;
        flag normal;
        flag error;
      }
      interface at-0/0/0.131 {
        lsp-interval 50;
        level 2 disable;
        level 1 {
          metric 3;
          hello-interval 5;
          hold-time 60;
        }
      }
      interface lo0.0 {
      passive;
    }
}
```

Configuring Firewalls

To configure firewall policies, configure the trusted source addresses with which each protocol or service wants to communicate. Once you define the prefix list, you apply it in the filter definition at the [edit firewall] hierarchy level.

For more information about configuring firewalls, see the *JUNOS Policy Framework Configuration Guide*.

```
[edit]
policy-options {
    prefix-list ssh-addresses {
      1.1.9.0/24;
    }
    prefix-list bgp-addresses {
      10.2.1.0/24;
    }
    prefix-list ntp-addresses {
      10.1.4.0/24;
    }
    prefix-list snmp-addresses {
      10.1.6.0/24;
    }
    prefix-list dns-address {
      10.1.1.0/24;
    }
    prefix-list radius-address {
      10.1.2.0/24;
    }
}
```

The following firewall filter protects the Routing Engine. To protect the Routing Engine, it is important to constrain the traffic load from each of the allowed services. Rate-limiting control traffic helps protect the Routing Engine from attack packets that are forged such that they appear to be legitimate traffic and are then sent at such a high rate as to cause a DoS attack.

Routing and control traffic are essential to proper functioning of the router, and rapid convergence of routing protocols is crucial for stabilizing the network during times of network instability. While it might seem desirable to limit the amount of routing protocol traffic to protect against various types of attacks, it is very difficult to determine a fixed maximum rate for protocol traffic, because it depends upon the number of peers and adjacencies, which varies over time. Therefore, it is best not to rate-limit routing protocol traffic.

By contrast, because management traffic is less essential and more deterministic than routing protocol traffic, it can be policed to a fixed rate, to prevent it from consuming resources necessary for less flexible traffic. We recommend allocating a fixed amount of bandwidth to each type of management traffic so that an attacker cannot consume all the router's CPU if an attack is launched using any single service.

```
[edit]
firewall {
    filter protect-routing-engine {
      policer ssh-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer small-bandwidth-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer snmp-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer ntp-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer dns-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer radius-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer tcp-policer {
        if-exceeding {
          bandwidth-limit 500k;
          burst-size-limit 15k;
        }
        then discard;
      }

/* The following terms accept traffic only from the trusted sources. The trusted
traffic is rate-limited with the exception of the routing protocols. */
```

```
/* The following term protects against ICMP flooding attacks against the Routing
Engine. */
      term icmp {
        from {
          protocol icmp;
          icmp-type [ echo-request echo-reply unreachable time-exceeded ];
        }
        then {
          policer small-bandwidth-policer;
          accept;
        }
      }
      term tcp-connection {
        from {
          source-prefix-list {
            ssh-addresses;
            bgp-addresses;
          }
          protocol tcp;
          tcp-flags "(syn & !ack) | fin | rst";
        }
        then {
          policer tcp-policer;
          accept;
        }
      }
/* The following term protects ssh traffic destined for the Routing Engine. */
      term ssh {
        from {
          source-prefix-list {
            ssh-addresses;
          }
          protocol tcp;
          port [ ssh telnet ];
        }
        policer ssh-policer;
        then accept;
      }
/* The following term protects BGP traffic destined for the Routing Engine. */
      term bgp {
        from {
          source-prefix-list {
            bgp-sessions-addresses;
          }
          protocol tcp;
          port bgp;
        }
        then accept;
      }
      term snmp {
        from {
          source-prefix-list {
            snmp-addresses;
          }
          protocol udp;
          port snmp;
        }
        then {
            policer snmp-policer;
            accept;
        }
    }
    term ntp {
        from {
            source-prefix-list {
                ntp-addresses;
            }
            protocol udp;
            port ntp;
        }
        then {
            policer ntp-policer;
            accept;
        }
    }
    term dns {
        from {
            source-address {
                dns-addresses;
            }
            protocol udp;
            port domain;
        }
        then {
            policer dns-policer;
            accept;
        }
    }
    term radius {
        from {
            source-address {
                radius-addresses;
            }
            protocol udp;
            port radius;
        }
        then {
            policer radius-policer;
            accept;
        }
    }
    term trace-route {
        from {
            protocol udp;
            destination-port 33434-33523;
        }
        then {
            policer small-bandwidth-policer;
            accept;
    }
```

/* All other traffic that is not trusted is silently dropped. We recommend logging the denied traffic for analysis purposes. */

```
            term everything-else {
                then {
                    syslog;
                    log;
                    discard;
                }
            }
        }
}
```

Example: Consolidated Security Configuration

Basic System Information

```
system {
        host-name Secure-Router;
        domain-name company.com;
        default-address-selection;
```

RADIUS

```
        authentication-order [ radius password ];
        root-authentication {
            encrypted-password "$9$aH1j8gqQ1gjyjgjhgjgiiiii"; # SECRET-DATA
        }
        name-server {
            10.1.1.1;
            10.1.1.2;
        }
        radius-server {
            10.1.2.1 {
                secret "$9$aH1j8gqQ1sdjerrrhser"; # SECRET-DATA
                timeout 5;
            }
            10.1.2.2 {
                secret "$9$aH1j8gqQ1csdoiuardwefoiud"; # SECRET-DATA
                timeout 5;
            }
        }
```

Login Classes

```
        login {
            class observation {
                idle-timeout 5;
                permissions [ view ];
            }
            class operation {
            idle-timeout 5;
                permissions [ admin clear configure interface interface-control network
                    reset routing routing-control snmp snmp-control trace-control
                    firewall-control rollback ];
            }
            class engineering {
                idle-timeout 5;
                permissions all;
            }
```

User Login Accounts

```
        user admin {
            uid 1000;
            class engineering;
            authentication {
                encrypted-password "<PASSWORD>"; # SECRET-DATA
            }
        }
```

RADIUS Template Accounts

```
        user observation {
            uid 1001;
            class observation;
        }
        user operation {
            uid 1002;
             class operation;
        }
        user engineering {
            uid 1003;
            class engineering;
        }
    }
```

Connection Services

```
    services {
        ssh connection-limit 10 rate-limit 4;
    }
```

System Logging

```
    syslog {
        file messages {
            any notice;
            authorization info;
            daemon any;
            kernel any;
            archive size 10m files 5 no-world-readable;
        }
        file authorization-commands {
            authorization any;
            interactive-commands any;
        }
        file firewall-logs {
            firewall any;
        }
    }
```

Time Source

```
    ntp {
        authentication-key 2 type md5 value "$9$aH1j8gqQ1gjyjgjhgjgiiiii";
            # SECRET-DATA
        boot-server 10.1.4.1;
        server 10.1.4.2;
    }
}
```

Interfaces

```
interfaces {
    at-4/0/0 {
        description core router;
        atm-options {
            vpi 0 maximum-vcs 1024;
            ilmi;
        }
        unit 131 {
            description to-other-core-router;
            encapsulation atm-snap;
            point-to-point;
            vci 0.131;
            family inet {
                address 12.1.1.1/30;
            }
            family iso;
        }
    }
    fxp0 {
        disable;
    }
    lo0 {
        unit 0 {
            family inet {
                filter {
                    input protect-routing-engine;
                }
                address 10.10.5.1/32;
            }
            family iso {
                address 48.0005.80dd.f900.0000.0001.0001.0000.0000.011.00;
            }
        }
    }
    so-2/0/0 {
        description To-other-router;
        clocking external;
        sonet-options {
            fcs 32;
            payload-scrambler;
        }
        unit 0 {
            family inet {
                address 10.1.5.1/30;
            }
            family iso;
        }
    }
}
```

SNMP

```
[edit snmp]
engine-id {
    use-fxp0-mac-address;
}
view jnxAlarms; {
    oid .1.3.6.1.4.1.2636.3.4 include;
)
view interfaces {
    oid .1.3.6.1.2.1.2 include;
)
view ping-mib; {
    oid .1.3.6.1.2.1.80 include;
)

[edit snmp v3]
notify n1 {
    tag router1;                 # Identifies a set of target addresses
    type trap;                   # Defines type of notification
}
notify n2 {
    tag host1;
    type trap;
}
notify-filter nf1 {
  oid 1 include;                # Defines which traps (or which objects for which
traps)
}                               # that will be sent. In this case, include all traps.
notify-filter nf2 {
  oid 1.3.6.1.4.1 include;      # Send enterprise-specific traps only
}
notify-filter nf3 {
  oid 1.3.6.1.2.1.1.5 include;  # Send BGP traps only
}
snmp-community index1 {
    community-name "$9$JOZi.QF/AtOz3"; # SECRET-DATA
    security-name john;           # Matches the security name at the
    tag host1;                  # target-parameters  Finds the addresses that are
}                               # allow to be used with this community string
target-address ta1 {           # Associates the target address with the group
san-francisco
    address 10.1.1.1;
    address-mask 255.255.255.0;# Defines the range of addresses
    port 162;
    tag-list router1;
    target-parameters tp1;        # Apply configured target parameters
}
target-address ta2 {
    address 10.1.1.2;
    address-mask 255.255.255.0;
    port 162;
    tag-list host1;
    target-parameters tp2;
}
```

```
target-address ta3 {
    address 10.1.1.3;
    address-mask 255.255.255.0;
    port 162;
    tag-list [router1 host1];
    target-parameters tp3;
}
target-parameters tp1 {                # Define the target parameters
    notify-filter nf1;                 # Specify which notify filter to apply
    parameters {
      message-processing-model v1;
      security-model v1‘;
      security-level none;
      security-name john; # Matches the security name configured at the [edit
    }                    #snmp v3  snmp-community community-index] hierarchy level
}
target-parameters tp2 {
    notify-filter nf2;
    parameters {
      message-processing-model v1;
      security-model v1‘;
      security-level none;
      security-name john;
    }
}
target-parameters tp3 {
    notify-filter nf3;
    parameters {
      message-processing-model v1;
      security-model v1‘;
      security-level none;
      security-name john;
    }
}
usm {
    local-engine {       #Define authentication and encryption for SNMP3 users.
      user user1 {
        authentication-md5 {
          authentication-password authentication-password;
        }
        privacy-des {
          privacy-password password;
          }
        }
      user user2 {
        authentication-sha {
          authentication-password authentication-password;
        }
        privacy-none;
        }
      user user3 {
        authentication-none;
        privacy-none;
      }
      user user4 {
        authentication-md5 {
          authentication-password authentication-password;
        }
        privacy-3des {
          privacy-password password;
        }
      }
      user user5 {
        authentication-sha {
          authentication-password authentication-password;
        }
        privacy-aes128 {
          privacy-password password;
        }
      }
    }
}
vacm {
     access {
     group san-francisco {      # Defines the access privileges for the group
        default-context-prefix { #san-francisco
          security-model v1 {
            security-level none {
              notify-view ping-mib;
              read-view interfaces;
              write-view jnxAlarms;
            }
          }
        }
      }
    }
    security-to-group {
     security-model v1 {
        security-name john {      # Assigns john to the security group san-francisco
          group san-francisco;
        }
        security-name bob {
          group new-york;
        }
        security-name elizabeth {
          group chicago;
        }
     }
    }
}
```

Protocol-Independent Routing Properties

```
routing-options {
    router-id 10.1.7.1;
    autonomous-system 222;
    martians {
      1.0.0.0/8 exact;
      10.0.0.0/8 exact;
      19.255.0.0/16 exact;
      59.0.0.0/8 exact;
      129.156.0.0/16 exact;
      172.16.0.0/12 exact;
      192.0.2.0/24 exact;
      192.5.0.0/24 exact;
      192.9.200.0/24 exact;
      192.9.99.0/24 exact;
      192.168.0.0/16 exact;
      224.0.0.0/3 exact;
    }
}
```

Routing Protocols

```
protocols {
```

BGP

```
    bgp {
      group ibgp {
        type internal;
        traceoptions {
          file bgp-trace size 1m files 10;
          flag state;
          flag general;
        }
        local-address 10.10.5.1;
        log-updown;
        neighbor 10.2.1.1;
        authentication-key "$9$aH1j8gqQ1gjyjgjhgjgiiiii";
      }
      group ebgp {
        type external;
        traceoptions {
          file ebgp-trace size 10m files 10;
          flag state;
          flag general;
        }
        local-address 10.10.5.1;
        log-updown;
        peer-as 2;
        neighbor 10.2.1.2;
        authentication-key "$9$aH1j8gqQ1gjyjgjhgjgiiiii";
      }
    }
```

IS-IS

```
    isis {
      authentication-key "$9$aH1j8gqQ1gjyjgjhgjgiiiii"; # SECRET-DATA
      authentication-type md5;
      traceoptions {
        file isis-trace size 10m files 10;
        flag normal;
        flag error;
      }
      interface at-0/0/0.131 {
        lsp-interval 50;
        level 2 disable;
        level 1 {
          metric 3;
          hello-interval 5;
          hold-time 60;
        }
      }
      interface lo0.0 {
      passive;
    }
}
```

Firewall Policies

```
policy-options {
    prefix-list ssh-addresses {
      1.1.9.0/24
    }
    prefix-list bgp-addresses {
      10.2.1.0/24;
    }
    prefix-list ntp-addresses {
      10.1.4.0/24
    }
    prefix-list snmp-addresses {
      10.1.6.0/24;
    }
    prefix-list dns-address {
      10.1.1.0/24;
    }
    prefix-list radius-address {
      10.1.2.0/24;
    }
}
```

Firewall Filters

```
firewall {
    filter protect-routing-engine {
      policer ssh-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer small-bandwidth-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer snmp-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer ntp-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer dns-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer radius-policer {
        if-exceeding {
          bandwidth-limit 1m;
          burst-size-limit 15k;
        }
        then discard;
      }
      policer tcp-policer {
        if-exceeding {
          bandwidth-limit 500k;
          burst-size-limit 15k;
        }
        then discard;
      }
      term icmp {
        from {
          protocol icmp;
          icmp-type [ echo-request echo-reply unreachable time-exceeded ];
        }
    then {
        policer small-bandwidth-policer;
        accept;
    }
}
term tcp-connection {
    from {
        source-prefix-list {
            ssh-addresses;
            bgp-addresses;
        }
        protocol tcp;
        tcp-flags "(syn & !ack) | fin | rst";
    }
    then {
        policer tcp-policer;
        accept;
    }
}
term ssh {
    from {
        source-prefix-list {
            ssh-addresses;
        }
        protocol tcp;
        port [ ssh telnet ];
    }
    policer ssh-policer;
    then accept;
}
term bgp {
    from {
        source-prefix-list {
            bgp-sessions-addresses;
        }
        protocol tcp;
        port bgp;
    }
    then accept;
}
term snmp {
    from {
        source-prefix-list {
            snmp-addresses;
        }
        protocol udp;
        port snmp;
    }
    then {
        policer snmp-policer;
        accept;
    }
}
```

```
      term ntp {
        from {
          source-prefix-list {
            ntp-addresses;
          }
          protocol udp;
          port ntp;
        }
        then {
          policer ntp-policer;
          accept;
        }
      }
      term dns {
        from {
          source-address {
            dns-addresses;
          }
          protocol udp;
          port domain;
        }
        then {
          policer dns-policer;
          accept;
        }
      }
      term radius {
        from {
          source-address {
            radius-addresses;
          }
          protocol udp;
          port radius;
        }
        then {
          policer radius-policer;
          accept;
        }
      }
      term trace-route {
        from {
          protocol udp;
          destination-port 33434-33523;
        }
        then {
          policer small-bandwidth-policer;
          accept;
      }
      term everything-else {
        then {
          syslog;
          log;
          discard;
        }
      }
    }
}
```

Chapter 27

Summary of System Management Configuration Statements

The following sections explain each of the system management configuration statements. The statements are organized alphabetically.

accounting

Syntax

```
accounting {
    events [ login change-log interactive-commands ];
    destination {
        radius {
            server {
                server-address {
                    accounting-port port-number;
                    secret password;
                    source-address address;
                    retry number;
                    timeout seconds;
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure audit of TACACS+ or RADIUS authentication events, configuration changes, and interactive commands.

The remaining statements are explained separately.

accounting-port

Syntax accounting-port *port-number*;

Hierarchy Level [edit system accounting destination radius server *server-address*]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure the accounting port number on which to contact the RADIUS server.

Options *number*—Port number on which to contact the RADIUS server.
Default: 1646

Usage Guidelines See "Configuring RADIUS System Accounting" on page 485.

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

Usage Guidelines See Configuring RADIUS System Accounting on page 485 and "Configuring TACACS+ System Accounting" on page 487.

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

allow-commands

Syntax allow-commands "*regular-expression*";

Hierarchy Level [edit system login class *class-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the operational mode commands that members of a login class can use.

Default If you omit this statement and the deny-commands statement, users can issue only those commands for which they have access privileges through the permissions statement.

Options *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Usage Guidelines See “Specifying Operational Mode Commands” on page 405.

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

See Also deny-commands on page 543, user on page 608

allow-configuration

Syntax allow-configuration "*regular-expression*";

Hierarchy Level [edit system login class *class-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the configuration mode commands that members of a login class can use.

Default If you omit this statement and the deny-configuration statement, users can issue only those commands for which they have access privileges through the permissions statement.

Options *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Usage Guidelines See “Specifying Operational Mode Commands” on page 405.

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

See Also deny-commands on page 543, user on page 608

allow-transients

Syntax allow-transients;

Hierarchy Level [edit system scripts commit]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, enable transient configuration changes to be committed.

Default Transient changes are disabled by default. If you do not include the allow-transients statement, and an enabled script generates transient changes, the CLI generates an error message and the commit operation does not succeed.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

announcement

Syntax announcement *text*;

Hierarchy Level [edit system login]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a system login announcement. This announcement appears after a user logs in.

Options *text*—Text of the announcement. If the text contains any spaces, enclose it in quotation marks.

Usage Guidelines See “Configuring a System Login Announcement” on page 480.

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

See Also message on page 566

archival

Syntax

```
archival {
    configuration {
        transfer-interval interval;
        transfer-on-commit;
        archive-sites {
            ftp://<username>:<password>@<host>:<port>/<url-path>;
        }
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4

Description Configure copying of the currently active configuration to an archive site.The remaining statements are described separately.

Usage Guidelines See “Configuring a Router to Transfer its Configuration to an Archive Site” on page 483.

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

archive

Syntax

```
archive {
    files number;
    size size;
    (world-readable | no-world-readable);
}
```

Hierarchy Level [edit system syslog],
[edit system syslog file *filename*]

Release Information Statement introduced before JUNOS Release 7.4

Description Configure how to archive system log files. The remaining statements are described separately.

Usage Guidelines See “Configuring Log File Archiving” on page 438.

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

archive-sites

Syntax

```
archive-sites {
    ftp://username@host:<port>url-path password password;
}
```

Hierarchy Level [edit system archival configuration]

Release Information Statement introduced before JUNOS Release 7.4

Description Specifies where to transfer the current configuration files. If you specify more than one archive site, the router attempts to transfer to the first archive site in the list, moving to the next only if the transfer fails. The format for the destination file name is < *router-name* > _juniper.conf[.gz]_*YYYYMMDD_HHMMSS*.

Usage Guidelines See "Configuring Archive Sites" on page 484.

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

See Also transfer-on-commit on page 607, transfer-on-commit on page 607, and configuration on page 539.

arp

Syntax

```
arp {
    passive-learning;
    aging-timer minutes;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4

Description Specifies ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

Options

- passive-learning—Configures backup VRRP routers to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the master router fails, the backup router must learn all entries present in the ARP cache of the master router. Configuring passive learning reduces transition delay when the backup router is activated.

- aging-timer—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, metro Ethernet environments), increasing the time between updates can improve system performance.
  Default: 20 minutes
  Range: 20 to 240 minutes

Usage Guidelines *JUNOS Network Interfaces Configuration Guide*.

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

See Also “Configuring ARP Learning and Aging” on page 490.

authentication

Syntax

```
authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-dsa "public-key";
    ssh-rsa "public-key";
}
```

Hierarchy Level [edit system login user *username*]

Release Information Statement introduced before JUNOS Release 7.4

Description Authentication methods that a user can use to log in to the router. You can assign multiple authentication methods to a single user.

Options encrypted-password "*password*"—Use Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.

plain-text-password—Use a plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it. For information about how to create plain-text passwords, see “Specifying Plain-Text Passwords” on page 17.

ssh-dsa "*public-key*"—Secure shell (SSH version 2) authentication. Specify the SSH public key. You can specify one or more public keys for each user.

ssh-rsa "*public-key*"—Secure shell (SSH version 1) authentication. Specify the SSH public key. You can specify one or more public keys for each user.

Usage Guidelines See “Configuring User Accounts” on page 413.

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

See Also root-authentication on page 582

authentication-key

Syntax authentication-key *key-number* type *type* value *password*;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure Network Time Protocol (NTP) authentication keys so that the router can send authenticated packets. If you configure the router to operate in authenticated mode, you must configure a key.

Both the keys and the authentication schemes (DES or MD5) must be identical between a set of peers sharing the same key number.

Options *key-number*—Positive integer that identifies the key.

type *type*—Authentication type. It can be either md5 or des.

value *password*—The key itself, which can be from 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.

Usage Guidelines See "Configuring NTP Authentication Keys" on page 424.

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

See Also broadcast on page 535, peer on page 572, server on page 587, trusted-key on page 607

authentication-order

Syntax authentication-order [*authentication-methods*];

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.

Default If you do not include the authentication-order statement, users are verified based on their configured passwords.

Options *authentication-methods*—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:

- password—Verify the user using the password configured for the user with the authentication statement at the [edit system login user] hierarchy level.
- radius—Verify the user using RADIUS authentication services.
- tacplus—Verify the user using TACACS+ authentication services.

Usage Guidelines See “Configuring the Authentication Order” on page 395.

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

auxiliary

Syntax

```
auxiliary {
    type terminal-type;
}
```

Hierarchy Level [edit system ports]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the characteristics of the auxiliary port, which is on the router’s craft interface.

Default The auxiliary port is disabled.

Options type *terminal-type*—Type of terminal that is connected to the port.
Values: ansi, vt100, small-xterm, xterm
Default: The terminal type is unknown, and the user is prompted for the terminal type.

Usage Guidelines See “Configuring Console and Auxiliary Port Properties” on page 456.

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

autoinstallation

Syntax

```
autoinstallation {
    interfaces {
        interface-name {
            bootp;
            rarp;
            slarp;
        }
    }
    configuration-servers {
        url;
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only, enables you to download a configuration file automatically from an FTP, Hypertext Transfer Protocol (HTTP), or Trivial FTP (TFTP) server. When you power on a J-series Services Router configured for autoinstallation, it requests an IP address from a Dynamic Host Configuration Protocol (DHCP) server. Once the router has an address, it sends a request to a configuration server and downloads and installs a configuration.

The remaining statements are explained separately in this chapter.

Usage Guidelines See the *J-series Services Router Configuration Guide*.

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

See Also configuration-servers on page 540, idle-timeout on page 557

backup-router

Syntax backup-router *address* <destination *destination-address*>;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Set a default router (running IP version 4 [IPv4]) to use while the local router (running IPv4) is booting and if the routing protocol processes fail to start. The JUNOS software removes the route to this router as soon as the software starts.

Options *address*—Address of the default router.

destination *destination-address*—(Optional) Destination address that is reachable through the backup router. Include this option to achieve network reachability while loading, configuring, and recovering the router, but without the risk of installing a default route in the forwarding table.
Default: All hosts (default route) are reachable through the backup router.

Usage Guidelines See “Configuring a Backup Router” on page 376.

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

boot-file

Syntax boot-file *filename*;

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup.

Options *filename*—The location of the boot file on the boot server. The filename can include a pathname.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

See also boot-server (DHCP) on page 534

boot-server

See the following sections:

- boot-server (DHCP) on page 534
- boot-server (NTP) on page 534

boot-server (DHCP)

Syntax boot-server *address*;

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup.

Options *address*—Address of a boot server. You must specify an IPv4 address, not a hostname.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

See also boot-file on page 533

boot-server (NTP)

Syntax boot-server *address*;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the server that NTP queries when the router boots to determine the local date and time.

When you boot the router, it issues an ntpdate request, which polls a network server to determine the local date and time. You need to configure a server that the router uses to determine the time when the router boots. Otherwise, NTP will not be able to synchronize to a time server if the server’s time appears to be very far off of the local router’s time.

Options *address*—Address of an NTP server. You must specify an address, not a hostname.

Usage Guidelines See “Configuring the NTP Boot Server” on page 419.

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

broadcast

Syntax broadcast *address* <key *key-number*> <version *value*> <ttl *value*>;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the local router to operate in broadcast mode with the remote system at the specified *address*. In this mode, the local router sends periodic broadcast messages to a client population at the specified broadcast or multicast *address*. Normally, you include this statement only when the local router is operating as a transmitter.

Options *address*—Address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. Currently, the multicast address must be 224.0.1.1.

key *key-number*—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.
Values: Any unsigned 32-bit integer

ttl *value*—(Optional) Time-to-live (TTL) value to use.
Range: 1 through 255
Default: 1

version *value*—(Optional) Specify the version number to be used in outgoing NTP packets.
Values: 1, 2, 3
Default: 3

Usage Guidelines See “Configuring the NTP Time Server and Time Services” on page 420.

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

broadcast-client

Syntax broadcast-client;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the local router to listen for broadcast messages on the local network to discover other servers on the same subnet.

Usage Guidelines See “Configuring the Router to Listen for Broadcast Messages” on page 424.

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

change-type

Syntax change-type (character-sets | set-transitions);

Hierarchy Level [edit system login passwords]

Release Information Statement introduced in JUNOS Release 7.4.

Description Set requirements for using character sets in plain-text passwords. When combined with the minimum-changes statement, you can check for the total number of character sets included in the password or for the total number of character set changes in the password. Newly created passwords must meet these requirements.

Options One of the following:

- character-sets—The number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters.
- set-transitions—The number of transitions between character sets.

Usage Guidelines See "Configuring Special Requirements for Plain-Text Passwords" on page 381.

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

See also minimum-changes on page 566.

class

See the following sections:

- class (Assign a Class to an Individual User) on page 536
- class (Define Login Classes) on page 537

class (Assign a Class to an Individual User)

Syntax class *class-name*;

Hierarchy Level [edit system login user *username*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a user's login class. You must configure one class for each user.

Options *class-name*—One of the classes defined at the [edit system login class] hierarchy level.

Usage Guidelines See "Configuring User Accounts" on page 413.

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

class (Define Login Classes)

Syntax

```
class class-name {
    allow-commands "regular-expression";
    allow-configuration "regular-expression";
    deny-commands "regular-expression";
    deny-configuration "regular-expression";
    idle-timeout minutes;
    no-world-readable;
    permissions [ permissions ];
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define login classes.

Options *class-name*—A name you choose for the login class.

The remaining statements are explained separately in this chapter.

Usage Guidelines See "Defining Login Classes" on page 401.

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

See Also user on page 608

client-identifier

Syntax client-identifier (ascii *client-id* | hexadecimal *client-id*);

Hierarchy Level [edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure the client's unique identifier. This identifier is used by the DHCP server to index its database of address bindings. Either a client identifier or the client's MAC address is required to uniquely identify the client on the network.

Options *client-id*—A name or number that uniquely identifies the client on the network. The client identifier can be an ASCII string or hexadecimal digits.

Usage Guidelines See "Configuring a DHCP Server" on page 459.

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

commit

Syntax

```
commit {
      allow-transients;
      file filename.xsl {
         optional;
         refresh;
         refresh-from url;
         source url;
      }
      refresh;
      refresh-from url;
      source url;
      traceoptions {
         file filename <files number> <size size>;
         flag flag;
      }
}
```

Hierarchy Level [edit system scripts]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, configure commit-time scripting mechanism.

The statements are explained separately.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

commit synchronize

Syntax

```
commit synchronize;
```

Hierarchy Level [edit system]

Release Information Statement introduced in JUNOS Release 7.4.

Description For multiple Routing Engines only. Configure a commit command to automatically result in a commit synchronize. The Routing Engine on which you execute the commit command (requesting Routing Engine) copies and loads its candidate configuration to the other (responding) Routing Engines. All Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on all Routing Engines.

Usage Guidelines See “Configuring Multiple Routing Engines to Synchronize Configurations Automatically” on page 383.

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

compress-configuration-files

Syntax

```
compress-configuration-files;
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Compress the current operational configuration file. By default, the current operational configuration file is uncompressed, and is stored in the file juniper.conf, in the /config file system, along with the last three committed versions of the configuration. However, with large networks, the current configuration file might exceed the available space in the /config file system. Compressing the current configuration file allows the file to fit in the file system, typically reducing the size of the file by 90 percent. The current configuration file is compressed on the second commit of the configuration after the first commit is made to include the compress-configuration-files statement.

NOTE: We recommend that you enable compression of the router configuration files to minimize the amount of disk space that they require.

Default The current operational configuration file is uncompressed.

Usage Guidelines See “Compressing the Current Configuration File” on page 383.

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

configuration

Syntax

```
configuration {
    transfer-interval interval;
    transfer-on-commit;
    archive-sites {
      ftp://<username>:<password>@<host>:<port>/<url-path>;
    }
}
```

Hierarchy Level [edit system archival]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the router to transfer its currently active configuration by means of FTP periodically or after each commit. The remaining statements are explained separately.

Usage Guidelines See “Configuring a Router to Transfer its Configuration to an Archive Site” on page 483.

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

See Also transfer-interval on page 606, transfer-on-commit on page 607, and archive on page 527.

configuration-servers

Syntax

```
configuration-servers {
        url;
    }
}
```

Hierarchy Level [edit system autoinstallation]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only, configure the URL address of a server from which to obtain configuration files.
Example URLs: tftp://tftpconfig.sp.com/config.conf;
ftp://*user*:*password*@sftpconfig.sp.com/*path*/*file-name*

Usage Guidelines See the *J-series Services Router Getting Started Guide*.

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

See Also autoinstallation on page 532, idle-timeout on page 557.

connection-limit

Syntax connection-limit *limit*;

Hierarchy Level [edit system services finger],
[edit system services ftp],
[edit system services ssh],
[edit system services telnet],
[edit system services xnm-clear-text],
[edit system services xnm-ssl]

Options *limit*—(Optional) Maximum number of established connections.
Range: 1 through 250
Default: 75

Usage Guidelines See "Configuring System Services" on page 458.

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

console

See the following sections:

- console (Physical Port) on page 541
- console (System Logging) on page 542

console (Physical Port)

Syntax

```
console {
    insecure;
    log-out-on-disconnect;
    type terminal-type;
}
```

Hierarchy Level [edit system ports]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the characteristics of the console port, which is on the router's craft interface.

Default The console port is enabled and its speed is 9600 baud.

Options insecure—Disable root login connections to the console and auxiliary ports.

log-out-on-disconnect—Logs out the session when the data carrier on the console port is lost.

type *terminal-type*—Type of terminal that is connected to the port.
Values: ansi, vt100, small-xterm, xterm
Default: The terminal type is unknown, and the user is prompted for the terminal type.

Usage Guidelines See "Configuring Console and Auxiliary Port Properties" on page 456.

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

console (System Logging)

Syntax

```
console {
    facility level;
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the types of messages to log to the system console.

Options *facility*—Class of messages to log. To specify multiple classes, include multiple *facility level* statements. For a list of the facilities, see Table 22 on page 430.

level—Severity of the messages that belong to the facility specified by the paired *facility* name. For a list of the severities, see Table 23 on page 431.

Usage Guidelines See “Directing Messages to the Console” on page 433.

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

See Also *JUNOS System Log Messages Reference*

default-address-selection

Syntax default-address-selection;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Use the loopback interface, lo0, as the source address for all locally generated IP packets. The lo0 interface is the interface to the router’s Routing Engine.

Default The outgoing interface is used as the source address.

Usage Guidelines See “Configuring the Source Address for Locally Generated TCP/IP Packets” on page 457 and the *JUNOS Network Interfaces Configuration Guide*.

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

default-lease-time

Syntax default-lease-time *seconds*;

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Specify the length of time in seconds that a client holds the lease for an IP address assigned by a DHCP server. This setting is used if a lease time is not requested by the client.

Options *seconds*— The number of seconds the lease can be held.

Default: 86400 (1day)

Usage Guidelines See "Configuring a DHCP Server" on page 459.

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

See Also maximum-lease-time on page 565

deny-commands

Syntax deny-commands "*regular-expression*";

Hierarchy Level [edit system login class]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the operational mode commands that the user is denied permission to issue, even though the permissions set with the permissions statement would allow it.

Default If you omit this statement and the allow-commands statement, users can issue only those commands for which they have access privileges through the permissions statement.

Options *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Usage Guidelines See "Specifying Operational Mode Commands" on page 405.

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

See Also allow-commands on page 525, user on page 608

deny-configuration

Syntax deny-configuration "*regular-expression*";

Hierarchy Level [edit system login class]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the configuration mode commands that the user is denied permission to issue, even though the permissions set with the permissions statement would allow it.

Default If you omit this statement and the allow-configuration statement, users can issue only those commands for which they have access privileges through the permissions statement.

Options *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Usage Guidelines See “Specifying Operational Mode Commands” on page 405.

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

See Also allow-configuration on page 525, user on page 608

destination

Syntax

```
destination {
    radius {
      server {
        server-address {
          accounting-port port-number;
          secret password;
          source-address address;
          retry number;
          timeout seconds;
        }
      }
    }
    tacplus {
      server {
        server-address {
          secret password;
          single-connection;
          timeout seconds;
          port port-number;
        }
      }
    }
}
```

Hierarchy Level [edit system accounting]

Description Configure the authentication server.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring RADIUS System Accounting” on page 485 and “Configuring TACACS+ System Accounting” on page 487.

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

dhcp

Syntax

```
dhcp {
    boot-file filename;
    boot-server (address | hostname);
    domain-name domain-name;
    domain-search [domain-list];
    default-lease-time seconds;
    maximum-lease-time seconds;
    name-server {
        address;
    }
    option {
        [ (id-number option-type option-value) | (id-number array option-type
            option-values) ];
    }
    pool (address /prefix-length) {
        address-range {
            low address;
            high address;
        }
        exclude-address {
            address;
        }
    }
    router {
        address;
    }
    static-binding MAC-address {
        fixed-address {
            address;
        }
        host hostname;
        client-identifier (ascii client-id | hexadecimal client-id);
    }
    server-identifier address;
    wins-server {
        address;
    }
}
```

Hierarchy Level [edit system services]

Description For J-series Services Routers only. Configure a router or interface as a DHCP server. A DHCP server can allocate network addresses and deliver configuration information to client hosts on a TCP/IP network.

The remaining statements are explained separately.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

diag-port-authentication

Syntax diag-port-authentication (encrypted-password "*password* " | plain-text-password);

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a password for performing diagnostics on the router's System Control Board (SCB), System and Switch Board (SSB), Switching and Forwarding Module (SFM), or Forwarding Engine Board (FEB) port.

For routers that have more than one SSB, the same password is used for both SSBs.

NOTE: Do not run diagnostics on the SCB, SSB, SFM, or FEB unless you have been instructed to do so by Customer Support personnel.

Default No password is configured on the diagnostics port.

Options encrypted-password "*password*"—Use MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.

plain-text-password—Use a plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password for each user. For information about how to create plain-text passwords, see "Specifying Plain-Text Passwords" on page 17.

Usage Guidelines See "Configuring the Password on the Diagnostics Port" on page 482.

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

domain-name

See the following sections:

- domain-name (DHCP) on page 548
- domain-name (Router) on page 548

domain-name (DHCP)

Syntax domain-name *domain-name*;

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified.

Options *domain-name*—Name of the domain.

Usage Guidelines See "Configuring a DHCP Server" on page 459.

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

domain-name (Router)

Syntax domain-name *domain-name*;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the name of the domain in which the router is located. This is the default domain name that is appended to hostnames that are not fully qualified.

Options *domain-name*—Name of the domain.

Usage Guidelines See "Configuring the Router's Domain Name" on page 374.

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

domain-search

Syntax

```
domain-search [domain-list];
```

Hierarchy Level

```
[edit system],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-bindings]
```

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a list of domains to be searched.

Options *domain-list*—A list of domain names to search. The list can contain up to 6 domain names, with a total of up to 256 characters.

Usage Guidelines See “Configuring Which Domains to Search” on page 374 and “Configuring a DHCP Server” on page 459.

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

dump-device

Syntax

```
dump-device {
    compact-flash;
    removable-compact-flash;
    usb;
}
```

Hierarchy Level

```
[edit system]
```

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure the medium used for storing memory snapshots of system failure. When you specify the storage and an operating system fails, the operating system writes a snapshot of the state of the router when it failed to the storage medium. When the operating system is rebooted, the storage device is checked for a snapshot. If found, the snapshot of memory is written to the /var/crash directory on the router and can be examined by Juniper Networks customer support to help determine the cause of failure.

If the swap partition on the device medium is not large enough for the system memory snapshot, the snapshot not successfully written to the directory. Use the request system snapshot command to specify the swap partition.

Options compact-flash—The primary compact flash.

removable-compact-flash—The compact flash device on the front of the router (J4300 and J6300 only) as the system software failure memory snapshot device.

usb—The device attached to the universal serial bus (USB) port.

Usage Guidelines See the *J-series Services Router Getting Started Guide*.

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

events

Syntax events [*events*];

Hierarchy Level [edit system accounting]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the types of events to track and log.

Options *events*—Event types; can be one or more of the following:

- login—Audit logins.
- change-log—Audit configuration changes.
- interactive-commands—Audit interactive commands (any command-line input).

Usage Guidelines See "Specifying Events" on page 488.

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

explicit-priority

Syntax explicit-priority;

Hierarchy Level [edit system syslog file *filename*],
[edit system syslog host]

Release Information Statement introduced before JUNOS Release 7.4.

Description Record the priority (facility and severity level) in each message logged to a system log file or directed to a remote destination.

Usage Guidelines See "Including Priority in System Log Messages" on page 439.

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

See Also *JUNOS System Log Messages Reference*

facility-override

Syntax facility-override *facility*;

Hierarchy Level [edit system syslog host]

Release Information Statement introduced before JUNOS Release 7.4.

Description Substitute an alternate facility for the default facilities used when messages are directed to a remote destination.

Options *facility*—Alternate facility to substitute for the default facilities. For a list of the possible facilities, see Table 25 on page 436.

Usage Guidelines See "Changing the Alternate Facility Name for Remote Messages" on page 435.

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

See Also *JUNOS System Log Messages Reference*

file

See the following sections:

- file (Commit Scripts) on page 551
- file (System Logging) on page 552

file (Commit Scripts)

Syntax

```
file filename.xsl {
    optional;
    refresh;
    refresh-from url;
    source url;
}
```

Hierarchy Level [edit system scripts commit]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, enable a commit script that is located in the /var/db/scripts/commit directory.

Options *filename*.xsl—The name of an XSLT file containing a commit script.

The statements are explained separately.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

file (System Logging)

Syntax

```
file filename {
    facility level;
    explicit-priority;
    match "regular-expression";
    archive {
      files number;
      size size;
      (world-readable | no-world-readable);
    }
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the types of system logging messages to log to a file.

Options *facility*—Class of messages to log. To specify multiple classes, include multiple *facility level* statements. For a list of the facilities, see Table 22 on page 430.

filename—File in the /var/log directory in which to log messages from the specified facility. To log messages to more than one file, include more than one file statement.

level—Severity of the messages that belong to the facility specified by the paired *facility* name. For a list of the severities, see Table 23 on page 431.

The remaining statements are explained separately.

Usage Guidelines See “Directing Messages to a Log File” on page 432.

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

See Also *JUNOS System Log Messages Reference*

files

Syntax files *number*;

Hierarchy Level [edit system syslog archive],
[edit system syslog file *filename* archive]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the maximum number of archived log files to retain. When the JUNOS logging utility has written a defined maximum amount of data to a log file *logfile*, it closes the file, compresses it, and renames it to *logfile*.0.gz (for information about the maximum file size, see size on page 593). The utility then opens and writes to a new file called *logfile*. When the new file reaches the maximum size, the *logfile*.0.gz file is renamed to *logfile*.1.gz, and the new file is closed, compressed, and renamed *logfile*.0.gz. By default, the logging facility creates up to 10 archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).

Options *number*—Maximum number of archived files.
Range: 1 through 1000
Default: 10 files

Usage Guidelines See “Configuring Log File Archiving” on page 438.

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

See Also size on page 593, *JUNOS System Log Messages Reference*

finger

Syntax

```
finger {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before JUNOS Release 7.4.

Description Allow finger requests from remote systems to the local router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Finger Service” on page 475.

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

format

Syntax format (md5 | sha1 | des);

Hierarchy Level [edit system login passwords]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure the authentication algorithm for plain-text passwords.

Default For JUNOS software, the default encryption format is md5. For JUNOS-FIPS software, the default encryption format is sha1.

Options The hash algorithm that authenticates the password can be one of three algorithms:

- md5—Produces a 128-bit digest.
- sha1—Produces a 160-bit digest.
- des—Has a block size of 8 bytes; its key size is 48 bits long.

Usage Guidelines See “Configuring Special Requirements for Plain-Text Passwords” on page 381.

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

ftp

Syntax

```
ftp {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before JUNOS Release 7.4.

Description Allow FTP requests from remote systems to the local router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring FTP Service” on page 476.

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

full-name

Syntax full-name *complete-name*;

Hierarchy Level [edit system login user]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the complete name of a user.

Options *complete-name*—Full name of the user. If the name contains spaces, enclose it in quotation marks.

Usage Guidelines See "Configuring User Accounts" on page 413.

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

host

Syntax

```
host (hostname | other-routing-engine| scc-master) {
      facility level;
      explicit-priority;
      facility-override facility;
      log-prefix string;
      match "regular-expression";
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the types of system log messages to log to a remote destination.

Options *facility*—Class of messages to log. To specify multiple classes, include multiple *facility level* statements. For a list of the facilities, see Table 22 on page 430.

hostname—IP address or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a host statement for each one.

level—Severity of the messages that belong to the facility specified by the paired *facility* name. For a list of the severities, see Table 23 on page 431.

other-routing-engine—Direct messages to the other Routing Engine on a routing platform with two Routing Engines installed and operational.

scc-master—On a T640 routing node that is part of a routing matrix, direct messages to the TX Matrix platform.

The remaining statements are explained separately.

Usage Guidelines See “Directing Messages to a Remote Machine or the Other Routing Engine” on page 433 and “Directing Messages to a Remote Destination from the Routing Matrix” on page 452.

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

See Also *JUNOS System Log Messages Reference*

host-name

Syntax

```
host-name host-name;
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Set the hostname of the router.

Options *host-name*—Name of the router.

Usage Guidelines See “Configuring the Router’s Name and Addresses” on page 372.

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

http

Syntax

```
http {
    interfaces [ interface-names ];
    port port;
}
```

Hierarchy Level [edit system services web-management]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure port and interfaces for HTTP service, which is unencrypted.

Options interfaces [*interface-names*]—Name of one or more interfaces on which to allow the HTTP service. By default, HTTP access is allowed through built-in Fast Ethernet interfaces only.

The remaining statement is explained separately in this chapter.

Usage Guidelines See the *J-Web Interface User Guide*.

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

See Also https on page 557, port (HTTP/HTTPS) on page 575, web-management on page 610

https

Syntax

```
https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
}
```

Hierarchy Level [edit system services web-management]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the secure version of HTTP (HTTPS) service, which is encrypted.

Options interfaces [*interface-names*]—Name of one or more interfaces on which to allow the HTTPS service. By default, HTTPS access is allowed through any ingress interface, but HTTP access is allowed through built-in Fast Ethernet interfaces only.

local-certificate *name*—Name of the X.509 certificate for a secure sockets layer (SSL) connection. An SSL connection is configured at the [edit security certificates local] hierarchy.

The remaining statement is explained separately in this chapter.

Usage Guidelines See the *J-Web Interface User Guide*.

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

See Also http on page 556, port (HTTP/HTTPS) on page 575, web-management on page 610

idle-timeout

Syntax idle-timeout *minutes*;

Hierarchy Level [edit system login class *class-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For a login class, configure the maximum time that a session can be idle before the user is logged off the router. The session times out after remaining at the CLI operational mode prompt for the specified time.

Default If you omit this statement, a user is never forced off the system after extended idle times.

Options *minutes*—Maximum idle time.
Range: 0 through 100,000 minutes

Usage Guidelines See “Configuring the Timeout Value for Idle Login Sessions” on page 412.

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

See Also user on page 608

inet6-backup-router

Syntax

```
inet6-backup-router address <destination destination-address>;
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Set a default router (running IP version 6 [IPv6]) to use while the local router (running IPv6) is booting and if the routing protocol processes fail to start. The JUNOS software removes the route to this router as soon as the software starts.

Options *address*—Address of the default router.

destination *destination-address*—(Optional) Destination address that is reachable through the backup router. Include this option to achieve network reachability while loading, configuring, and recovering the router, but without the risk of installing a default route in the forwarding table.
Default: All hosts (default route) are reachable through the backup router.

Usage Guidelines See "Configuring a Backup Router" on page 376.

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

interfaces

Syntax

```
interfaces {
        interface-name {
           bootp;
           rarp;
           slarp;
        }
     }
}
```

Hierarchy Level [edit system autoinstallation]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure the interface on which to perform autoinstallation. A request for an IP address is sent from the interface. Specify the IP address procurement protocol.

Options bootp—Sends requests over all available interfaces.

rarp—Sends requests over Ethernet interfaces.

slarp—Sends requests over serial interfaces.

Usage Guidelines See the *J-series Services Router Configuration Guide*.

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

See Also autoinstallation on page 532

internet-options

Syntax

```
internet-options {
    path-mtu-discovery;
    source-port upper-limit upper-limit;
    source-quench;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure path maximum transmission rate (MTU) discovery, source quench, and port range.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Path MTU Discovery” on page 489, “Configuring Source Quench” on page 489, and “Configuring the Range of Port Addresses” on page 490.

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

load-key-file

Syntax

```
load-key-file;
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Load RSA (SSH version 1) and DSA (SSH version 2) public keys from a file. The file is a URL containing one or more SSH keys.

Usage Guidelines See “Configuring the Root Password” on page 379 and “Configuring User Accounts” on page 413.

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

local-certificate

Syntax

```
local-certificate;
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Import or reference a SSL certificate.

Usage Guidelines See “Configuring JUNOScript SSL Service” on page 477 and “Using JUNOScript SSL Service” on page 757.

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

location

Syntax

```
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the system location in various formats.

Options altitude *feet*—Number of feet above sea level.

building *name*—Name of building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").

country-code *code*—Two-letter country code.

floor *number*—Floor in the building.

hcoord *horizontal-coordinate*—Bellcore Horizontal Coordinate.

lata *service-area*—Long-distance service area.

latitude *degrees*—Latitude in degree format.

longitude *degrees*—Longitude in degree format.

npa-nxx *number*—First six digits of the phone number (area code and exchange).

postal-code *postal-code*—Postal code.

rack *number*—Rack number.

vcoord *vertical-coordinate*—Bellcore Vertical Coordinate.

Usage Guidelines See “Configuring the System Location” on page 378.

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

log-prefix

Syntax log-prefix *string*;

Hierarchy Level [edit system syslog host]

Release Information Statement introduced before JUNOS Release 7.4.

Description Include a text string in each message directed to a remote destination.

Options *string*—Text string to include in each message.

Usage Guidelines See “Adding a String to System Log Messages” on page 437.

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

See Also *JUNOS System Log Messages Reference*

login

Syntax

```
login {
    announcement text;
    class class-name {
        allow-commands "regular-expression";
        allow-configuration "regular-expression";
        deny-commands "regular-expression";
        deny-configuration "regular-expression";
        idle-timeout minutes;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    passwords {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
        user username {
        full-name complete-name;
        uid uid-value;
        class class-name;
        authentication authentication;
            (encrypted-password “password” | plain-text-password);
            ssh-rsa “public-key”;
            ssh-dsa “public-key”;
        }
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure user access to the router.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See “Configuring User Access” on page 401.

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

login-alarms

Syntax login-alarms;

Hierarchy Level [edit system login class admin]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Show system alarms automatically when an admin user logs on to the router.

Usage Guidelines See “Configuring System Alarms to Show Automatically” on page 491.

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

See Also *J-series Services Router Administration Guide*

login-tip

Syntax login-tip;

Hierarchy Level [edit system login class *class-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable CLI tips at login.

Default Disabled.

Usage Guidelines See “Configuring Tips” on page 412.

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

See Also “Displaying Tips About CLI Commands” on page 185.

match

Syntax match "*regular-expression*";

Hierarchy Level [edit system syslog file *filename*],
[edit system syslog host (*hostname* | other-routing-engine| scc-master)],
[edit system syslog user (*username* | *)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.

Usage Guidelines See “Using Regular Expressions to Refine the Set of Logged Messages” on page 442.

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

max-configurations-on-flash

Syntax max-configurations-on-flash *number*;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the number of configurations stored on the the flash drive.

Options *number*—The number of configurations stored on the flash drive.
Range: 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.

Usage Guidelines See “Specifying the Number of Configurations Stored on the Flash Drive” on page 484.

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

maximum-lease-time

Syntax maximum-lease-time *seconds*;

Hierarchy Level [edit system services dhcp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Specify the maximum length of time in seconds for which a client can request and hold a lease on a DHCP server.

Exception: Dynamic BOOTP lease length can exceed the maximum lease length specified.

Options *seconds*—The maximum number of seconds the lease can be held.

Usage Guidelines See "Configuring a DHCP Server" on page 459.

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

See Also default-lease-time on page 543

maximum-length

Syntax maximum-length *length*;

Hierarchy Level [edit system login passwords]

Release Information Statement introduced in JUNOS Release 7.4.

Description Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.

Default For JUNOS-FIPS software, the maximum number of characters for plain-text passwords is 20. For JUNOS software, no maximum is set.

Options length—The maximum number of characters the password can include.
Range: 20 to 128 characters

Usage Guidelines See "Configuring Special Requirements for Plain-Text Passwords" on page 381.

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

message

Syntax message *text*;

Hierarchy Level [edit system login]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a system login message. This message appears before a user logs in.

Options *text*—Text of the message.

Usage Guidelines See “Configuring a System Login Message” on page 480.

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

See Also announcement on page 526

minimum-changes

Syntax minimum-changes *number*;

Hierarchy Level [edit system login passwords]

Release Information Statement introduced in JUNOS Release 7.4.

Description Specify the minimum number of character sets (or character set changes) required in plain-text passwords. Newly created passwords must meet this requirement.

This statement is used in combination with the change-type statement. If the change-type is character-sets, then the number of character sets included in the password is checked against the specified minimum. If change-type is set-transitions, then the number of character set changes in the password is checked against the specified minimum.

Default For JUNOS software, the minimum number of changes is 1. For JUNOS-FIPS software, the minimum number of changes is 3.

Options *number*—The minimum number of character sets (or character set changes) required for the password.

Usage Guidelines See “Configuring Special Requirements for Plain-Text Passwords” on page 381.

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

See Also change-type on page 536.

minimum-length

Syntax minimum-length *length*;

Hierarchy Level [edit system login passwords]

Release Information Statement introduced in JUNOS Release 7.4.

Description Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.

Default For JUNOS software, the minimum number of characters for plain-text passwords is 6. For JUNOS-FIPS software, the minimum number of characters for plain-text passwords is 10.

Options length—The minimum number of characters the password must include.
Range: 6 to 20 characters

Usage Guidelines See "Configuring Special Requirements for Plain-Text Passwords" on page 381.

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

See Also maximum-length on page 565.

mirror-flash-on-disk

Syntax mirror-flash-on-disk;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the hard disk to automatically mirror the contents of the compact flash. The hard disk maintains a synchronized mirror copy of the compact-flash contents. Data written to the compact flash is simultaneously updated in the mirrored copy of the hard disk. If the flash drive fails to read data, the hard disk automatically retrieves its mirrored copy of the flash disk.

CAUTION: We recommend that you disable flash disk mirroring when you upgrade or downgrade the router.

You cannot issue the request system snapshot command while flash disk mirroring is enabled.

NOTE: After you have enabled or disabled the mirror-flash-on-disk statement, you must reboot the router for your changes to take effect. To reboot, issue the request system reboot command.

Usage Guidelines See “Configuring Flash Disk Mirroring” on page 377.

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

multicast-client

Syntax multicast-client <*address*>;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For NTP, configure the local router to listen for multicast messages on the local network to discover other servers on the same subnet.

Options *address*—(Optional) One or more IP addresses. If you specify addresses, the router joins those multicast groups.
Default: 224.0.1.1.

Usage Guidelines See “Configuring the Router to Listen for Multicast Messages” on page 425.

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

name-server

Syntax

```
name-server {
    address;
}
```

Hierarchy Level [edit system],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure one or more Domain Name System (DNS) name servers.

Options *address*—Address of the name server. To configure multiple name servers, include multiple *address* options.

Usage Guidelines See “Configuring a DNS Name Server” on page 375 and “Configuring a DHCP Server” on page 459.

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

no-compression-configuration-files

Syntax no-compress-configuration-files;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the router so that it does not compress the current operational configuration.

Usage Guidelines See “Compressing the Current Configuration File” on page 383.

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

no-redirects

Syntax no-redirects;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Disable the sending of protocol redirect messages by the router.

To disable the sending of redirect messages on a per-interface basis, include the no-redirects statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family*] hierarchy level.

Default The router sends redirect messages.

Usage Guidelines See “Disabling the Sending of Redirect Messages on the Router” on page 456.

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

See Also *JUNOS Network Interfaces Configuration Guide*.

no-saved-core-context

Syntax no-saved-core-context;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Do not save core files generated by internal JUNOS processes, or the associated context information.

Default Each core file is saved along with contextual information (configuration and system log files) for debugging purposes in a compressed tar file named /var/tmp/*process-name*.core.*core-number*.tgz.

Usage Guidelines See "Saving Core Files from JUNOS Processes" on page 482.

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

See Also saved-core-context on page 584, saved-core-files on page 585

no-world-readable

See world-readable on page 611

ntp

Syntax

```
ntp {
    authentication-key number type type value password;
    boot-server (NTP) address;
    broadcast <address> <key key-number> <version value> <ttl value>;
    broadcast-client;
    multicast-client <address>;
    peer address <key key-number> <version value> <prefer>;
    server address <key key-number> <version value> <prefer>;
    source-address source-address;
    trusted-key [ key-numbers ];
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure NTP on the router.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring the Network Time Protocol" on page 418.

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

optional

Syntax

```
optional;
```

Hierarchy Level [edit system scripts commit file *filename*]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, allow a commit operation to succeed even if the script specified in the file statement is missing from the /var/db/scripts/commit directory on the routing platform. The optional statement allows the commit operation to progress as if the commit script were not enabled in the configuration.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

passwords

Syntax

```
passwords {
    change-type (set-transitions | character-set);
    format (md5 | sha1 | des);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirement.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring Special Requirements for Plain-Text Passwords" on page 381.

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

See Also maximum-length on page 565.

path-mtu-discovery

Syntax path-mtu-discovery

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure path MTU discovery on outgoing Transmission Control Protocol (TCP) connections.

Usage Guidelines See "Configuring the Path MTU Discovery" on page 489.

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

peer

Syntax peer *address* <key *key-number*> <version *value*> <prefer>;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For NTP, configure the local router to operate in symmetric active mode with the remote system at the specified address. In this mode, the local router and the remote system can synchronize each other. This configuration is useful in a network in which either the local router or the remote system might be a better source of time.

Options *address*—Address of the remote system. You must specify an address, not a hostname.

key *key-number*—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.
Values: Any unsigned 32-bit integer

prefer—(Optional) Mark the remote system as the preferred host, which means that, if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.

version *value*—(Optional) Specify the NTP version number to be used in outgoing NTP packets.
Values: 1, 2, 3
Default: 3

Usage Guidelines See "Configuring the NTP Time Server and Time Services" on page 420.

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

permissions

Syntax permissions [*permissions*];

Hierarchy Level [edit system login class]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the login access privileges to be provided on the router.

Options *permissions*—Privilege type. For a list of types, see Table 17 on page 403.

Usage Guidelines See “Configuring Access Privilege Levels” on page 402.

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

See Also user on page 608

pic-console-authentication

Syntax

```
pic-console authentication {
    (encrypted-password "password" | plain-text-password);
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure console access to Physical Interface Cards (PICs).

Default Disabled. By default, there is no password setting for console access.

Options encrypted-password "*password*"—Use MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.

plain-text-password—Use a plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password. For information about how to create plain-text passwords, see “Specifying Plain-Text Passwords” on page 17.

Usage Guidelines See “Specifying Plain-Text Passwords” on page 17 and “Configuring Console Access to PICs” on page 479.

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

See Also “Configuring Console and Auxiliary Port Properties” on page 456.

pool

Syntax

```
pool (address /prefix-length) {
    address-range {
      low address;
      high address;
    }
    exclude-address {
      address;
    }
}
```

Hierarchy Level [edit system services dhcp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure a pool of IP addresses for DHCP clients on a subnet. When a client joins the network, the DHCP server dynamically allocates an IP address from this pool.

Options address-range—Lowest and highest IP addresses in the pool that are available for dynamic address assignment. If no range is specified, the pool will use all available addresses within the subnet specified. (Broadcast addresses, interface addresses, and excluded addresses are not available.)

exclude-address—Addresses within the range that are not used for dynamic address assignment. You can exclude one or more addresses within the range.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

port

See the following sections:

- port (HTTP/HTTPS) on page 575
- port (RADIUS Server) on page 575
- port (SDX Server) on page 575
- port (TACACS+ Server) on page 576

port (HTTP/HTTPS)

Syntax port *port*;

Hierarchy Level [edit system services web-management]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the port on which the HTTP or HTTPS service is connected.

Options *port*—The TCP port number on which the specified service listens.

Usage Guidelines See the *J-Web Interface User Guide*.

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

See Also web-management on page 610, http on page 556, https on page 557

port (RADIUS Server)

Syntax port *port-number*;

Hierarchy Level [edit system radius-server *address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the port number on which to contact the RADIUS server.

Options *number*—Port number on which to contact the RADIUS server.
Default: 1812 (as specified in RFC 2865)

Usage Guidelines See "Configuring RADIUS Authentication" on page 386.

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

port (SDX Server)

Syntax port *port-number*;

Hierarchy Level [edit system services service-deployment servers *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the port number on which to contact the SDX server.

Options *port-number*—(Optional) The TCP port number for the SDX server.
Default: 3333

Usage Guidelines See "Enabling the SDX Software" on page 489.

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

port (TACACS+ Server)

Syntax port *number*;

Hierarchy Level [edit system accounting destination tacplus server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the port number on which to contact the TACACS+ server.

Options *number*—Port number on which to contact the TACACS+ server.
Default: 49

Usage Guidelines See "Configuring TACACS+ System Accounting" on page 487.

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

ports

Syntax

```
ports {
    auxiliary {
      type terminal-type;
    }
    console {
      type terminal-type;
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the properties of the console and auxiliary ports, which are located on the router's craft interface.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring Console and Auxiliary Port Properties" on page 456.

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

processes

Syntax

```
processes {
    adaptive-services (enable | disable) failover failover-option;
    alarm-control (enable | disable) failover failover-option;
    chassis-control (enable | disable) failover failover-option;
    class-of-service (enable | disable) failover failover-option;
    craft-control (enable | disable) failover failover-option;
    dhcp (enable | disable) failover failover-option;
    disk-monitoring (enable | disable) failover failover-option;
    ecc-error-logging (enable | disable) failover failover-option;
    firewall (enable | disable) failover failover-option;
    inet-process (enable | disable) failover failover-option;
    interface-control (enable | disable) failover failover-option;
    kernel-replication (enable | disable) failover failover-option;
    l2tp-service (enable | disable) failover failover-option;
    link-management (enable | disable) failover failover-option;
    mib-process (enable | disable) failover failover-option;
    network-access-service (enable | disable) failover failover-option;
    ntp (enable | disable) failover failover-option;
    pgm (enable | disable) failover failover-option;
    pic-services-logging (enable | disable) failover failover-option;
    redundancy-device (enable | disable) failover failover-option;
    remote-operations (enable | disable) failover failover-option;
    routing (enable | disable) failover failover-option;
    sampling (enable | disable) failover failover-option;
    service-deployment (enable | disable) failover failover-option;
    snmp (enable | disable) failover failover-option;
    usb-control (enable | disable) failover failover-option;
    timeout seconds;
    watchdog (enable | disable) failover failover-option;
    web-management (enable | disable) failover failover-option;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure which JUNOS software processes are running on the router.

Default All processes are enabled by default

CAUTION: Never disable any of the software processes unless instructed to do so by a Customer Support engineer.

Options failover (alternate-media | other-routing-engine)—(Optional) For routers with redundant Routing Engines only, switch to backup media if a process fails repeatedly. If a process fails three times in quick succession, the router reboots from the alternate media or the other Routing Engine.

timeout *seconds*—(Optional) How often the system checks the watchdog timer, in seconds. If the watchdog timer has not been checked in the specified number of seconds, the system reloads. If you set the time value too low, it is possible for the system to reboot immediately after it loads.
Values: 15, 60, 180
Default: 180 seconds (rounded up to 291 seconds by the JUNOS kernel)

Usage Guidelines See “Disabling JUNOS Software Processes” on page 481.

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

protocol-version

Syntax protocol-version *version*;

Hierarchy Level [edit system services ssh]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the secure shell (SSH) protocol version.

Options *version*—[v1], [v2], or [v1 v2]
Default: [v2]

Usage Guidelines See “Configuring the SSH Protocol Version” on page 478.

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

radius

Syntax

```
radius {
    server {
      server-address {
        accounting-port port-number;
        secret password;
        source-address address;
        retry number;
        timeout seconds;
      }
    }
}
```

Hierarchy Level [edit system accounting destination]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure the RADIUS accounting server.

Options *server-address*—Address of the RADIUS accounting server.

The remaining statements are explained separately.

Usage Guidelines See “Configuring RADIUS System Accounting” on page 485.

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

radius-server

Syntax

```
radius-server server-address {
    port number ;
    retry number ;
    routing-instance routing-instance-name;
    secret password ;
    source-address source-address;
    timeout seconds ;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the Remote Authentication Dial-In User Service (RADIUS) for Point-to-Point Protocol (PPP).

To configure multiple RADIUS servers, include multiple radius-server statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options *server-address*—Address of the RADIUS authentication server.

The remaining statements are explained separately in this chapter.

Usage Guidelines See “Configuring RADIUS Authentication” on page 386.

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

rate-limit

Syntax rate-limit *limit*;

Hierarchy Level [edit system services finger],
[edit system services ftp],
[edit system services ssh],
[edit system services telnet],
[edit system services xnm-clear-text],
[edit system services xnm-ssl]

Description Maximum number of connection attempts on an access service.

Options rate-limit *limit*—(Optional) Maximum number of connection attempts allowed per minute.
Range: 1 through 250
Default: 150

Usage Guidelines See “Configuring System Services” on page 458.

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

refresh

Syntax refresh;

Hierarchy Level [edit system scripts commit],
[edit system scripts commit file *filename*]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the /var/db/scripts/commit directory with the copy located at the source URL, as specified in the source statement.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

See Also refresh-from on page 581, source on page 593

refresh-from

Syntax refresh-from *url*;

Hierarchy Level [edit system scripts commit],
[edit system scripts commit file *filename*]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the /var/db/scripts/commit directory with the copy located at a URL other than the URL specified in the source statement.

Options *url*—The source specified as an HTTP URL, FTP URL, or SCP-style remote file specification.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

See Also refresh on page 580, source on page 593

retry

Syntax retry *number*;

Hierarchy Level [edit system radius-server *server-address*],
[edit system accounting destination radius server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Number of times that the router attempts to contact a RADIUS authentication or accounting server.

Options *number*—Number of times to retry contacting a RADIUS server.
Range: 1 through 10
Default: 3

Usage Guidelines See “Configuring RADIUS Authentication” on page 386 and “Configuring RADIUS System Accounting” on page 485.

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

See Also timeout on page 602

root-authentication

Syntax

```
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-dsa "public-key";
    ssh-rsa "public-key";
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the authentication methods for the root-level user, whose username is “root.”

Options encrypted-password "*password*"—Use MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.

plain-text-password—Use a plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.For information about how to create plain-text passwords, see “Specifying Plain-Text Passwords” on page 17.

ssh-dsa "*public-key*"—SSH version 2 authentication. Specify the DSA (SSH version 2) public key. You can specify one or more public keys.

ssh-rsa "*public-key*"—SSH version 1 authentication. Specify the RSA (SSH version 1) public key. You can specify one or more public keys.

Usage Guidelines See “Configuring the Root Password” on page 379.

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

See Also authentication on page 529

root-login

Syntax

```
root-login (allow | deny | deny-password);
```

Hierarchy Level [edit system services ssh]

Release Information Statement introduced before JUNOS Release 7.4.

Description Control user access through SSH.

Options allow—Allow users to log in to the router as root through SSH.
Default: allow

deny—Disable users from logging in to the router as root through SSH.

deny-password—Allow users to log in to the router as root through SSH when the authentication method (for example, RSA authentication) does not require a password.

Usage Guidelines See “Configuring the Root Login” on page 478.

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

See Also “Configuring SSH Service” on page 477.

router

Syntax

```
router {
    address;
}
```

Hierarchy Level [edit system services dhcp-service],
[edit system services dhcp-service pool],
[edit system services dhcp-service static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Specify IPv4 addresses for one or more routers available to a DHCP client. List routers in order of preference.

Options address—IPv4 address of the router. To configure multiple routers, include multiple address options.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit system radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the routing instance used to send RADIUS packets to the RADIUS server.

Options *routing-instance-name*—Routing instance name.

Usage Guidelines See "Configuring RADIUS Authentication" on page 386.

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

See also *JUNOS Network Interfaces Configuration Guide*

saved-core-context

Syntax saved-core-context;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Save each core file generated by internal JUNOS processes, along with contextual information (configuration and system log files), for debugging purposes in a compressed tar file named /var/tmp/*process-name*.core.*core-number*.tgz.

Usage Guidelines See "Saving Core Files from JUNOS Processes" on page 482.

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

See Also no-saved-core-context on page 570, saved-core-files on page 585

saved-core-files

Syntax save-core-files *number*;

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Save core files generated by internal JUNOS processes, but not the associated contextual information (configuration and system log files).

Options *number*—Maximum number of core files to save.

Range: 1 through 64

Usage Guidelines See "Saving Core Files from JUNOS Processes" on page 482.

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

See Also saved-core-context on page 584.

scripts

Syntax

```
scripts {
    commit {
      file filename.xsl {
        optional;
        refresh;
        refresh-from url;
        source url;
      }
      traceoptions {
        file filename <files number> <size size>;
        flag flag;
      }
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, configure scripting mechanisms.

The statements are explained separately.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

secret

Syntax secret *password* ;

Hierarchy Level [edit system accounting destination radius server *server-address*],
[edit system accounting destination tacplus server *server-address*],
[edit system radius-server *server-address*],
[edit system tacplus-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local router must match that used by the server.

Options *password*—Password to use; can include spaces.

Usage Guidelines See "Configuring RADIUS Authentication" on page 386, "Configuring TACACS+ Authentication" on page 388, "Configuring TACACS+ System Accounting" on page 487, and "Configuring RADIUS System Accounting" on page 485.

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

server

See the following sections:

- server (NTP) on page 587
- server (RADIUS Accounting) on page 588
- server (TACACS+ Accounting) on page 588

server (NTP)

Syntax server *address* <key *key-number*> <version *value*> <prefer>;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For NTP, configure the local router to operate in client mode with the remote system at the specified *address*. In this mode, the local router can be synchronized to the remote system, but the remote system never can be synchronized to the local router.

Options *address*—Address of the remote system. You must specify an address, not a hostname.

key *key-number*—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.
Values: Any unsigned 32-bit integer

prefer—(Optional) Mark the remote system as preferred host, which means that, if all other are equal, this remote system is chosen for synchronization among a set of correctly operating systems.

version *value*—(Optional) Specify the version number to be used in outgoing NTP packets.
Values: 1, 2, 3
Default: 3

Usage Guidelines See “Configuring the NTP Time Server and Time Services” on page 420.

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

server (RADIUS Accounting)

Syntax

```
server {
    server-address {
        accounting-port port-number;
        secret password;
        source-address address;
        retry number
        timeout seconds;
    }
}
```

Hierarchy Level [edit system accounting destination radius]

Description Configure RADIUS logging.

The remaining statements are explained separately.

Usage Guidelines See “Configuring RADIUS System Accounting” on page 485.

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

server (TACACS+ Accounting)

Syntax

```
server {
    server-address {
        port port-number;
        secret password;
        single-connection;
        timeout seconds;
    }
}
```

Hierarchy Level [edit system accounting destination tacplus]

Description Configure TACACS+ logging.

The remaining statements are explained separately.

Usage Guidelines See “Configuring TACACS+ System Accounting” on page 487.

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

server-identifier

Syntax server-identifier *address*;

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Configure a server identifier. This is an optional setting that can be used to identify a DHCP server in a DHCP message. It can also be used as a destination address from clients to servers (for example, when the boot file is set, but not the boot server).

Servers include the server identifier in DHCPOFFER messages so that clients can distinguish between multiple lease offers. Clients include the server identifier in DHCPREQUEST messages to select a lease and indicate which offer is accepted from multiple lease offers. Also, clients can use the server identifier to send unicast request messages to specific DHCP servers to renew a current lease.

This address must be a manually assigned, static IP address. The server cannot send a request and receive an IP address from itself or another DHCP server.

Options *address*—The IPv4 address of the server. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

Default If no server identifier is set, the DHCP server sets the server identifier based on the primary interface address used by the server to receive a client request. For example, if the client sends a DHCP request and the server receives it on fe-0/0/0 and the primary interface address is 1.1.1.1, then the server identifier is set to 1.1.1.1.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

servers

Syntax

```
servers server-address {
      port port-number;
}
```

Hierarchy Level [edit system services service-deployment]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure an IPv4 address for the Service Deployment System (SDX) server.

Options *server-address*—The TCP port number.
Default: 3333

The remaining statement is explained separately in this chapter.

Usage Guidelines See "Enabling the SDX Software" on page 489.

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

service-deployment

Syntax

```
service-deployment {
      servers server-address {
         port port-number;
      }
      source-address source-address;
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable JUNOS software to work with the SDX software.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See "Enabling the SDX Software" on page 489.

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

services

Syntax

```
services {
    dchp {
        dhcp_services;
    }
    finger {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    ftp {
            <connection-limit limit>;
            <rate-limit limit>;
    }
    ssh {
        root-login (allow | deny | deny-password);
        protocol-version [v1 v2];
        <connection-limit limit>;
        <rate-limit limit >;
    }
    service-deployment {
        servers server-address {
            port-number port-number;
        }
        source-address source-address;
    }
    telnet {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    web-management {
        http {
            interfaces [ interface-names ];
            port port;
        }
        https {
            interfaces [ interface-names ];
            local-certificate name;
            port port;
        }
    }
    xnm-clear-text {
        <connection-limit limit>;
        <rate-limit limit>;
    }
    xnm-ssl {
        <connection-limit limit>;
        <rate-limit limit>;
        <local-certificate name>
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the router so that users on remote systems can access the local router through the DHCP server, finger, rlogin, SSH, telnet, Web management, JUNOScript clear-text, JUNOScript SSL, and network utilities or enable JUNOS software to work with the SDX software.

The statements are explained separately.

Usage Guidelines See “Configuring System Services” on page 458 and “Enabling the SDX Software” on page 489.

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

single-connection

Syntax single-connection;

Hierarchy Level [edit system tacplus-server *server-address*],
[edit system accounting destination tacplus-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests, rather than opening a connection for each connection attempt.

Usage Guidelines See “Configuring TACACS+ Authentication” on page 388 and “Configuring TACACS+ System Accounting” on page 487.

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

size

Syntax size *size*;

Hierarchy Level [edit system syslog archive],
[edit system syslog file *filename* archive]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the maximum amount of data that the JUNOS logging utility writes to a log file *logfile* before archiving it (closing it, compressing it, and changing its name to *logfile*.0.gz). The utility then opens and writes to a new file called *logfile*. For information about the number of archive files that the utility creates in this way, see files on page 553.

Options *size*—Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).
Syntax: *x*k to specify the number of kilobytes, *x*m for the number of megabytes, or *x*g for the number of gigabytes
Range: 64 KB through 1 GB
Default: 128 KB

Usage Guidelines See "Configuring System Log Messages" on page 427.

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

See Also files on page 553, *JUNOS System Log Messages Reference*

source

Syntax source *url*;

Hierarchy Level [edit system scripts commit],
[edit system scripts commit file *filename*]

Release Information Statement introduced in JUNOS Release 7.4.

Description For JUNOS commit scripts, specify the location of the source file for all enabled commit scripts or a single enabled script located in the /var/db/scripts/commit directory. When you include the refresh statement, the source URL is the location from which the local copy is refreshed.

Options *url*—The source specified as an HTTP URL, FTP URL, or SCP-style remote file specification.

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

See Also refresh on page 580, refresh-from on page 581

source-address

- source-address (NTP, RADIUS, System Logging, or TACACS+) on page 594
- source-address (SDX Software) on page 594

source-address (NTP, RADIUS, System Logging, or TACACS+)

Syntax source-address *source-address*;

Hierarchy Level [edit system accounting destination radius server *server-address*],
[edit system accounting destination tacplus server *server-address*],
[edit system ntp],
[edit system radius-server *server-address*],
[edit system syslog],
[edit system tacplus-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify a source address for each configured TACACS+ server, RADIUS server, NTP server, or the source address to use when directing system log messages to a remote machine.

Options *source-address*—A valid IP address configured on one of the router interfaces. For system logging, the address is used for messages sent to the remote machines specified in all the host *hostname* statements at the [edit system syslog] hierarchy level, but not for messages directed to the other Routing Engine or to the TX Matrix platform in a routing matrix.

Usage Guidelines See "Specifying a Source Address for RADIUS and TACACS+ Servers" on page 390, "Specifying a Source Address for an NTP Server" on page 420, and "Specifying an Alternate Source Address" on page 434.

See Also set date ntp source-address on page 291

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

source-address (SDX Software)

Syntax source-address *source-address*;

Hierarchy Level [edit system services service-deployment]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable JUNOS software to work with the SDX software.

Options *source-address*—(Optional) The local IPv4 address to be used as source address for traffic to the SDX server. The source address restricts traffic within the out-of-band network.

Usage Guidelines See “Enabling the SDX Software” on page 489.

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

source-port

Syntax source-port upper-limit <*upper-limit*>;

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the range of port addresses.

Options upper-limit *upper-limit*—(Optional) The range of port addresses and can be a value from 5000 through 65,355.

Usage Guidelines See “Configuring the Range of Port Addresses” on page 490.

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

source-quench

Syntax source-quench

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the JUNOS software to ignore Internet Control Message Protocol (ICMP) source quench messages.

Usage Guidelines See “Configuring Source Quench” on page 489.

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

ssh

Syntax

```
ssh {
      root-login (allow | deny | deny-password);
      protocol-version [v1 v2];
      <connection-limit limit>;
      <rate-limit limit>;
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before JUNOS Release 7.4.

Description Allow ssh requests from remote systems to the local router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring SSH Service” on page 477.

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

static-binding

Syntax

```
static-binding MAC-address {
      fixed-address {
         address;
      }
      host client-hostname;
      client-identifier (ascii client-id | hexadecimal client-id);
}
```

Hierarchy Level [edit system services dhcp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Set static bindings for DHCP clients. A static binding is a mapping between a fixed IP address and the client’s MAC address or client identifier.

Options *MAC-address*—The MAC address of the client. This is a hardware address that uniquely identifies a client on the network.

fixed-address *address*—Fixed IP address assigned to the client. Typically a client has one address assigned, but you can assign more.

host client-*hostname*—Hostname of the client requesting the DHCP server. The name can include the local domain name. Otherwise, the name is resolved based on the domain-name statement.

client-identifier (ascii *client-id* | hexadecimal *client-id*);—Used by the DHCP server to index the database of address bindings. The client identifier is an ASCII string or hexadecimal number and can include a type-value pair as specified in RFC 1700, *Assigned Numbers*. Either a client identifier or the client's MAC address must be configured to uniquely identify the client on the network.

Usage Guidelines See "Configuring a DHCP Server" on page 459.

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

static-host-mapping

Syntax

```
static-host-mapping {
    host-name {
        inet [ address ];
        sysid system-identifier ;
        alias [ alias ];
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Map a hostname to one or more IP addresses and aliases, and configure an International Organization for Standardization (ISO) system identifier (system ID).

Options alias *alias*—(Optional) Alias for the hostname.

host-name—Fully qualified hostname.

inet *address*—IP address. You can specify one or more IP addresses for the host.

sysid *system-identifier*—ISO system identifier (system ID). This is the 6-byte portion of the Intermediate System-to-Intermediate System (IS-IS) network service access point (NSAP). We recommend that you use the host's IP address represented in binary-coded decimal (BCD) format. For example, the IP address 208.197.169.18 is 2081.9716.9018 in BCD.

Usage Guidelines See "Configuring the Router's Name and Addresses" on page 372.

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

syslog

Syntax

```
syslog {
    archive {
        files number;
        size size;
        (world-readable | no-world-readable);
    }
    console {
        facility severity;
    }
    file filename {
        facility severity;
        explicit-priority;
        match "regular-expression";
        archive {
            files number;
            size size;
            (world-readable | no-world-readable);
        }
    }
    host (hostname | other-routing-engine | scc-master) {
        facility severity;
        explicit-priority;
        facility-override facility;
        log-prefix string;
        match "regular-expression";
    }
    source-address source-address;
    time-format (year | millisecond | year millisecond);
    user (username | *) {
        facility severity;
        match "regular-expression";
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the types of system log messages to log to files, a remote destination, user terminals, or the system console.

The statements are explained separately.

Usage Guidelines See “Configuring System Log Messages” on page 427.

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

See Also *JUNOS System Log Messages Reference*

system

Syntax system { ... }

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure system management properties.

Usage Guidelines See “System Management Configuration Statements” on page 365.

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

tacplus

Syntax

```
tacplus {
    server {
      server-address {
        port port-number;
        secret password;
        single-connection;
        timeout seconds;
      }
    }
}
```

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the Terminal Access Controller Access Control System Plus (TACACS+).

Hierarchy Level [edit system accounting destination]

Options *server-address*—Address of the TACACS+ authentication server.

The remaining statements are explained separately.

Usage Guidelines See “Configuring TACACS+ System Accounting” on page 487.

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

tacplus-options

Syntax

```
tacplus-options service-name service-name;
```

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure multiple TACACS+ servers to use the same authentication service.

Hierarchy Level [edit system]

Options service-name *service-name*—The name of the authentication service.

Default: junos-exec

Usage Guidelines See “Configuring the Same Authentication Service for Multiple TACACS+ Servers” on page 391.

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

tacplus-server

Syntax

```
tacplus-server server-address {
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
```

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the TACACS+ server.

Hierarchy Level [edit system]

Options *server-address*—Address of the TACACS+ authentication server.

The remaining statements are explained separately.

Usage Guidelines See “Configuring TACACS+ Authentication” on page 388.

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

telnet

Syntax

```
telnet {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

Hierarchy Level [edit system services]

Description Allow telnet connections from remote systems to the local router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring System Services” on page 458.

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

time-format

Syntax

```
time-format (year | millisecond | year millisecond);
```

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.

Description Include the year, the millisecond, or both, in the timestamp on every system log message. The additional information is included for messages directed to each destination configured by a file, console, or user statement at the [edit system syslog] hierarchy level, but not to destinations configured by a host statement.

By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged, for example, Aug 21 12:36:30.

Options millisecond—Include the millisecond in the timestamp.

year—Include the year in the timestamp.

Usage Guidelines See “Including the Year or Millisecond in Timestamps” on page 442.

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

See Also *JUNOS System Log Messages Reference*

timeout

Syntax timeout *seconds*;

Hierarchy Level [edit system radius-server *server-address*],
[edit system tacplus-server *server-address*],
[edit system accounting destination radius server server-address],
[edit system accounting destination tacplus server server-address]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the amount of time that the local router waits to receive a response from a RADIUS or TACACS+ server.

Options *seconds*—Amount of time to wait.
Range: 1 through 90 seconds
Default: 3 seconds

Usage Guidelines See "Configuring RADIUS Authentication" on page 386 and "Configuring TACACS+ Authentication" on page 388.

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

See Also retry on page 581

time-zone

Syntax time-zone (GMT*hour-offset* | *time-zone*);

Hierarchy Level [edit system]

Release Information Statement introduced before JUNOS Release 7.4.
GMT*hour-offset* option added in JUNOS Release 7.4

Description Set the local time zone. To have the time zone change take effect for all processes running on the router, you must reboot the router.

Default UTC

Options GMT*hour-offset*—Set the time zone relative to GMT time.
Range: -14 through +12
Default: 0

time-zone—Specify the time zone either as UTC, which is the default time zone, or use one of the following continents and major cities:

Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako, Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali, Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi, Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu, Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis, Africa/Windhoek

America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires, America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit, America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung, America/Paramaribo, America/Phoenix, America/Port-au-Prince, America/Port_of_Spain, America/Porto_Acre, America/Puerto_Rico, America/Rainy_River, America/Rankin_Inlet, America/Regina, America/Rosario, America/Santiago, America/Santo_Domingo, America/Sao_Paulo, America/Scoresbysund, America/Shiprock, America/St_Johns, America/St_Kitts, America/St_Lucia, America/St_Thomas, America/St_Vincent, America/Swift_Current, America/Tegucigalpa, America/Thule, America/Thunder_Bay, America/Tijuana, America/Tortola, America/Vancouver, America/Whitehorse, America/Winnipeg, America/Yakutat, America/Yellowknife
Antarctica/Casey, Antarctica/DumontDUrville, Antarctica/Mawson, Antarctica/McMurdo, Antarctica/Palmer, Antarctica/South_Pole

Arctic/Longyearbyen

Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqtau, Asia/Aqtobe, Asia/Ashkhabad, Asia/Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut, Asia/Bishkek, Asia/Brunei, Asia/Calcutta, Asia/Chungking, Asia/Colombo, Asia/Dacca, Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/Harbin, Asia/Hong_Kong, Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/Jerusalem, Asia/Kabul, Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Krasnoyarsk, Asia/Kuala_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/Manila, Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom_Penh, Asia/Pyongyang, Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul, Asia/Shanghai, Asia/Singapore, Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran, Asia/Thimbu, Asia/Tokyo, Asia/Ujung_Pandang, Asia/Ulan_Bator, Asia/Urumqi, Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk, Asia/Yekaterinburg, Asia/Yerevan

Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape_Verde, Atlantic/Faeroe, Atlantic/Jan_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South_Georgia, Atlantic/St_Helena, Atlantic/Stanley

Australia/Adelaide, Australia/Brisbane, Australia/Broken_Hill, Australia/Darwin, Australia/Hobart, Australia/Lindeman, Australia/Lord_Howe, Australia/Melbourne, Australia/Perth, Australia/Sydney

Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade, Europe/Berlin, Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest, Europe/Chisinau, Europe/Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki, Europe/Istanbul, Europe/Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana, Europe/London, Europe/Luxembourg, Europe/Madrid, Europe/Malta, Europe/Minsk, Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/Paris, Europe/Prague, Europe/Riga, Europe/Rome, Europe/Samara, Europe/San_Marino, Europe/Sarajevo, Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn, Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius, Europe/Warsaw, Europe/Zagreb, Europe/Zurich

Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro, Indian/Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte, Indian/Reunion

Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate, Pacific/Enderbury, Pacific/Fakaofo, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos, Pacific/Gambier, Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston, Pacific/Kiritimati, Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas, Pacific/Midway, Pacific/Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea, Pacific/Pago_Pago, Pacific/Palau, Pacific/Pitcairn, Pacific/Ponape, Pacific/Port_Moresby, Pacific/Rarotonga, Pacific/Saipan, Pacific/Tahiti, Pacific/Tarawa, Pacific/Tongatapu, Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/Yap

Usage Guidelines See “Setting the Time Zone” on page 417.

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

traceoptions

Syntax

```
traceoptions {
    file filename <files number> <size size>;
    flag flag;
}
```

Hierarchy Level [edit system scripts commit]

Description Define tracing operations for the commit scripts.

Default If you do not include this statement, no commit-script-specific tracing operations are performed.

Options *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, commit script process tracing output is placed in the file cscript.log.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file*.0, then *trace-file*.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the size option.

Range: 2 through 1000
Default: 10 files

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

- all—Log all operations
- events—Log important events
- input—Log commit script input data
- offline—Generate data for offline development
- output—Log commit script output data
- rpc—Log commit script RPCs
- xslt—Log the XSLT library

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file*.0. When the *trace-file* again reaches its maximum size, *trace-file*.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.

Syntax: *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB
Range: 128 KB through 1 GB
Default: 128 KB

Usage Guidelines See the *JUNOS Configuration Scripting Guide*.

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

transfer-interval

Syntax transfer-interval *interval*;

Hierarchy Level [edit system archival configuration]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the router to periodically transfer its currently active configuration to an archive site.

Option *interval*—Time interval to transfer the current configuration to an archive site.
Range: 15 through 2880 minutes

Usage Guidelines See "Configuring the Transfer Interval" on page 483.

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

See Also configuration on page 539, transfer-on-commit on page 607, archive on page 527

transfer-on-commit

Syntax transfer-on-commit;

Hierarchy Level [edit system archival configuration]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the router to transfer its currently active configuration to an archive site each time you commit a candidate configuration.

Usage Guidelines See “Configuring Transfer on Commit” on page 483.

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

See Also configuration on page 539, transfer-interval on page 606, archive on page 527

trusted-key

Syntax trusted-key [*key-numbers*];

Hierarchy Level [edit system ntp]

Release Information Statement introduced before JUNOS Release 7.4.

Description For NTP, configure the keys you are allowed to use when you configure the local router to synchronize its time with other systems on the network.

Options *key-numbers*—One or more key numbers. Each key can be any 32-bit unsigned integer except 0.

Usage Guidelines See “Configuring NTP Authentication Keys” on page 424.

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

See Also authentication-key on page 530, broadcast on page 535, peer on page 572, server on page 587

uid

Syntax uid *uid-value*;

Hierarchy Level [edit system login user]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a user identifier for a login account.

Options *uid-value*—Number associated with the login account. This value must be unique on the router.
Range: 100 through 64,000

Usage Guidelines See "Configuring User Access" on page 401.

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

user

See the following sections:

- user (Access) on page 608
- user (System Logging) on page 609

user (Access)

Syntax

```
user username {
    full-name complete-name;
    uid uid-value;
    class class-name;
    authentication {
        (encrypted-password "password" | plain-text-password);
        ssh-rsa "public-key";
        ssh-dsa "public-key";
    }
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure access permission for individual users.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring User Access" on page 401.

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

See Also class on page 536

user (System Logging)

Syntax

```
user (username | *) {
    facility level;
    match "regular-expression";
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the types of system log messages to log to user terminals.

Options * (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.

facility—Class of messages to log. To specify multiple classes, include multiple *facility level* statements. For a list of the facilities, see Table 22 on page 430.

level—Severity of the messages that belong to the facility specified by the paired *facility* name. For a list of the severities, see Table 23 on page 431.

username—JUNOS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one user statement.

The remaining statement is explained separately.

Usage Guidelines See "Directing Messages to a User Terminal" on page 433.

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

See Also *JUNOS System Log Messages Reference*

web-management

Syntax

```
web-management {
    http {
        interfaces [ interface-names ];
        port port;
        }
    https {
        interfaces [ interface-names ];
        local-certificate name;
        port port;
    }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure settings for the HTTP or HTTPS access. HTTP access allows management of the router via the browser-based J-Web graphical user interface. HTTPS access allows secure management of the router via the J-Web interface. With HTTPS access, communication between the router Web server and your browser is encrypted.

Options The remaining statements are explained separately in this chapter.

Usage Guidelines See the *J-Web Interface User Guide*.

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

See Also http on page 556, https on page 557, port (HTTP/HTTPS) on page 575

wins-server

Syntax

```
wins-server {
    address;
}
```

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before JUNOS Release 7.4.

Description For J-series Services Routers only. Specify one or more NetBIOS Name Servers. When a DHCP client is added to the network and assigned an IP address, the NetBIOS Name Server manages the Windows Internet Name Service (WINS) database that matches IP addresses (such as 192.168.1.3) to Windows NetBIOS names (such as \\Marketing). List servers in order of preference.

Options *address*—IPv4 address of the NetBIOS Name Server running WINS. To configure multiple servers, include multiple *address* options.

Usage Guidelines See “Configuring a DHCP Server” on page 459.

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

world-readable

Syntax world-readable | no-world-readable;

Hierarchy Level [edit system syslog]

Release Information Statement introduced before JUNOS Release 7.4.

Description Grant all users permission to read log files, or restrict the permission only to the root user and users who have the JUNOS maintenance permission.

Default no-world-readable

Usage Guidelines See “Configuring Log File Archiving” on page 438.

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

See Also *JUNOS System Log Messages Reference*

xnm-clear-text

Syntax

```
xnm-clear-text {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

Hierarchy Level [edit system services]

Description Allow JUNOScript clear-text requests from remote systems to the local router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring JUNOScript Clear-Text Service” on page 476.

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

xnm-ssl

Syntax

```
xnm-ssl {
    <connection-limit limit>;
    <rate-limit limit>;
}
```

Hierarchy Level [edit system services]

Description Allow JUNOScript SSL requests from remote systems to the local router.

The remaining statements are explained separately.

Usage Guidelines See “Configuring JUNOScript SSL Service” on page 477.

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

Part 5
Configuration Groups

Chapter 28
Configuration Groups

This chapter discusses the following topics:

Overview

Configuration groups allow you to create a group containing configuration statements and to direct the inheritance of that group's statements in the rest of the configuration. The same group can be applied to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.

Configuration groups allow you to create smaller, more logically constructed configuration files, making it easier to configure and maintain the JUNOS software. For example, you can group statements that are repeated in many places in the configuration, such as when configuring interfaces, and thereby limit updates to just the group.

You can also use wildcards in a configuration group to allow configuration data to be inherited by any object that matches a wildcard expression.

The configuration group mechanism is separate from the grouping mechanisms used elsewhere in the configuration, such as Border Gateway Protocol (BGP) groups. Configuration groups provide a generic mechanism that can be used throughout the configuration but that are known only to the JUNOS software command-line interface (CLI). The individual software processes that perform the actions directed by the configuration receive the expanded form of the configuration; they have no knowledge of configuration groups.

Inheritance Model

Configuration groups use true inheritance, which involves a dynamic, ongoing relationship between the source of the configuration data and the target of that data. Data values changed in the configuration group are automatically inherited by the target. The target need not contain the inherited information, although the inherited values can be overridden in the target without affecting the source from which they were inherited.

This inheritance model allows you to see only the instance-specific information without seeing the inherited details. A command pipe in configuration mode allows you to display the inherited data.

Configuration Groups Configuration Statements

To configure configuration groups and inheritance, you can include the groups statement at the [edit] hierarchy level:

```
[edit]
groups {
    group-name {
      configuration-data;
    }
}
```

Include the apply-groups [*group-names*] statement anywhere in the configuration that the configuration statements contained in a configuration group are needed.

Configuration Groups Configuration Guidelines

For areas of your configuration to inherit configuration statements, you must first put the statements into a configuration group and then apply that group to the levels in the configuration hierarchy that require the statements. This section covers the following topics:

- Creating a Configuration Group on page 618
- Applying a Configuration Group on page 619
- Disabling Inheritance of a Configuration Group on page 622
- Displaying Inherited Values on page 623
- Using Wildcards on page 624

Creating a Configuration Group

To create a configuration group, include the groups statement at the [edit] hierarchy level:

```
[edit]
groups {
    group-name {
      configuration-data;
    }
    lccn-re0 {
      configuration-data;
    }
    lccn-re1 {
      configuration-data;
    }
}
```

group-name is the name of a configuration group. You can configure more than one configuration group by specifying multiple *group-name* statements. However, you cannot use the prefix junos- in a group name because it is reserved for use by the JUNOS software.

One reason for the naming restriction is a configuration group called junos-defaults. This preset configuration group is applied to the configuration automatically. You cannot modify or remove the junos-defaults configuration group. For more information about the JUNOS default configuration group, see "Using JUNOS Default Groups" on page 639.

On routers that support multiple Routing Engines, you can also specify two special group names:

- re0—Configuration statements applied to the Routing Engine in slot 0.
- re1—Configuration statements applied to the Routing Engine in slot 1.

The configuration specified in group re0 is only applied if the current Routing Engine is in slot 0; likewise, the configuration specified in group re1 is only applied if the current Routing Engine is in slot 1. Therefore, both Routing Engines can use the same configuration file, each using only the configuration statements that apply to it. Each re0 or re1 group contains at a minimum the configuration for the hostname and the management interface (fxp0). If each Routing Engine uses a different management interface, the group also should contain the configuration for the backup router and static routes.

configuration-data contains the configuration statements applied elsewhere in the configuration with the apply-groups statement. To have a configuration inherit the statements in a configuration group, include the apply-groups statement. For information about the apply-groups statement, see "Applying a Configuration Group" on page 619.

In addition, the TX Matrix platform supports group names for the Routing Engines in each T640 routing node. Providing special group names for all Routing Engines in the routing matrix allows you to configure the individual Routing Engines in each T640 routing node differently. Parameters that are not configured at the [edit groups] hierarchy level apply to all Routing Engines in the routing matrix.

The group names have the following formats:

- lcc*n*-re0—Configuration statements applied to the Routing Engine in slot 0 in a specified T640 routing node.
- lcc*n*-re1—Configuration statements applied to the Routing Engine in slot 1 in a specified T640 routing node.

n identifies the T640 routing node and can be from 0 through 3. For example, to configure Routing Engine 1 properties for lcc3, you include statements at the [edit groups lcc3–re1] hierarchy level. For information about the TX Matrix platform and routing matrix, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

Applying a Configuration Group

To have a configuration inherit the statements in a configuration group, include the apply-groups statement:

```
apply-groups [ group-names ];
```

If you specify more than one group name, list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.

For routers that support multiple Routing Engines, you can specify re0 and re1 group names. The configuration specified in group re0 is only applied if the current Routing Engine is in slot 0; likewise, the configuration specified in group re1 is only applied if the current Routing Engine is in slot 1. Therefore, both Routing Engines can use the same configuration file, each using only the configuration statements that apply to it. Each re0 or re1 group contains at a minimum the configuration for the hostname and the management interface (fxp0). If each Routing Engine uses a different management interface, the group also should contain the configuration for the backup router and static routes.

You can include the apply-groups statement at any level of the configuration hierarchy, listing group names within each apply-groups statement in priority order.

You can include only one apply-groups statement at each specific level of the configuration hierarchy. The apply-groups statement at a specific hierarchy level lists the configuration groups to be added to the containing statement's list of configuration groups.

Values specified at the specific hierarchy level override values inherited from the configuration group.

Groups listed in nested apply-groups statements take priority over groups in outer statements. In the following example, the BGP neighbor 10.0.0.1 inherits configuration data from group one first, then from groups two and three. Configuration data in group one overrides data in any other group. Data from group ten is used only if a statement is not contained in any other group.

```
apply-groups [ eight nine ten ];
protocols {
    apply-groups seven;
    bgp {
        apply-groups [ five six ];
        group some-bgp-group {
            apply-groups four;
            neighbor 10.0.0.1 {
                apply-groups [ one two three ];
            }
        }
    }
}
```

Example: Configuring and Applying Configuration Groups

In this example, the Simple Network Management Protocol (SNMP) configuration is divided between the group basic and the normal configuration hierarchy.

There are a number of advantages to placing the system-specific configuration (SNMP contact) into a configuration group and thus separating it from the normal configuration hierarchy—the user can replace (using the load replace command) either section without discarding data from the other.

In addition, setting a contact for a specific box is now possible because the group data would be hidden by the router-specific data.

```
[edit]
groups {
    basic {         # User-defined group name
        snmp {      # This group contains some SNMP data
            contact "My Engineering Group";
            community BasicAccess {
                authorization read-only;
            }
        }
    }
}
apply-groups basic;    # Enable inheritance from group "basic"
snmp {                 # Some normal (non-group) configuration
    location "West of Nowhere";
}
```

This configuration is equivalent to the following:

```
[edit]
snmp {
    location "West of Nowhere";
    contact "My Engineering Group";
    community BasicAccess {
        authorization read-only;
    }
}
```

For information about how to disable inheritance of a configuration group, see “Disabling Inheritance of a Configuration Group” on page 622.

Example: Creating and Applying Configuration Groups on a TX Matrix Platform

```
[edit]
groups {
    re0 { # Routing Engine 0 on TX Matrix platform
        system {
            host-name <host-name>;
            backup-router <ip-address>;
        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address <ip-address>;
                    }
                }
            }
        }
    }
    re1 { # Routing Engine 1 on TX Matrix platform
        system {
            host-name <host-name>;
            backup-router <ip-address>;
        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address <ip-address>;
                    }
                }
            }
        }
    }
```

```
        lcc0-re0 { # Routing Engine 0 on T640 routing node numbered 0
          system {
            host-name <host-name>;
            backup-router <ip-address>;
          }
          interfaces {
            fxp0 {
              unit 0 {
              family inet {
              address <ip-address>;
            }
          }
        }
        lcc0-re1 { # Routing Engine 1 on T640 routing node numbered 0
          system {
          host-name <host-name>;
          backup-router <ip-address>;
        }
          interfaces {
            fxp0 {
              unit 0 {
                family inet {
                address <ip-address>;
                }
            }
        }
    }
    apply-groups [ re0 re1 lcc0-re0 lcc0-re1 ];
```

Disabling Inheritance of a Configuration Group

To disable inheritance of a configuration group at any level except the top level of the hierarchy, include the apply-groups-except statement:

```
apply-groups-except [ group-names ];
```

This useful when you use the apply-group statement at a specific hierarchy level but also want to override the values inherited from the configuration group for a specific parameter.

Example: Disabling Inheritance on Interface s0-1/1/0

In the following example, the apply-groups statement is applied globally at the interfaces level. The apply-groups-except statement is also applied at interface s0-1/1/0 so that it uses the default values hold-time and link-mode.

```
[edit]
groups {              # "groups" is a top-level statement
    global {          # User-defined group name
        interfaces {
            <*> {
            hold-time down 640;
            link-mode full-duplex;
        }
    }
}
```

```
apply-groups global;
interfaces {
    so-1/1/0 {
      apply-groups-except global;  # Disables inheritance from group "global":
    }                               # so-1/1/0 uses default values for “hold-time”
                                    # and “link-mode”
}
```

For information about applying a configuration group, see “Applying a Configuration Group” on page 619.

Displaying Inherited Values

Configuration groups can add some confusion regarding the actual values used by the router, because configuration data can be inherited from configuration groups. To view the actual values used by the router, use the display inheritance command after the pipe in a show command. This command displays the inherited statements at the level at which they are inherited and the group from which they have been inherited.

```
[edit]
user@host# show | display inheritance
snmp {
    location "West of Nowhere";
    ##
    ## ’My Engineering Group’ was inherited from group ’basic’
    ##
    contact "My Engineering Group";
    ##
    ## ’BasicAccess’ was inherited from group ’basic’
    ##
    community BasicAccess {
      ##
      ## ’read-only’ was inherited from group ’basic’
      ##
      authorization read-only;
    }
}
```

To display the expanded configuration (the configuration, including the inherited statements) without the ## lines, use the except command after the pipe in a show command:

```
[edit]
user@host# show | display inheritance | except ##
snmp {
    location "West of Nowhere";
    contact "My Engineering Group";
    community BasicAccess {
      authorization read-only;
    }
}
```

Using Wildcards

You can use wildcards to identify names and allow one statement to provide data for a variety of statements. For example, grouping the configuration of the sonet-options statement over all SONET/SDH interfaces or the dead interval for Open Shortest Path First (OSPF) over all Asynchronous Transfer Mode (ATM) interfaces simplifies configuration files and eases their maintenance.

Wildcarding in normal configuration data is done in a style that is consistent with traditional UNIX shell name wildcarding. In this style of wildcarding, you can use the following metacharacters:

- Asterisk (*)—Matches any string of characters.
- Question mark (?)—Matches any single character.
- Open bracket ([ )—Introduces a character class.
- Close bracket ( ])—Indicates the end of a character class. If the close bracket is missing, the open bracket matches a [rather than introduce a character class.
- A character class matches any of the characters between the square brackets. Character classes must be enclosed in quotation marks (“ ”).
- Hyphen (-)—Specifies a range of characters.
- Exclamation point (!)—The character class can be complemented by making an exclamation point the first character of the character class. To include a] in a character class, make it the first character listed (after the !, if any). To include a minus sign, make it the first or last character listed.

Wildcarding in configuration groups follows the same rules, but the wildcard pattern must be enclosed in angle brackets (<*pattern*>) to differentiate it from other wildcarding in the configuration file. For example:

```
[edit]
groups {
    sonet-default {
        interfaces {
            <so-*> {
                sonet-options {
                    payload-scrambler;
                    rfc-2615;
                }
            }
        }
    }
}
```

Wildcard expressions match (and provide configuration data for) existing statements in the configuration that match their expression only. In the example above, the expression <so-*> passes its sonet-options statement to any interface that matches the expression so-*.

Angle brackets allow you to pass normal wildcarding through without modification. In all matching within the configuration, whether it is done with or without wildcards, the first item encountered in the configuration that matches is used. In the following example, data from the wildcarded BGP groups is inherited in the order in which the groups are listed. The preference value from <*a*> overrides the preference in <*b*>, just as the p value from <*c*> overrides the one from <*d*>. Data values from any of these groups override the data values from abcd.

```
[edit]
user@host# show
groups {
    one {
        protocols {
            bgp {
                group <*a*> {
                    preference 1;
                }
                group <*b*> {
                    preference 2;
                }
                group <*c*> {
                    out-delay 3;
                }
                group <*d*> {
                    out-delay 4;
                }
                group abcd {
                    preference 10;
                    hold-time 10;
                    out-delay 10;
                }
            }
        }
    }
}
protocols {
    bgp {
        group abcd {
            apply-groups one;
        }
    }
}
```

```
[edit]
user@host# show | display inheritance
protocols {
    bgp {
      group abcd {
        ##
        ## '1' was inherited from group 'one'
        ##
        preference 1;
        ##
        ## '10' was inherited from group 'one'
        ##
        hold-time 10;
        ##
        ## '3' was inherited from group 'one'
        ##
        out-delay 3;
      }
    }
}
```

Example: Using Wildcards

The following example demonstrates the use of wildcarding. The interface so-0/0/0 inherits data from the various SONET/SDH interface wildcard patterns in group one.

```
[edit]
user@host# show
groups {
    one {
      interfaces {
        <so-*> {
          sonet-options {
            rfc-2615;
          }
        }
        <so-0/*> {
          sonet-options {
            fcs 32;
          }
        }
        <so-*/0/*> {
          sonet-options {
            fcs 16;
          }
        }
        <so-*/*/0> {
          sonet-options {
            payload-scrambler;
          }
        }
      }
    }
}
```

```
apply-groups one;
interfaces {
    so-0/0/0 {
      unit 0 {
        family inet {
          address 10.0.0.1/8;
        }
      }
    }
}
[edit]
user@host# show | display inheritance
interfaces {
    so-0/0/0 {
      ##
      ## 'sonet-options' was inherited from group 'one'
      ##
      sonet-options {
        ##
        ## '32' was inherited from group 'one'
        ##
        fcs 32;
        ##
        ## 'payload-scrambler' was inherited from group 'one'
        ##
        payload-scrambler;
        ##
        ## 'rfc-2615' was inherited from group 'one'
        ##
        rfc-2615;
      }
      unit 0 {
        family inet {
          address 10.0.0.1/8;
        }
      }
    }
}
```

Examples: Configuration Groups

The following examples illustrate ways to use configuration groups and inheritance:

- Configuring Sets of Statements on page 628
- Configuring Interfaces on page 629
- Configuring a Consistent Management IP Address on page 631

- Configuring Peer Entities on page 633
- Establishing Regional Configurations on page 635
- Selecting Wildcard Names on page 637

Configuring Sets of Statements

When sets of statements exist in configuration groups, all values are inherited. For example:

```
[edit]
user@host# show
groups {
    basic {
        snmp {
            interface so-1/1/1.0;
        }
    }
}
apply-groups basic;
snmp {
    interface so-0/0/0.0;
}
[edit]
user@host# show | display inheritance
snmp {
    ##
    ## 'so-1/1/1.0' was inherited from group 'basic'
    ##
    interface [ so-0/0/0.0 so-1/1/1.0 ];
}
```

For sets that are not displayed within brackets, all values are also inherited. For example:

```
[edit]
user@host# show
groups {
    worldwide {
        system {
            name-server {
                10.0.0.100;
                10.0.0.200;
            }
        }
    }
}
apply-groups worldwide;
system {
    name-server {
        10.0.0.1;
        10.0.0.2;
    }
}
```

```
[edit]
user@host# show | display inheritance
system {
    name-server {
    ##
    ## '10.0.0.100' was inherited from group 'worldwide'
    ##
      10.0.0.100;
    ##
    ## '10.0.0.200' was inherited from group 'worldwide'
    ##
      10.0.0.200;
    }
}
```

Configuring Interfaces

You can use configuration groups to separate the common interface media parameters from the interface-specific addressing information. The following example places configuration data for ATM interfaces into a group called atm-options:

```
[edit]
user@host# show
groups {
    atm-options {
      interfaces {
        <at-*> {
          atm-options {
            vpi 0 maximum-vcs 1024;
          }
          unit <*> {
            encapsulation atm-snap;
            point-to-point;
            family iso;
          }
        }
      }
    }
}
apply-groups atm-options;
interfaces {
    at-0/0/0 {
      unit 100 {
        vci 0.100;
        family inet {
          address 10.0.0.100/30;
        }
      }
      unit 200 {
        vci 0.200;
        family inet {
          address 10.0.0.200/30;
        }
      }
    }
}
```

```
[edit]
user@host# show | display inheritance
interfaces {
    at-0/0/0 {
      ##
      ## "atm-options" was inherited from group "atm-options"
      ##
      atm-options {
        ##
        ## "1024" was inherited from group "atm-options"
        ##
        vpi 0 maximum-vcs 1024;
      }
      unit 100 {
        ##
        ## "atm-snap" was inherited from group "atm-options"
        ##
        encapsulation atm-snap;
        ##
        ## "point-to-point" was inherited from group "atm-options"
        ##
        point-to-point;
        vci 0.100;
        family inet {
          address 10.0.0.100/30;
        }
        ##
        ## "iso" was inherited from group "atm-options"
        ##
        family iso;
      }
      unit 200 {
        ##
        ## "atm-snap" was inherited from group "atm-options"
        ##
        encapsulation atm-snap;
        ##
        ## "point-to-point" was inherited from group "atm-options"
        ##
        point-to-point;
        vci 0.200;
        family inet {
          address 10.0.0.200/30;
        }
        ##
        ## "iso" was inherited from group "atm-options"
        ##
        family iso;
      }
    }
}
```

```
[edit]
user@host# show | display inheritance | except ##
interfaces {
    at-0/0/0 {
      atm-options {
        vpi 0 maximum-vcs 1024;
      }
      unit 100 {
        encapsulation atm-snap;
        point-to-point;
        vci 0.100;
        family inet {
          address 10.0.0.100/30;
        }
        family iso;
      }
      unit 200 {
        encapsulation atm-snap;
        point-to-point;
        vci 0.200;
        family inet {
          address 10.0.0.200/30;
        }
        family iso;
      }
    }
}
```

Configuring a Consistent Management IP Address

On platforms with multiple Routing Engines, each Routing Engine is configured with a separate IP address for the management interface (fxp0). To access the master Routing Engine, you must know which Routing Engine is active and use the appropriate IP address.

Optionally, for consistent access to the master Routing Engine, you can configure an additional IP address and use this address for the management interface regardless of which Routing Engine is active. This additional IP address is active only on the management interface for the master Routing Engine. During switchover, the address moves to the new master Routing Engine.

In the following example, address 10.17.40.131 is configured for both Routing Engines and includes a master-only statement. With this configuration, the 10.17.40.131 address is active only on the master Routing Engine. The address remains consistent regardless of which Routing Engine is active. Address 10.17.40.132 is assigned to fxp0 on re0, and 10.17.40.133 is assigned to fxp0 on re1.

```
[edit groups re0 interfaces fxp0]
unit 0 {
    family inet {
      address 10.17.40.131/25 {
              master_only;
       }
      address 10.17.40.132/25;
    }
}

[edit groups re1 interfaces fxp0]
unit 0 {
    family inet {
      address 10.17.40.131/25 {
              master_only;
       }
      address 10.17.40.133/25;
    }
}
```

This feature is available on all platforms that include dual Routing Engines. On the TX Matrix platform, this feature is applicable to the switch-card chassis (SCC) only.

Configuring Peer Entities

In this example, we create a group some-isp that contains configuration data relating to another Internet service provider (ISP). We can then insert apply-group statements at any point to allow any location in the configuration hierarchy to inherit this data.

```
[edit]
user@host# show
groups {
    some-isp {
      interfaces {
        <ge-*> {
          gigether-options {
            flow-control;
          }
        }
      }
      protocols {
        bgp {
          group <*> {
            neighbor <*> {
              remove-private;
            }
          }
        }
        pim {
          interface <*> {
            version 1;
          }
        }
      }
    }
}
interfaces {
    ge-0/0/0 {
      apply-groups some-isp;
      unit 0 {
        family inet {
          address 10.0.0.1/24;
        }
      }
    }
}
protocols {
    bgp {
      group main {
        neighbor 10.254.0.1 {
          apply-groups some-isp;
        }
      }
    }
    pim {
      interface ge-0/0/0.0 {
        apply-groups some-isp;
      }
    }
}
[edit]
user@host# show | display inheritance
interfaces {
    ge-0/0/0 {
      ##
      ## "gigether-options" was inherited from group "some-isp"
      ##
      gigether-options {
        ##
        ## "flow-control" was inherited from group "some-isp"
        ##
        flow-control;
      }
      unit 0 {
        family inet {
          address 10.0.0.1/24;
        }
      }
    }
}
protocols {
    bgp {
      group main {
        neighbor 10.254.0.1 {
          ##
          ## "remove-private" was inherited from group "some-isp"
          ##
          remove-private;
        }
      }
    }
    pim {
      interface ge-0/0/0.0 {
        ##
        ## "1" was inherited from group "some-isp"
        ##
        version 1;
      }
    }
}
```

Establishing Regional Configurations

In this example, one group is populated with configuration data that is standard throughout the company, while another group contains regional deviations from this standard:

```
[edit]
user@host# show
groups {
    standard {
        interfaces {
            <t3-*> {
                t3-options {
                    compatibility-mode larscom subrate 10;
                    idle-cycle-flag ones;
                }
            }
        }
    }
    northwest {
        interfaces {
            <t3-*> {
                t3-options {
                    long-buildout;
                    compatibility-mode kentrox;
                }
            }
        }
    }
}
apply-groups standard;
interfaces {
    t3-0/0/0 {
        apply-groups northwest;
    }
}
```

```
[edit]
user@host# show | display inheritance
interfaces {
    t3-0/0/0 {
        ##
        ## "t3-options" was inherited from group "northwest"
        ##
        t3-options {
            ##
            ## "long-buildout" was inherited from group "northwest"
            ##
            long-buildout;
            ##
            ## "kentrox" was inherited from group "northwest"
            ##
            compatibility-mode kentrox;
            ##
            ## "ones" was inherited from group "standard"
            ##
            idle-cycle-flag ones;
        }
    }
}
```

Selecting Wildcard Names

You can combine wildcarding and thoughtful use of names in statements to tailor statement values:

```
[edit]
user@host# show
groups {
    mpls-conf {
      protocols {
        mpls {
          label-switched-path <*-major> {
            retry-timer 5;
            bandwidth 155m;
            optimize-timer 60;
          }
          label-switched-path <*-minor> {
            retry-timer 15;
            bandwidth 64k;
            optimize-timer 120;
          }
        }
      }
    }
}
apply-groups mpls-conf;
protocols {
    mpls {
      label-switched-path metro-major {
        to 10.0.0.10;
      }
      label-switched-path remote-minor {
        to 10.0.0.20;
      }
    }
}
```

```
[edit]
user@host# show | display inheritance
protocols {
    mpls {
      label-switched-path metro-major {
        to 10.0.0.10;
        ##
        ## "5" was inherited from group "mpls-conf"
        ##
        retry-timer 5;
        #
        ## "155m" was inherited from group "mpls-conf"
        ##
        bandwidth 155m;
        ##
        ## "60" was inherited from group "mpls-conf"
        ##
        optimize-timer 60;
      }
      label-switched-path remote-minor {
        to 10.0.0.20;
        ##
        ## "15" was inherited from group "mpls-conf"
        ##
        retry-timer 15;
        ##
        ## "64k" was inherited from group "mpls-conf"
        ##
        bandwidth 64k;
        ##
        ## "120" was inherited from group "mpls-conf"
        ##
        optimize-timer 120;
      }
    }
}
```

Using JUNOS Default Groups

The JUNOS software provides a hidden and immutable configuration group called junos-defaults that is automatically applied to the configuration of your routing platform. The junos-defaults group contains preconfigured statements that contain predefined values for common applications. Some of the statements must be referenced to take effect, such as definitions for applications (for example, FTP or telnet settings). Other statements are applied automatically, such as terminal settings.

NOTE: Many identifiers included in the junos-defaults configuration group begin with the name junos-. Because identifiers beginning with the name junos- are reserved for use by Juniper Networks, you cannot define any configuration objects using this name.

You cannot specify identifiers contained in the junos-defaults configuration group in an apply-groups statement.

To view the full set of available preset statements from the JUNOS default group, issue the show groups junos-defaults configuration mode command at the top level of the configuration. The following example displays a partial list of JUNOS default groups:

```
user@host# show groups junos-defaults
#
# Make vt100 the default for the console port
#
system {
    ports {
        console type vt100;
    }
}
applications {
    #
    # File Transfer Protocol
    #
    application junos-ftp {
        application-protocol ftp;
        protocol tcp;
        destination-port 21;
    }
    #
    # Trivial File Transfer Protocol
    #
    application junos-tftp {
        application-protocol tftp;
        protocol udp;
        destination-port 69;
    }
    #
    # RPC port mapper on TCP
    #
```

```
        application junos-rpc-portmap-tcp {
            application-protocol rpc-portmap;
            protocol tcp;
            destination-port 111;
        }
        #
        # RPC port mapper on UDP
        #
        }
    }
```

To reference statements available from the junos-defaults group, include the selected junos- *default-name* statement at the applicable hierarchy level.

Example: Referencing the Preset Statement from the JUNOS Defaults Group

The following example is a preset statement from the JUNOS defaults group that is available for FTP in a stateful firewall:

```
[edit]
groups {
    junos-defaults {
        applications {
            application junos-ftp {          # Use FTP default configuration
                application-protocol ftp;
                protocol tcp;
                destination-port 21;
            }
        }
    }
}
```

To reference a preset JUNOS default statement from the JUNOS defaults group, include the junos- *default-name* statement at the applicable hierarchy level. For example, to reference the JUNOS default statement for FTP in a stateful firewall, include the junos-ftp statement at the [edit services stateful-firewall rule *rule-name* term *term-name* from applications] hierarchy level:

```
[edit]
services {
    stateful-firewall {
        rule my-rule {
            term my-term {
                from {
                    applications junos-ftp; #Reference predefined statement, junos-ftp,
                }                           #for FTP in the stateful firewall configuration
            }
        }
    }
}
```

Example: Viewing Default Statements That Have Been Applied to the Configuration

To view the JUNOS defaults that have been applied to the configuration, issue the show | display inheritance defaults command. For example, to view the inherited JUNOS defaults at the [edit system ports] hierarchy level:

```
user@host# show system ports | display inheritance defaults
## ## 'console' was inherited from group 'junos-defaults'
## 'vt100' was inherited from group 'junos-defaults'
## console type vt100;
```

If you choose not to use existing JUNOS default statements, you can create your own configuration groups manually. For more information about manually creating of configuration groups, see "Overview" on page 616 and "Configuration Groups Configuration Statements" on page 617.

Summary of Configuration Group Statements

The following sections explain each of the configuration group statements. The statements are organized alphabetically.

apply-groups

Syntax apply-groups [*group-names*];

Hierarchy Level All hierarchy levels

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.

You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.

For routers that support multiple Routing Engines, you can specify re0 and re1 as group names. The configuration specified in group re0 is applied only if the current Routing Engine is in slot 0; likewise, the configuration specified in group re1 is applied only if the current Routing Engine is in slot 1. Therefore, both Routing Engines can use the same configuration file, each using only the configuration statements that apply to it. Each re0 or re1 group contains at a minimum the configuration for the hostname and the management interface (fxp0). If each Routing Engine uses a different management interface, the group also should contain the configuration for the backup router and static routes.

You can include the **apply-groups** statement at any level of the configuration hierarchy.

You can include only one **apply-groups** statement at each specific level of the configuration hierarchy. The **apply-groups** statement at a specific hierarchy level lists the configuration groups to be added to the containing statement's list of configuration groups.

Options *group-name*—One or more names specified in the **groups** statement.

Usage Guidelines See "Applying a Configuration Group" on page 619.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also groups on page 642

apply-groups-except

Syntax

```
apply-groups-except [ group-names ];
```

Hierarchy Level All hierarchy levels except the top level

Description Disables inheritance of a configuration group.

Options *group-names*—One or more names specified in the **groups** statement.

Usage Guidelines See "Disabling Inheritance of a Configuration Group" on page 622.

Required Privilege Level configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.

See Also groups on page 642

groups

Syntax

```
groups {
    group-name {
      configuration-data;
    }
    lccn-re0 {
      configuration-data;
    }
    lccn-re1 {
      configuration-data;
    }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Create a configuration group.

Options *configuration-data*—The configuration statements that are to be applied elsewhere in the configuration with the apply-groups statement, to have the target configuration inherit the statements in the group.

group-name—Name of the configuration group. To configure multiple groups, specify more than one *group-name*. On routers that support multiple Routing Engines, you can also specify two special group names:

- re0—Configuration statements that are to be applied to the Routing Engine in slot 0.
- re1—Configuration statements that are to be applied to the Routing Engine in slot 1.

The configuration specified in group re0 is applied only if the current Routing Engine is in slot 0; likewise, the configuration specified in group re1 is applied only if the current Routing Engine is in slot 1. Therefore, both Routing Engines can use the same configuration file, each using only the configuration statements that apply to it. Each re0 or re1 group contains at a minimum the configuration for the hostname and the management interface (fxp0). If each Routing Engine uses a different management interface, the group also should contain the configuration for the backup router and static routes.

(Routing matrix only) The TX Matrix platform supports group names for the Routing Engines in each connected T640 routing node in the following formats:

- lcc*n*-re0—Configuration statements applied to the Routing Engine in slot 0 of the specified T640 routing node that is connected to a TX Matrix platform.
- lcc*n*-re1—Configuration statements applied to the specified to the Routing Engine in slot 1 of the specified T640 routing node that is connected to a TX Matrix platform.

n identifies the T640 routing node and can be from 0 through 3.

Usage Guidelines See “Creating a Configuration Group” on page 618.

Required Privilege Level configure—To enter configuration mode.

See Also apply-groups on page 641, apply-groups-except on page 642

Part 6

Access

Chapter 29

Configuring Access

To configure access, include the following statements at the [edit access] hierarchy level:

```
[edit access]
address-pool pool-name {
    address address-or-prefix;
    address-range low <lower-limit> high <upper-limit >;
}
group-profile profile-name {
    l2tp {
        interface-id interface-id;
        lcp-renegotiation;
        local-chap;
        maximum-sessions-per-tunnel number;
        multilink {
            drop-timeout milliseconds;
            fragmentation-threshold bytes;
        }
    }
    ppp {
        framed-pool pool-id;
        idle-timeout seconds;
        interface-id interface-id;
        keepalive seconds;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
    }
}
profile profile-name {
    authentication-order [ authentication-methods ];
    client client-name {
        chap-secret chap-secret;
        group-profile profile-name;
        ike {
            allowed-proxy pair {
                remote remote-proxy-address local local-proxy-address;
            }
            pre-shared-key [ascii-text key-string] [hexadecimal key-string];
            interface-id interface-id;
        }
```

```
          l2tp {
            interface-id interface-id;
            lcp-renegotiation;
            local-chap;
            maximum-sessions-per-tunnel number;
            multilink {
              drop-timeout milliseconds;
              fragmentation-threshold bytes;
            }
            ppp-authentication (chap | pap);
            ppp-profile profile-name;
            shared-secret shared-secret;
          }
          pap-password pap-password;
          ppp {
            framed-ip-address ip-address;
            framed-pool framed-pool;
            idle-timeout seconds;
            interface-id interface-id;
            keepalive seconds;
            primary-dns primary-dns;
            primary-wins primary-wins;
            secondary-dns secondary-dns;
            secondary-wins secondary-wins;
          }
          user-group-profile profile-name;
        }
        radius-server server-address {
          accounting-port port-number;
          port port-number;
          retry attempts;
          routing-instance routing-instance-name;
          secret password;
          source-address source-address;
          timeout seconds;
        }
    }
    radius-disconnect {
        client-address {
          secret password;
        }
    }
    radius-disconnect-port port-number;
    radius-server server-address {
        accounting-port port-number;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
    }
```

```
traceoptions {
    flag all;
    flag authentication;
    flag chap;
    flag configuration;
    flag kernel;
    flag radius;
}
```

This chapter discusses the following topics:

- Configuring the Point-to-Point Protocol on page 649
- Tracing Access Processes on page 654
- Configuring the Layer 2 Tunneling Protocol on page 655
- Configuring RADIUS Authentication for L2TP on page 677
- Configuring an Internet Key Exchange (IKE) Access Profile on page 681

Configuring the Point-to-Point Protocol

To configure the Point-to-Point Protocol (PPP), do the following:

- Configuring the Challenge Handshake Authentication Protocol on page 649
- Example: CHAP Authentication with RADIUS on page 651

Configuring the Challenge Handshake Authentication Protocol

The Challenge Handshake Authentication Protocol (CHAP) allows each end of a PPP link to authenticate its peer, as defined in RFC 1994. The authenticator sends its peer a randomly generated challenge that the peer must encrypt using a one-way hash; the peer must then respond with that encrypted result. The key to the hash is a secret known only to the authenticator and authenticated. When the response is received, the authenticator compares its calculated result with the peer's response. If they match, the peer is authenticated.

Each end of the link identifies itself to its peer by including its name in the CHAP challenge and response packets it sends to the peer. This name defaults to the local hostname, or you can explicitly set it using the local-name option. When a host receives a CHAP challenge or CHAP response packet on a particular interface, it uses the peer identity to look up the CHAP secret key to use. For more information about the local-name option, see the *JUNOS Network Interfaces Configuration Guide*.

To configure CHAP, include the profile statement at the [edit access] hierarchy level:

```
[edit access]
profile profile-name {
    client client-name chap-secret chap-secret;
}
```

Then reference the CHAP profile name at the [edit interfaces] hierarchy level. For more information about how to reference CHAP, see the *JUNOS Network Interfaces Configuration Guide*.

You can configure multiple profiles. You can also configure multiple clients for each profile.

profile is the mapping between peer identifiers and CHAP secret keys. The identity of the peer contained in the CHAP challenge or response queries the profile for the secret key to use.

client is the peer identity.

chap-secret secret is the secret key associated with that peer.

The following examples show authentication using CHAP:

- Example: PPP Challenge Handshake Authentication Protocol on page 650
- Example: CHAP Authentication with RADIUS on page 651

Example: PPP Challenge Handshake Authentication Protocol

Configure the profile pe-A-ppp-clients at the [edit access] hierarchy level, then reference it at the [edit interfaces] hierarchy level:

```
[edit]
access {
    profile pe-A-ppp-clients {
        client cpe-1 chap-secret "$1$dQYsZ$B5ojUeUjDsUo.yKwcCZ0";
            # SECRET-DATA
        client cpe-2 chap-secret "$1$kdAsfaDAfkdjDsASxfafdKdFKJ";
            # SECRET-DATA
    }
}
interfaces {
    so-1/1/1 {
        encapsulation ppp;
        ppp-options {
            chap {
                access-profile pe-A-ppp-clients;
                local-name "pe-A-so-1/1/1";
            }
        }
    }
    so-1/1/2 {
        encapsulation ppp;
        ppp-options {
            chap {
                passive;
                access-profile pe-A-ppp-clients;
                local-name "pe-A-so-1/1/2";
            }
        }
    }
}
```

Example: CHAP Authentication with RADIUS

You can send RADIUS messages through a routing instance to customer RADIUS servers in a private network.To configure, include the routing-instance statement at the [edit access profile profile-name radius-server] hierarchy level and apply the profile to an interface with the access-profile statement at the [edit interfaces *interface-name* unit *logical-unit-number* ppp-options chap] hierarchy level.

In this example, PPP peers of interfaces at-0/0/0.0 and at-0/0/0.1 are authenticated by a RADIUS server reachable via routing instance A. PPP peers of interfaces at-0/0/0.2 and at-0/0/0.3 are authenticated by a RADIUS server reachable via routing instance B.

For more information on RADIUS authentication, see “Configuring RADIUS Authentication” on page 386.

```
system {
  radius-server {
    1.1.1.1 secret $9$dalkfj;
    2.2.2.2 secret $9$adsfaszx;
  }
}
routing-instances {
  A {
    instance-type vrf;
    ...
  }
  B {
    instance-type vrf;
    ...
  }
}
access {
 profile A-PPP-clients {
    authentication-order radius;
    radius-server {
      3.3.3.3 {
        port 3333;
        secret "$9$LO/7NbDjqmPQGDmT"; # # SECRET-DATA
        timeout 3;
        retry 3;
        source-address 99.99.99.99;
        routing-instance A;
      }
      4.4.4.4 {
        routing-instance A;
        secret $9$adsfaszx;
      }
    }
  }
  profile B-PPP-clients {
    authentication-order radius;
    radius-server {
      5.5.5.5 {
        routing-instance B;
        secret $9$kljhlkhl;
      }
```

```
            6.6.6.6 {
                routing-instance B;
                secret $9$kljhlkhl;
            }
        }
    }
}
interfaces {
    at-0/0/0 {
        atm-options {
            vpi 0;
        }
        unit 0 {
            encapsulation atm-ppp-llc;
            ppp-options {
                chap {
                    access-profile A-PPP-clients;
                }
            }
            keepalives {
                interval 20;
                up-count 5;
                down-count 5;
            }
            vci 0.128;
            family inet {
                address 21.21.21.21/32 {
                    destination 21.21.21.22;
                }
            }
        }
        unit 1 {
            encapsulation atm-ppp-llc;
            ...
            ppp-options {
                chap {
                    access-profile A-PPP-clients;
                }
            }
            ...
        }
        unit 2 {
            encapsulation atm-ppp-llc;
            ...
            ppp-options {
                chap {
                    access-profile B-PPP-clients;
                }
            }
            ...
        }
```

```
            unit 3 {
                encapsulation atm-ppp-llc;
                ...
                ppp-options {
                    chap {
                        access-profile B-PPP-clients;
                    }
                }
                ...
            }
            ...
        }
        ...
    }
```

Users who log in to the router with telnet or SSH connections are authenticated by the RADIUS server 1.1.1.1. The backup RADIUS server for these users is 2.2.2.2.

Each profile may contain one or more backup RADIUS servers. In this example, PPP peers are CHAP authenticated by the RADIUS server 3.3.3.3 (with 4.4.4.4. as the backup server) or RADIUS server 5.5.5.5 (with 6.6.6.6 as the backup server).

Configuring the Authentication Order

You can configure the order in which the JUNOS software tries different authentication methods when authenticating peers. For each access attempt, the software tries the authentication methods in order, from first to last.

To configure the authentication order, include the authentication-order statement at the [edit access profile *name*] hierarchy level:

```
[edit access profile profile-name]
authentication-order [ authentication-methods ];
```

In *authentication-methods*, specify one or more of the following in the preferred order, from first tried to last tried:

- radius—Verify the client using RADIUS authentication services.
- password—Verify the client using the information configured at the [edit access profile *profile-name* client *client-name*] hierarchy level.

If you do not include the authentication-order statement, clients are verified by means of password authentication.

The CHAP authentication sequence cannot take more than 30 seconds. If it takes longer to authenticate a client, the authentication is abandoned and a new sequence is initiated.

For example, if you configure three RADIUS servers so that the router attempts to contact each server three times, and with each retry the server times out after 3 seconds, then the maximum time given to the RADIUS authentication method before CHAP considers it a failure is 27 seconds. If you add more RADIUS servers to this configuration, they might not be contacted because the authentication process might be abandoned before these servers are tried.

The JUNOS software enforces a limit to the number of standing authentication server requests that the CHAP authentication can have at one time. Thus, an authentication server method—RADIUS, for example—may fail to authenticate a client when this limit is exceeded. In the above example, any authentication method following this method is tried. If it fails, the authentication sequence is reinitiated by the router until authentication succeeds and the link is brought up.

Tracing Access Processes

To trace access processes, you can specify options in the traceoptions statement at the [edit access] hierarchy level:

```
[edit access]
traceoptions {
    flag all;
    flag authentication;
    flag chap;
    flag configuration;
    flag radius;
}
```

You can specify the following access tracing flags:

- all—All tracing operations
- authentication—All authentication module handling
- chap—All CHAP messages and handling
- configuration—Reading of configuration
- radius—All RADIUS messages and handling

Configuring the Layer 2 Tunneling Protocol

For M7i and M10i routers, you can configure Layer 2 Tunneling Protocol (L2TP) tunneling security services on an Adaptive Services Physical Interface Card (PIC). The L2TP protocol allows the PPP to be tunneled within a network.

NOTE: For information about how to configure L2TP service, see the *JUNOS Services Interfaces Configuration Guide* and the *JUNOS Network Interfaces Configuration Guide*.

To configure L2TP, include the following statements at the [edit access] hierarchy level:

```
[edit access]
address-pool pool-name {
    address address-or-prefix;
    address-range low <lower-limit> high <upper-limit>;
}
group-profile profile-name {
    l2tp {
      interface-id interface-id;
      lcp-renegotiation;
      local-chap;
      maximum-sessions-per-tunnel number;
    ppp {
      framed-pool pool-id;
      idle-timeout seconds;
      interface-id interface-id;
      keepalive seconds;
      primary-dns primary-dns;
      primary-wins primary-wins;
      secondary-dns secondary-dns;
      secondary-wins secondary-wins;
    }
}
profile profile-name {
    authentication-order [ authentication-methods ];
    client client-name {
      chap-secret chap-secret;
      group-profile profile-name;
      l2tp {
        interface-id interface-id;
        lcp-renegotiation;
        local-chap;
        maximum-sessions-per-tunnel number;
        ppp-authentication (chap | pap);
        shared-secret shared-secret;
      }
      pap-password pap-password;
      ppp {
        framed-ip-address ip-address;
        framed-pool framed-pool;
        idle-timeout seconds;
        interface-id interface-id;
        keepalive seconds;
```

```
                primary-dns primary-dns;
                primary-wins primary-wins;
                secondary-dns secondary-dns;
                secondary-wins secondary-wins;
            }
            user-group-profile profile-name;
        }
    }
    radius-disconnect-port port-number {
        radius-disconnect {
            client-address {
                secret password;
            }
        }
    }
    radius-server server-address {
        accounting-port port-number;
        port port-number;
        retry attempts;
        secret password;
        source-address source-address;
        timeout seconds;
    }
```

This section includes the following topics:

- Minimum L2TP Configuration on page 657
- Configuring the Address Pool on page 658
- Configuring the Group Profile on page 659
- Configuring the Profile on page 662
- Example: Configuring L2TP on page 675
- Configuring RADIUS Authentication for L2TP on page 677
- Configuring the RADIUS Disconnect Server for L2TP on page 679

Minimum L2TP Configuration

To define L2TP, include at least the following statements at the [edit access] hierarchy level:

```
[edit access]
address-pool pool-name {
    address address-or-prefix;
    address-range low <lower-limit> high <upper-limit>;
}
profile profile-name {
    authentication-order [ authentication-methods ];
    client client-name {
      chap-secret chap-secret;
      l2tp {
        interface-id interface-id;
        maximum-sessions-per-tunnel number;
        ppp-authentication (chap | pap);
        shared-secret shared-secret;
      }
      pap-password pap-password;
      ppp {
        framed-ip-address ip-address;
        framed-pool framed-pool;
        interface-id interface-id;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
      }
    }
}
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    secret password;
}
```

Configuring the Address Pool

With an address pool, you configure an address or address range. When you define an address pool for a client, the L2TP network server (LNS) allocates IP addresses for clients from an address pool. If you do not want to use an address pool, you can specify an IP address by means of the framed-ip-address statement at the [edit access profile *profile-name* client *client-name* ppp] hierarchy level. For information about specifying an IP address, see "Configuring the PPP Properties for a Profile" on page 671.

NOTE: When an address pool is modified or deleted, all the sessions using that pool are deleted.

To define an address or a range of addresses, include the address-pool statement at the [edit access] hierarchy level:

```
[edit access]
address-pool pool-name;
```

pool-name is the name assigned to the address pool.

To configure an address, include the address statement at the [edit access address-pool *pool-name*] hierarchy level:

```
[edit access address-pool pool-name]
address address-or-prefix;
```

address-or-prefix is one address or a prefix value.

When you specify an address range, it cannot exceed 65,535 IP addresses. To configure the address range, include the address-range statement at the [edit access address-pool *pool-name*] hierarchy level:

```
[edit access address-pool pool-name]
address-range <low lower-limit> <high upper-limit>;
```

- low *lower-limit*—The lower limit of an address range.
- high *upper-limit*—The upper limit of an address range.

Configuring the Group Profile

You can optionally configure the group profile to define the PPP or L2TP attributes. Any client referencing the configured group profile inherits all the group profile attributes.

NOTE: The group-profile statement overrides the user-group-profile statement, which is configured at the [edit access profile *profile-name*] hierarchy level. The profile statement overrides the attributes configured at the [edit access group-profile *profile-name*] hierarchy level. For information about the user-group-profile statement, see "Applying a Configured PPP Group Profile to a Tunnel" on page 672.

To configure the group profile, include the group-profile statement at the [edit access] hierarchy level:

```
[edit access]
group-profile profile-name;
```

profile-name is the name assigned to the group profile.

To configure the L2TP properties for a group profile, include the following statements at the [edit access group-profile *profile-name*] hierarchy level:

```
[edit access group-profile profile-name]
l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
}
```

To configure the PPP properties for a group profile, include the following statements at the [edit access group-profile *profile-name*] hierarchy level:

```
[edit access group-profile profile-name]
ppp {
    framed-pool pool-id;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
```

This section describes how to configure the group profile:

- Configuring L2TP for a Group Profile on page 660
- Configuring the PPP Attributes for a Group Profile on page 660
- Example: Group Profile Configuration on page 661

Configuring L2TP for a Group Profile

To configure the L2TP for the group profile, include the following statements at the [edit access group-profile *profile-name* l2tp] hierarchy level:

```
[edit access group-profile profile-name l2tp]
interface-id interface-id;
lcp-renegotiation;
local-chap;
maximum-sessions-per-tunnel number;
```

interface-id is the identifier for the interface representing an L2TP session configured at the [edit interfaces *interface-name* unit *local-unit-number* dial-options] hierarchy level. For more information about the interface ID, see the *JUNOS Services Interfaces Configuration Guide*.

You can configure the LNS so that it renegotiates the link control protocol (LCP) with the PPP client (in the renegotiation statement). By default, the PPP client negotiates the LCP with the L2TP access concentrator (LAC). When you do this, the LNS discards the last sent and the last received LCP configuration request attribute value pairs (AVPs) from the LAC; for example, the LCP negotiated between the PPP client and the LAC.

You can configure the JUNOS software so that the LNS ignores proxy authentication AVPs from the LAC and reauthenticates the PPP client using a CHAP challenge (in the local-chap statement). When you do this, the LNS directly authenticates the PPP client. By default, the PPP client is not reauthenticated by the LNS.

number is the maximum number of sessions per L2TP tunnel.

Configuring the PPP Attributes for a Group Profile

To configure the PPP attributes for a group profile, include the following statements at the [edit access group-profile *profile-name* ppp] hierarchy level:

```
[edit access group-profile profile-name ppp]
framed-pool pool-id;
idle-timeout seconds;
interface-id interface-id;
keepalive seconds;
primary-dns primary-dns;
primary-wins primary-wins;
secondary-dns secondary-dns;
secondary-wins secondary-wins;
```

pool-id (in the framed-pool statement) is the name assigned to the address pool.

seconds (in the idle-timeout statement) is the number of seconds a user can remain idle before the session is terminated. By default, idle timeout is set to “0”. You can configure this to be a value in the range from 0 through 4,294,967,295**.**

interface-id (in the interface-id statement) is the identifier for the interface representing an L2TP session configured at the [edit interfaces *interface-name* unit *local-unit-number* dial-options] hierarchy level. For more information about the interface ID, see the *JUNOS Services Interfaces Configuration Guide*.

*second*s (in the keepalive statement) is the time period that must elapse before the JUNOS software checks the status of the PPP session by sending an echo request to the peer. For each session, JUNOS software sends out three keepalives at 10-second intervals and the session is close if there is no response. By default, the time to send a keepalive messages is set to 10 seconds. You configure this to be a value in the range from 0 through 32,767.

primary-dns (in the primary-dns statement) is an IP version 4 (IPv4) address.

secondary-dns (in the secondary-dns statement) is an IPv4 address.

primary-wins (in the primary-wins statement) is an IPv4 address.

secondary-wins (in the secondary-wins statement) is an IPv4 address.

Example: Group Profile Configuration

Configure an L2TP and PPP group profile:

```
[edit access]
group-profile westcoast_users {
    ppp {
      framed-pool customer_a;
      keepalive 15;
      primary-dns 192.120.65.1;
      secondary-dns 192.120.65.2;
      primary-wins 192.120.65.3;
      secondary-wins 192.120.65.4;
      interface-id west
    }
}
group-profile eastcoast_users {
    ppp {
      framed-pool customer_b;
      keepalive 15;
      primary-dns 192.120.65.5;
      secondary-dns 192.120.65.6;
      primary-wins 192.120.65.7;
      secondary-wins 192.120.65.8;
      interface-id east;
     }
}
group-profile westcoast_tunnel {
    l2tp {
      maximum-sessions-per-tunnel 100;
    }
}
group-profile east_tunnel {
    l2tp {
      maximum-sessions-per-tunnel 125;
    }
}
```

Configuring the Profile

You can configure multiple profiles. You can also configure multiple clients for each profile. To configure the profile, include the profile statement at the [edit access] hierarchy level:

```
[edit access]
profile profile-name;
```

profile-name is the name assigned to the profile.

NOTE: The group-profile statement overrides the user-group-profile statement, which is configured at the [edit access profile *profile-name*] hierarchy level. The profile statement overrides the attributes configured at the [edit access group-profile *profile-name*] hierarchy level. For information about the user-group-profile statement, see "Applying a Configured PPP Group Profile to a Tunnel" on page 672.

When you configure a profile, you can only configure L2TP or PPP parameters. You cannot configure both.

To configure the L2TP properties for a profile, include the following statements at the [edit access profile *profile-name*] hierarchy level:

```
[edit access profile profile-name]
authentication-order [ authentication-methods ];
client client-name {
      group-profile profile-name;
      l2tp {
         interface-id interface-id;
         lcp-renegotiation;
         local-chap;
         maximum-sessions-per-tunnel number;
         ppp-authentication (chap | pap);
         shared-secret shared-secret;
      }
}
user-group-profile profile-name;
```

To configure the PPP properties for a profile, include the following statements at the [edit access profile *profile-name*] hierarchy level:

```
[edit access profile profile-name]
authentication-order [ authentication-methods ];
client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    pap-password pap-password;
    ppp {
      framed-ip-address ip-address;
      framed-pool framed-pool;
      idle-timeout seconds;
      interface-id interface-id;
      keepalive seconds;
      primary-dns primary-dns;
      primary-wins primary-wins;
      secondary-dns secondary-dns;
      secondary-wins secondary-wins;
    }
}
```

NOTE: When you configure PPP properties for a profile, you typically configure the chap-secret statement or pap-password statement.

To configure the profile, do the following:

- Configuring the Authentication Order on page 663
- Configuring the Client on page 664

Configuring the Authentication Order

You can configure the order in which the JUNOS software tries different authentication methods when authenticating peers. For each access attempt, the software tries the authentication methods in order, from first to last.

To configure the authentication order, include the authentication-order statement at the [edit access profile *profile-name*] hierarchy level:

```
[edit access profile profile-name]
authentication-order [ authentication-methods ];
```

In *authentication-methods*, specify one or more of the following in the preferred order, from first tried to last tried:

- radius—Verify the client using RADIUS authentication services.
- password—Verify the client using the information configured at the [edit access profile *profile-name* client *client-name*] hierarchy level.

NOTE: When you configure the authentication methods for L2TP, only the first configured authentication method is used.

For L2TP, RADIUS authentication servers are configured at the [edit access radius-server] hierarchy level. For more information about configuring RADIUS authentication servers, see "Configuring RADIUS Authentication for L2TP" on page 677.

If you do not include the authentication-order statement, clients are verified by means of password authentication.

Configuring the Client

To configure the client, include the client statement at the [edit access profile *profile-name*] hierarchy level:

```
[edit access profile profile-name]
client client-name;
```

client-name is the peer identity.

For L2TP, you can optionally use the wildcard (*) to define a default tunnel client to authenticate multiple LACs with the same secret and L2TP attributes. If an LAC with a specific name is not defined in the configuration, the wildcard tunnel client authenticates it.

This section includes the following topics:

- Configuring the PPP Properties for a Profile on page 671
- Applying a Configured PPP Group Profile to a Tunnel on page 672
- Example: Applying a User Group Profile on page 673
- Example: Configuring the Profile on page 674

Example: Defining the Default Tunnel Client

Use the wildcard (*) to define a default tunnel client to authenticate multiple LACs with the same secret:

```
[edit access profile profile-name]
client * {
    l2tp {
      interface-id interface1;
      lcp-renegotiation;
      local-chap;
      maximum-sessions-per-tunnel 500;
      ppp-authentication chap;
      shared-secret "$1$dQYsZ$B5ojUeUjDsUo.yKwcCZ0";
    }
}
```

For any tunnel client, you can optionally use the user group profile to define default PPP attributes for all users coming in through a tunnel. The user group profile must define PPP attributes. If the user group profile is specified, all users (PPP sessions) use the PPP attributes specified in the user group profile. The PPP attributes specified in the local or RADIUS server take precedence over those specified in the user group profile.

You can optionally use a wildcard client to define a user group profile. When you do this, any client entering this tunnel uses the PPP attributes (defined user group profile attributes) as its default PPP attributes.

Example: Defining the User Group Profile

Use a wildcard client to define a user group profile:

```
[edit access profile profile]
client * {
    user-group-profile user-group-profile1;
}
```

For information about how to configure the user group profile, see “Applying a Configured PPP Group Profile to a Tunnel” on page 672.

Configuring the CHAP Secret

CHAP allows each end of a PPP link to authenticate its peer, as defined in RFC 1994. The authenticator sends its peer a randomly generated challenge that the peer must encrypt using a one-way hash; the peer must then respond with that encrypted result. The key to the hash is a secret known only to the authenticator and authenticated. When the response is received, the authenticator compares its calculated result with the peer's response. If they match, the peer is authenticated.

Each end of the link identifies itself to its peer by including its name in the CHAP challenge and response packets it sends to the peer. This name defaults to the local hostname, or you can explicitly set it using the local-name option. When a host receives a CHAP challenge or CHAP response packet on a particular interface, it uses the peer identity to look up the CHAP secret key to use. For more information about the local-name option, see the *JUNOS Network Interfaces Configuration Guide*.

NOTE: When you configure PPP properties for a profile, you typically configure the chap-secret statement or pap-password statement.

To configure CHAP, include the profile statement and specify a profile name at the [edit access] hierarchy level:

```
[edit access]
profile profile-name {
    client client-name chap-secret data;
}
```

Then reference the CHAP profile name at the [edit interfaces *interface-name* ppp-options chap] hierarchy level. For more information about how to reference CHAP, see the *JUNOS Network Interfaces Configuration Guide*.

You can configure multiple profiles. You can also configure multiple clients for each profile.

profile is the mapping between peer identifiers and CHAP secret keys. The identity of the peer contained in the CHAP challenge or response queries the profile for the secret key to use.

client is the peer identity.

chap-secret secret is the secret key associated with that peer.

Example: Configuring PPP CHAP

Configure the profile profile westcoast_bldg1 at the [edit access] hierarchy level, then reference it at the [edit interfaces] hierarchy level:

```
[edit]
access {
      profile westcoast_bldg1 {
         client cpe-1 chap-secret "$1$dQYsZ$B5ojUeUjDsUo.yKwcCZ0";
            # SECRET-DATA
         client cpe-2 chap-secret "$1$kdAsfaDAfkdjDsASxfafdKdFKJ";
            # SECRET-DATA
      }
}
```

Referencing the Group Profile

You can reference a configured group profile from the L2TP tunnel profile.

To reference the group profile configured at the [edit access group-profile *profile-name*] hierarchy level, include the group-profile statement at the [edit access profile *profile-name* client *client-name*] hierarchy level:

```
[edit access profile profile-name client client-name]
group-profile profile-name;
```

profile-name references a configured group profile from a PPP user profile.

Configuring L2TP Properties for a Profile

To define L2TP properties for a profile, include one or more of the following statements at the [edit access profile *profile-name* client *client-name* l2tp] hierarchy level:

NOTE: When you configure the profile, you can only configure L2TP or PPP parameters. You cannot configure both.

```
[edit access profile profile-name client client-name l2tp]
interface-id interface-id;
lcp-renegotiation;
local-chap;
maximum-sessions-per-tunnel number;
multilink {
      drop-timeout milliseconds;
      fragmentation-threshold bytes;
}
ppp-authentication (chap | pap);
shared-secret shared-secret;
```

interface-id (in the interface-id statement) is the identifier for the interface representing an L2TP session configured at the [edit interfaces *interface-name* unit *local-unit-number* dial-options] hierarchy level. For more information about the interface ID, see the *JUNOS Services Interfaces Configuration Guide*.

number (in the maximum-sessions-per-tunnel statement) is the maximum number of sessions for an L2TP tunnel.

shared-secret (in the shared-secret statement) is the shared secret for authenticating the peer.

You can specify PPP authentication (in the ppp-authentication statement). By default, the PPP authentication uses CHAP. You can configure this to use Password Authentication Protocol (PAP).

You can configure LNS so it renegotiates LCP with the PPP client (in the lcp-negotiation statement). By default, the PPP client negotiates the LCP with the LAC. When you do this, the LNS discards the last sent LCP configuration request and last received LCP configuration request AVPs from the LAC; for example, the LCP negotiated between the PPP client and LAC.

You can configure the JUNOS software so that the LNS ignores proxy authentication AVPs from the LAC and reauthenticates the PPP client using a CHAP challenge (in the local-chap statement). By default, the PPP client is not reauthenticated by the LNS. When you do this, the LNS directly authenticates the PPP client.

You can configure the PPP MP for L2TP if the PPP sessions that are coming into the LNS from the LAC have multilink PPP negotiated. When you do this, you join multilink bundles based on the endpoint discriminator (in the multilink statement).

- *milliseconds* (in the drop-timeout statement) specifies the number of milliseconds for the timeout that associated with the first fragment on the reassembly queue. If the timeout expires before all the fragments have been collected, the fragments at the beginning of the reassembly queue are dropped. If the drop timeout is not specified, the JUNOS software holds on to the fragments (fragments may still be dropped if the multilink reassembly algorithm determines that another fragment belonging to the packet on a reassembly queue has been lost).

NOTE: The drop timeout and fragmentation threshold for a bundled multilink might belong to different tunnels. The different tunnels might have different drop timeout and fragmentation thresholds. We recommend configuring group profiles instead of profiles when you have L2TP tunnels.

- *bytes* specifies the maximum size of a packet, in bytes (in the fragmentation-threshold statement). If a packet exceeds the fragmentation threshold, the JUNOS software fragments it into two or more multilink fragments.

Example: PPP MP for L2TP

Join multilink bundles based on the endpoint discriminator:

```
[edit access]
profile tunnel-profile {
  client remote-host {
    l2tp {
      multilink {
        drop-timeout 600;
        fragmentation-threshold 100;
      }
    }
  }
}
```

Example: L2TP Multilink PPP Support on Shared Interfaces

On M7i and M10i routers, L2TP multilink PPP sessions are supported on both dedicated and shared interfaces. This example shows how to configure many multilink bundles on a single ASP shared interface.

```
[edit}
interfaces {
    sp-1/3/0 {
        traceoptions {
            flag all;
        }
        unit 0 {
            family inet;
        }
        unit 20 {
            dial-options {
                l2tp-interface-id test;
                shared;
            }
            family inet;
        }
    }
}

access {
    profile t {
        client cholera {
            l2tp {
                interface-id test;
                multilink;
                shared-secret "$9$n8HX6A01RhlvL1R"; # SECRET-DATA
            }
        }
    }
    profile u {
        authentication-order radius;
    }
    radius-server {
        192.168.65.63 {
            port 1812;
            secret "$9$Vyb4ZHkPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
        }
    }
}
services {
    l2tp {
        tunnel-group 1 {
            tunnel-access-profile t;
            user-access-profile u;
            local-gateway {
                address 10.70.1.1;
            }
            service-interface sp-1/3/0;
        }
        traceoptions {
            flag all;
            debug-level packet-dump;
            filter {
                protocol l2tp;
                protocol ppp;
                protocol radius;
            }
        }
    }
}
```

Configuring the Password Authentication Protocol Password

When you configure PPP properties for a profile, you typically configure the chap-secret statement or pap-password statement. For information about how to configure the CHAP secret, see "Configuring the CHAP Secret" on page 666.

To configure the PAP password, include the pap-password statement at the [edit access profile *profile-name* client *client-name*] hierarchy level:

```
[edit access profile profile-name client client-name]
pap-password pap-password;
```

pap-password is the password for the PAP authentication protocol.

Example: PAP

```
[edit access]
profile sunnyvale_bldg_2 {
    client green {
        pap-password "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3";
        ppp {
            interface-id west;
        }
        group-profile sunnyvale_users;
    }
    client red {
        chap-secret "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3";
        group-profile sunnyvale_users;
    }
    authentication-order radius;
}
profile Sunnyvale_bldg_1_tunnel {
    client test {
        l2tp {
            shared-secret "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3";
            ppp-authentication pap;
        }
    }
}
```

Configuring the PPP Properties for a Profile

To define PPP properties for a profile, include one or more of the following statements at the [edit access profile *profile-name* client *client-name* ppp] hierarchy level. The properties defined in the profile take precedence over the values defined in the group profile.

```
[edit access profile profile-name client client-name ppp]
framed-ip-address ip-address;
framed-pool pool-id;
idle-timeout seconds;
interface-id interface-id;
keepalive seconds;
primary-dns primary-dns;
primary-wins primary-wins;
secondary-dns secondary-dns;
secondary-wins secondary-wins;
```

NOTE: When you configure a profile, you can only configure L2TP or PPP parameters. You cannot configure both.

ip-address (in the framed-ip-address statement) is the IPv4 prefix.

pool-id (in the framed-pool statement) is a configured address pool.

seconds (in the idle-timeout statement) is the number of seconds a user can remain idle before the session is terminated. By default, idle timeout is set to “0”. You can configure this to be a value in the range from 0 through 4,294,967,295.

interface-id (in the interface-id statement) is the identifier for the interface representing an L2TP session configured at the [edit interfaces *interface-name* unit *local-unit-number* dial-options] hierarchy level. For more information about the interface ID, see the *JUNOS Services Interfaces Configuration Guide*.

seconds (in the keepalive statement) is the time period that must elapse before the JUNOS software checks the status of the PPP session by sending an echo request to the peer. For each session, JUNOS software sends out three keepalives at 10-second intervals and the session is closed if there is no response.By default, the time to send a keepalive messages is set to 10 seconds. You configure this to be a value in the range from 0 through 32,767.

primary-dns (in the primary-dns statement) is an IPv4 address.

secondary-dns (in the secondary-dns statement) is an IPv4 address.

primary-wins (in the primary-wins statement) is an IPv4 address.

secondary-wins (in the secondary-wins statement) is an IPv4 address.

Applying a Configured PPP Group Profile to a Tunnel

You can optionally apply a configured PPP group profile to a tunnel. For any tunnel client, you can use the user-group-profile statement to define default PPP attributes for all users coming in through a tunnel. The user group profile must define PPP attributes. If the user group profile is specified, all users (PPP sessions) use the PPP attributes specified in the user group profile.

When a PPP client enters a tunnel, the JUNOS software first applies the PPP user group profile attributes and then any PPP attributes from the local or RADIUS server. The PPP attributes defined in the RADIUS or local server take precedence over the attributes defined in the user group profile.

To apply configured PPP attributes to a PPP client, include the user-group-profile statement at the [edit access profile *profile-name* client *client-name*] hierarchy level:

```
[edit access profile profile-name client client-name]
user-group-profile profile-name;
```

profile-name is a PPP group profile configured at the [edit access group-profile *profile-name*] hierarchy level. When a client enters this tunnel, it uses the user-group-profile attributes as the default attributes.

Example: Applying a User Group Profile

Apply a configured PPP group profile to a tunnel:

```
[edit access]
group-profile westcoast_users {
    ppp {
        idle-timeout 100;
    }
}
group-profile westcoast_default_configuration {
      ppp {
        framed-pool customer_b;
        idle-timeout 20;
        interface-id west;
        primary-dns 192.120.65.5;
        secondary-dns 192.120.65.6;
        primary-wins 192.120.65.7;
        secondary-wins 192.120.65.8;
       }
    }
}
profile westcoast_bldg_1_tunnel {
    client test {
      l2tp {
        interface-id west;
        shared-secret "$9$r3HKvLg4ZUDkX7JGjif5p0BIRS8LN";
          # SECRET-DATA
        maximum-sessions-per-tunnel 75;
        ppp-authentication chap;
      }
      user-group-profile westcoast_default_configuration; # Apply default PPP
    }                            # attributes for users coming through a tunnel
}
profile westcoast_bldg_1 {
    client white {
      chap-secret "$9$3s2690IeK8X7VKM7VwgaJn/Ctu1hclv87Ct87";
        # SECRET-DATA
      ppp {
        idle-timeout 22;
        primary-dns 192.120.65.9;
        framed-ip-address 12.12.12.12/32;
      }
      group-profile westcoast_users; # Reference the west_users group
    }                              # profile
}
```

Example: Configuring the Profile

Configure the profile:

```
[edit access]
profile westcoast_bldg_1 {
    client white {
      chap-secret "$9$3s2690IeK8X7VKM7VwgaJn/Ctu1hclv87Ct87";
        # SECRET-DATA
      ppp {
        idle-timeout 22;
        primary-dns 192.120.65.10;
        framed-ip-address 12.12.12.12/32;
      }
      group-profile westcoast_users;
    }
    client blue {
      chap-secret "$9$eq1KWxbwgZUHNdjqmTF3uO1Rhr-dsoJDNd";
        # SECRET-DATA
      group-profile sunnyvale_users;
    }
    authentication-order password;
}
profile westcoast_bldg_1_tunnel {
    client test {
      l2tp {
        shared-secret "$9$r3HKvLg4ZUDkX7JGjif5p0BIRS8LN";
          # SECRET-DATA
        maximum-sessions-per-tunnel 75;
        ppp-authentication chap;
      }
        group-profile westcoast_tunnel;
      }
    client production {
      l2tp {
        shared-secret "$9$R2QErv8X-goGylVwg4jiTz36/t0BEIeWFnRh
         rlXxbs2aJDHqf3nCP5";
          # SECRET-DATA
        ppp-authentication chap;
      }
      group-profile westcoast_tunnel;
    }
}
```

Example: Configuring L2TP

Configure L2TP:

```
[edit]
access {
    address-pool customer_a {
      address 1.1.1.1/32;
    }
    address-pool customer_b {
      address-range low 2.2.2.2 high 2.2.3.2;
    }
    group-profile westcoast_users {
      ppp {
        framed-pool customer_a;
        idle-timeout 15;
        primary-dns 192.120.65.1;
        secondary-dns 192.120.65.2;
        primary-wins 192.120.65.3;
        secondary-wins 192.120.65.4;
        interface-id west;
      }
    }
    group-profile eastcoast_users {
      ppp {
        framed-pool customer_b;
        idle-timeout 20;
        primary-dns 192.120.65.5;
        secondary-dns 192.120.65.6;
        primary-wins 192.120.65.7;
        secondary-wins 192.120.65.8;
        interface-id east;
       }
    }
    group-profile westcoast_tunnel {
      l2tp {
        maximum-sessions-per-tunnel 100;
      }
    }
    group-profile east_tunnel {
      l2tp {
        maximum-sessions-per-tunnel 125;
      }
    }
    profile westcoast_bldg_1 {
      client white {
        chap-secret "$9$3s2690IeK8X7VKM7VwgaJn/Ctu1hclv87Ct87";
          # SECRET-DATA
        ppp {
          idle-timeout 22;
          primary-dns 192.120.65.10;
          framed-ip-address 12.12.12.12/32;
        }
        group-profile westcoast_users;
      }
```

```
         client blue {
            chap-secret "$9$eq1KWxbwgZUHNdjqmTF3uO1Rhr-dsoJDNd";
               # SECRET-DATA
            group-profile sunnyvale_users;
         }
         authentication-order password;
      }
      profile west-coast_bldg_2 {
         client red {
            pap-password "$9$3s2690IeK8X7VKM8888Ctu1hclv87Ct87";
               # SECRET-DATA
            ppp {
               idle-timeout 22;
               primary-dns 192.120.65.11;
               framed-ip-address 12.12.12.12/32;
            }
            group-profile westcoast_users;
         }
      }
      profile westcoast_bldg_1_tunnel {
         client test {
            l2tp {
               shared-secret "$9$r3HKvLg4ZUDkX7JGjif5p0BIRS8LN";
               # SECRET-DATA
               maximum-sessions-per-tunnel 75;
               ppp-authentication chap;         #The default for PPP authentication
            }                                   # is CHAP
            group-profile westcoast_tunnel;
         }
         client production {
            l2tp {
               shared-secret "$9$R2QErv8X-goGylVwg4jiTz36/t0BEleWFnRh
                  rlXxbs2aJDHqf3nCP5"; # SECRET-DATA
               ppp-authentication chap;
            }
            group-profile westcoast_tunnel;
         }
      }
      profile westcoast_bldg_2_tunnel {
         client black {
            l2tp {
               shared-secret "$9$R2QErv8X-goGylVwg4jiTz36/t0BEleWFnRh
                  rlXxbs2aJDHqf3nCP5";
                     # SECRET-DATA
               ppp-authentication pap;
            }
            group-profile westcoast_tunnel;
         }
      }
}
```

Configuring RADIUS Authentication for L2TP

The LNS sends RADIUS authentication requests or accounting requests. Authentication requests are sent out to the authentication server port. Accounting requests are sent to the accounting port. To configure the RADIUS authentication for L2TP on an M10i or M7i routing platform, include following statements at the [edit access] hierarchy level:

```
[edit access]
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

NOTE: The RADIUS servers at the [edit access] hierarchy level are not used by the network access server process (NASD).

You can specify an accounting port number on which to contact the accounting server (in the accounting-port statement). Most RADIUS servers use port number 1813 (as specified in RFC 2866, *Radius Accounting*).

server-address specifies the address of the RADIUS authentication server (in the radius-server statement).

You can specify a port number on which to contact the RADIUS authentication server (in the port statement). Most RADIUS servers use port number 1812 (as specified in RFC 2865, *Remote Authentication Dial In User Service [RADIUS]*).

You must specify a password in the secret statement. Passwords can contain spaces. The secret used by the local router must match that used by the RADIUS authentication server.

Optionally, you can specify the amount of time that the local router waits to receive a response from a RADIUS server (in the timeout statement) and the number of times that the router attempts to contact a RADIUS authentication server (in the retry statement). By default, the router waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds. By default, the router retries connecting to the server three times. You can configure this to be a value in the range from 1 through 10 times.

In the source-address statement, specify a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 address configured on one of the router interfaces.

To configure multiple RADIUS servers, include multiple radius-server statements. For information about how to configure the RADIUS disconnect server for L2TP, see “Configuring the RADIUS Disconnect Server for L2TP” on page 679.

Example: RADIUS Authentication for L2TP

```
[edit access]
profile sunnyvale_bldg_2 {
    client green {
      chap-secret "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3";
      ppp {
        interface-id west;
      }
      group-profile sunnyvale_users;
    }
    client red {
      chap-secret "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3";
      group-profile sunnyvale_users;
    }
    authentication-order radius;
}
radius-server {
    192.168.65.213 {
        port 1812;
        accounting-port 1813;
        secret "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3"; # SECRET-DATA
    }
    192.168.65.223 {
      port 1812;
      accounting-port 1813;
      secret "$9$24gGiPfz6CuQFu1EyW8VwYgZUik.5z3"; # SECRET-DATA
    }
}
radius-disconnect-port 2500;
radius-disconnect {
    192.168.65.152 secret "$9$rtkl87ws4ZDkgokPT3tpEcylWL7-VY4a";
      # SECRET-DATA
    192.168.64.153 secret "$9$gB4UHf5F/A0z30Ihr8Lbs24GDHqmTFn";
      # SECRET-DATA
    192.168.64.157 secret "$9$Hk5FCA0IhruOrv87sYGDikfTFn/t0B";
      # SECRET-DATA
    192.168.64.173 secret "$9$Hk5FCA0IhruOrv87sYGDikfTFn/t0B";
      # SECRET-DATA
}
```

Configuring the RADIUS Disconnect Server for L2TP

To configure the RADIUS disconnect server to listen for disconnect requests from an administrator and process them, include the following statements at the [edit access] hierarchy level:

```
[edit access]
radius-disconnect-port port-number;
radius-disconnect {
   client-address {
      secret password;
   }
}
```

port-number is the server port to which the RADIUS client sends disconnect requests. The LNS, which accepts these disconnect requests, is the server. You can specify a port number on which to contact the RADIUS disconnect server. Most RADIUS servers use port number 1700.

NOTE: The JUNOS software accepts only disconnect requests from the client address configured at the [edit access radius-disconnect *client-address*] hierarchy level.

client-address is the host sending disconnect requests to the RADIUS server. The client address is a valid IP address configured on one of the router interfaces.

password authenticates the RADIUS client. Passwords can contain spaces. The secret used by the local router must match that used by the server.

For information about how to configure RADIUS authentication for L2TP, see “Configuring RADIUS Authentication for L2TP” on page 677.

Example: Configuring the RADIUS Disconnect Server

Configure the RADIUS disconnect server:

```
[edit access]
radius-disconnect-port 1700;
radius-disconnect {
      192.168.64.153 secret "$9$rtkl87ws4ZDkgokPT3tpEcylWL7-VY4a";
         # SECRET-DATA
      192.168.64.162 secret "$9$rtkl87ws4ZDkgokPT3tpEcylWL7-VY4a";
         # SECRET-DATA
}
```

Configuring RADIUS Authentication for an L2TP Profile

On an M10i or M7i routing platform, L2TP supports RADIUS authentication and accounting for users with one set of RADIUS servers at the [edit access] hierarchy level. You can also configure RADIUS authentication for each tunnel client or user profile.

To configure the RADIUS authentication for L2TP tunnel clients on an M10i or M7i routing platform, include the ppp-profile statement with the l2tp attributes for tunnel clients:

```
[edit access profile profile-name client client-name l2tp]
ppp-profile profile-name;
```

ppp-profile *profile-name* specifies the profile used to validate PPP session requests through L2TP tunnels. Clients of the referenced profile must have only PPP attributes. The referenced group profile must be defined.

To configure the RADIUS authentication for a profile, include the following statements at the [edit access profile *profile-name*] hierarchy level:

```
[edit access profile profile-name]
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

When a PPP user initiates a session and RADIUS authentication is configured for the user profile on the tunnel group, the following priority sequence is used to determine which RADIUS server is used for authentication and accounting:

- If the ppp-profile statement is configured under the tunnel client (LAC), the RADIUS servers configured under the specified ppp-profile are used.

- If RADIUS servers are configured under the user profile for the tunnel group, those servers will be used.

- If no RADIUS server is configured for the tunnel client (LAC) or user profile, then the RADIUS servers configured at the [edit access] hierarchy level are used.

Example: RADIUS Authentication for an L2TP Profile

Configure RADIUS authentication for an L2TP profile:

```
[edit access]
    profile t {
      client LAC_A {
        l2tp {
          ppp-profile u;
        }
      }
    }
    profile u {
      client client_1 {
        ppp {
        }
      }
    5.5.5.5 {
      port 3333;
      secret $9$dkafeqwrew;
      source-address 1.1.1.1;
      retry 3;
      timeout 3;
      }
    6.6.6.6 secret $9$fe3erqwrez;
    7.7.7.7 secret $9$f34929ftby;
    }
```

Configuring an Internet Key Exchange (IKE) Access Profile

An IKE access profile is used to negotiate IKE and IPSec security associations with dynamic peers. You can configure only one tunnel profile per service set for all dynamic peers. The configured preshared key in the profile is used for IKE authentication of all dynamic peers terminating in that service set.

The IKE tunnel profile specifies all the information needed to complete the IKE negotiation. Each protocol has its own statement hierarchy within the client statement to configure protocol-specific attribute value pairs, but only one client configuration is allowed for each profile. The following is the configuration hierarchy.

```
[edit access]
profile profile-name {
    client * {
      ike {
        allowed-proxy pair {
          remote remote-proxy-address local local-proxy-address;
        }
        pre-shared-key [ascii-text key-string] [hexadecimal key-string];
        interface-id interface-id;
      }
    }
}
```

NOTE: For dynamic peers, the JUNOS software supports only IKE main mode with the preshared key method of authentication. In this mode, IP address is used to identify a tunnel peer to get the preshared key information. The client value * (wildcard) means that configuration within this profile is valid for all dynamic peers terminating within the service set accessing this profile.

The following statements make up the IKE profile:

- allowed-proxy-pair—During phase 2 IKE negotiation, the remote peer supplies its network address (remote) and its peer's network address (local). Since multiple dynamic tunnels are authenticated through the same mechanism, this statement must include the list of possible combinations. If the dynamic peer does not present a valid combination, the phase 2 IKE negotiation fails.

  By default, remote 0.0.0.0/0 local 0.0.0.0/0 is used if no values are configured.

- pre-shared-key—Key used to authenticate the dynamic peer during IKE phase 1 negotiation. This key is known to both ends through an out-of-band secure mechanism. You can configure the value either in hexadecimal or ascii-text format. It is a mandatory value.

- interface-id—Interface identifier, a mandatory attribute used to derive the logical service interface information for the session.

For more information on configuring IPSec tunnels with dynamic peer security gateways, see the *JUNOS Feature Guide* and the *JUNOS Services Interfaces Configuration Guide*.

Chapter 30

Summary of Access Configuration Statements

The following sections explain each of the access configuration statements. The statements are organized alphabetically.

accounting-port

Syntax accounting-port *port-number*;

Hierarchy Level [edit access radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the port number on which to contact the accounting server.

Options *port-number*—The port number on which to contact the accounting server. Most RADIUS servers use port number 1813 (as specified in RFC 2866).

Usage Guidelines See “Configuring RADIUS Authentication for L2TP” on page 677.

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

address

Syntax address *address-or-prefix*;

Hierarchy Level [edit access address-pool *pool-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the IP address or prefix value for clients.

Options *address-or-prefix*—An address or prefix value.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Address Pool” on page 658.

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

address-pool

Syntax

```
address-pool pool-name {
      address address-or-prefix;
      address-range <low lower-limit> <high upper-limit>;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before JUNOS Release 7.4.

Description Allocate IP addresses for clients.

Options *pool-name*—Name assigned to an address pool.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Address Pool” on page 658.

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

address-range

Syntax

```
address-range <low lower-limit> <upper upper-limit>;
```

Hierarchy Level [edit access address-pool *pool-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the address range.

Options

- low *lower-limit*—The lower limit of an address range.
- high *upper-limit*—The upper limit of an address range.

Usage Guidelines See “Configuring the Address Pool” on page 658.

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

allowed-proxy pair

Syntax

```
allowed-proxy-pair {
    remote remote-proxy-address local local-proxy-address;
}
```

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced in JUNOS Release 7.4.

Description Specify the network address of the local and remote peer associated with an IKE access profile.

Default: remote 0.0.0.0/0 local 0.0.0.0/0

Options

- remote *remote-proxy-address*—Network address of the remote peer.
- local *local-proxy-address*—Network address of the local peer.

Usage Guidelines See "Configuring an Internet Key Exchange (IKE) Access Profile" on page 681.

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

authentication-order

Syntax authentication-order [*authentication-methods*];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Set the order in which the JUNOS software tries different authentication methods when verifying that a client can access the router. For each login attempt, the software tries the authentication methods in order, from first to last.

Options

- radius—Verify the client using RADIUS authentication services.
- password—Verify the client using the information configured at the [edit access profile *profile-name* client *client-name*] hierarchy level.

Usage Guidelines See "Example: CHAP Authentication with RADIUS" on page 651 and "Configuring the Authentication Order" on page 663.

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

chap-secret

Syntax chap-secret *chap-secret*;

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Options *chap-secret*—The secret key associated with a peer.

Usage Guidelines See "Configuring the CHAP Secret" on page 666.

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

client

Syntax

```
client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    ike {
      allowed-proxy pair {
        remote remote-proxy-address local local-proxy-address;
      }
      pre-shared-key [ascii-text key-string] [hexadecimal key-string];
      interface-id interface-id;
      }
    l2tp {
      interface-id interface-id;
      lcp-renegotiation;
      local-chap;
      maximum-sessions-per-tunnel number;
      multilink {
        drop-timeout milliseconds;
        fragmentation-threshold bytes;
      }
      ppp-authentication (chap | pap);
      ppp-profile profile-name;
      shared-secret shared-secret;
    }
    pap-password pap-password;
    ppp {
      framed-ip-address ip-address;
      framed-pool framed-pool;
      idle-timeout seconds;
      interface-id interface-id;
      keepalive seconds;
      primary-dns primary-dns;
      primary-wins primary-wins;
      secondary-dns secondary-dns;
      secondary-wins secondary-wins;
    }
    user-group-profile profile-name;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the peer identity.

Options *client-name*—A peer identity.

The other options are explained separately.

Usage Guidelines See “Configuring the Client” on page 664.

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

drop-timeout

Syntax drop-timeout *milliseconds*;

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp multilink]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the drop timeout for a multilink bundle.

Options *milliseconds*—Number of milliseconds for the timeout that is associated with the first fragment on the reassembly queue. If the timeout expires before all the fragments have been collected, the fragments at the beginning of the reassembly queue are dropped. If the drop timeout is not specified, the JUNOS software holds on to the fragments. (Fragments may still be dropped if the multilink reassembly algorithm determines that another fragment belonging to the packet on a reassembly queue has been lost.)

Usage Guidelines See “Configuring L2TP Properties for a Profile” on page 667.

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

fragmentation-threshold

Syntax fragmentation-threshold *bytes*;

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp multilink]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the fragmentation threshold for multilink bundle.

Options *bytes*—The maximum number of bytes in a packet. If a packet exceeds the fragmentation threshold, the JUNOS software fragments it into two or more multilink fragments.

Usage Guidelines See “Configuring L2TP Properties for a Profile” on page 667.

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

framed-ip-address

Syntax framed-ip-address *address*;

Hierarchy Level [edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify a framed IP address.

Options *address*—The IP version 4 (IPv4) prefix.

Usage Guidelines See “Configuring the PPP Properties for a Profile” on page 671.

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

framed-pool

Syntax framed-pool *framed-pool*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the address pool.

Options *framed-pool*—References a configured address pool.

Usage Guidelines See “Configuring the PPP Attributes for a Group Profile” on page 660 and “Configuring the PPP Properties for a Profile” on page 671.

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

group-profile

See the following sections:

- group-profile (Group Profile) on page 689
- group-profile (Profile) on page 690

group-profile (Group Profile)

Syntax

```
group-profile profile-name {
    l2tp {
      interface-id interface-id;
      lcp-renegotiation;
      local-chap;
      maximum-sessions-per-tunnel number;
    }
    ppp {
      framed-pool pool-id;
      idle-timeout seconds;
      interface-id interface-id;
      keepalive seconds;
      primary-dns primary-dns;
      primary-wins primary-wins;
      secondary-dns secondary-dns;
      secondary-wins secondary-wins;
    }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the group profile.

Options *profile-name*—Name assigned to the group profile.

The other options are explained separately.

Usage Guidelines See “Configuring the Group Profile” on page 659, “Configuring L2TP for a Group Profile” on page 660, “Configuring the PPP Attributes for a Group Profile” on page 660.

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

group-profile (Profile)

Syntax group-profile *profile-name*;

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Associate a group profile with a client.

Options *profile-name*—(Optional) Name assigned to the group profile.

Usage Guidelines See “Referencing the Group Profile” on page 667.

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

idle-timeout

Syntax idle-timeout *seconds*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the idle timeout for a user.

Options *seconds*—The number of seconds a user can remain idle before the session is terminated.
Range: 0 through 4,294,967,295 seconds
Default: 0

Usage Guidelines See “Configuring the PPP Attributes for a Group Profile” on page 660 and “Configuring the PPP Properties for a Profile” on page 671.

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

ike

Syntax

```
ike {
    allowed-proxy pair {
      remote remote-proxy-address local local-proxy-address;
    }
    pre-shared-key [ascii-text key-string] [hexadecimal key-string];
    interface-id <string-value>;
}
```

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure an IKE access profile.

The remaining statements are explained separately.

Usage Guidelines See "Configuring an Internet Key Exchange (IKE) Access Profile" on page 681.

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

interface-id

Syntax interface-id *interface-id*;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* l2tp],
[edit access profile *profile-name* client *client-name* ppp],
[edit access profile *profile-name* client *client-name* ike]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the interface identifier.

Options *interface-id*—The identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the [edit interfaces *interface-name* unit *local-unit-number* dial-options] hierarchy level. For more information about the interface ID, see the *JUNOS Services Interfaces Configuration Guide*.

Usage Guidelines See "Configuring L2TP for a Group Profile" on page 660, "Configuring the PPP Attributes for a Group Profile" on page 660, "Configuring L2TP Properties for a Profile" on page 667, "Configuring the PPP Properties for a Profile" on page 671, and "Configuring an Internet Key Exchange (IKE) Access Profile" on page 681.

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

keepalive

Syntax

```
keepalive seconds;
```

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the keepalive interval for an L2TP tunnel.

Options *seconds*—The time period that must elapse before the JUNOS software checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.
Range: 0 through 32,767 seconds

Usage Guidelines See "Configuring the PPP Attributes for a Group Profile" on page 660 and "Configuring the PPP Properties for a Profile" on page 671.

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

l2tp

See the following sections:

- l2tp (Group Profile) on page 692
- l2tp (Profile) on page 693

l2tp (Group Profile)

Syntax

```
l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
}
```

Hierarchy Level [edit access group-profile *profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the Layer 2 Tunneling Protocol for a group profile.

The remaining statements are explained separately.

Usage Guidelines See "Configuring L2TP for a Group Profile" on page 660.

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

l2tp (Profile)

Syntax

```
l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
    multilink {
      drop-timeout milliseconds;
      fragmentation-threshold bytes;
    }
    ppp-authentication (chap | pap);
    ppp-profile profile-name;
    shared-secret shared-secret;
}
```

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the L2TP properties for a profile.

The remaining statements are explained separately.

Usage Guidelines See "Configuring L2TP Properties for a Profile" on page 667.

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

lcp-renegotiation

Syntax lcp-renegotiation;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the L2TP network server (LNS) so it renegotiates the link control protocol (LCP) with the PPP client.

Usage Guidelines See "Configuring L2TP for a Group Profile" on page 660 and "Configuring L2TP Properties for a Profile" on page 667.

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

local-chap

Syntax local-chap;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the JUNOS software so that the LNS ignores proxy authentication attribute-value pairs (AVPs) from the L2TP access concentrator (LAC) and reauthenticates the PPP client using a Challenge Handshake Authentication Protocol (CHAP) challenge. When you do this, the LNS directly authenticates the PPP client.

Usage Guidelines See “Configuring L2TP for a Group Profile” on page 660 and “Configuring L2TP Properties for a Profile” on page 667.

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

maximum-sessions-per-tunnel

Syntax maximum-sessions-per-tunnel *number*;

Hierarchy Level [edit access group-profile l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the maximum sessions for a Layer 2 tunnel.

Options *number*—Maximum number of sessions for a Layer 2 tunnel.

Usage Guidelines See “Configuring L2TP for a Group Profile” on page 660 and “Configuring L2TP Properties for a Profile” on page 667.

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

multilink

Syntax

```
multilink {
    drop-timeout milliseconds;
    fragmentation-threshold bytes;
}
```

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the PPP MP for L2TP.

Options The statements are explained separately.

Usage Guidelines See "Configuring L2TP Properties for a Profile" on page 667.

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

pap-password

Syntax pap-password *password*;

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the Password Authentication Protocol (PAP) password.

Options *password*—PAP password.

Usage Guidelines See "Configuring the PPP Properties for a Profile" on page 671.

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

port

Syntax port *port-number*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the port number on which to contact the RADIUS server.

Options *port-number*—Port number on which to contact the RADIUS server.
Default: 1812 (as specified in RFC 2865)

Usage Guidelines See “Configuring RADIUS Authentication for L2TP” on page 677, “Configuring the RADIUS Disconnect Server for L2TP” on page 679, and “Example: CHAP Authentication with RADIUS” on page 651.

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

ppp

See the following sections:

- ppp (Group Profile) on page 696
- ppp (Profile) on page 697

ppp (Group Profile)

Syntax

```
ppp {
      framed-pool framed-pool;
      idle-timeout seconds;
      interface-id interface-id;
      keepalive seconds;
      primary-dns primary-dns;
      primary-wins primary-wins;
      secondary-dns secondary-dns;
      secondary-wins secondary-wins;
}
```

Hierarchy Level [edit access group-profile *profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure PPP properties for a group profile.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the PPP Attributes for a Group Profile” on page 660.

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

ppp (Profile)

Syntax

```
ppp {
    framed-ip-address address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
```

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure PPP properties for a client profile.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the PPP Properties for a Profile" on page 671.

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

ppp-authentication

Syntax ppp-authentication (chap | pap);

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure PPP authentication.

Options

- chap— The Challenge Handshake Authentication Protocol.
- pap— The Password Authentication Protocol.

Usage Guidelines See "Configuring L2TP Properties for a Profile" on page 667.

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

ppp-profile

Syntax ppp-profile *profile-name*;

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced in JUNOS Release 7.4.

Description Specify the profile used to validate PPP session requests through L2TP tunnels.

Options *profile-name*—Identifier for the PPP profile.

Usage Guidelines See "Configuring RADIUS Authentication for an L2TP Profile" on page 680.

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

pre-shared-key

Syntax pre-shared key [ascii-text *key-string*] [hexadecimal *key-string*];

Hierarchy Level [edit access group-profile *profile-name* client *client-name*]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure the key used to authenticate a dynamic peer during IKE phase 1 negotiation. Specify the key in either ASCII or hexadecimal format.

Options

- ascii-text *key-string*—Authentication key in ASCII format.
- hexadecimal *key-string*—Authentication key in hexadecimal format.

Usage Guidelines See "Configuring an Internet Key Exchange (IKE) Access Profile" on page 681.

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

primary-dns

Syntax primary-dns *primary-dns*;

Hierarchy Level [edit access group-profile *profile-name* client *client-name* ppp],
[edit access profile *profile-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the primary Domain Name System (DNS) server.

Options *primary-dns*—An IPv4 address.

Usage Guidelines See “Configuring the PPP Attributes for a Group Profile” on page 660 and “Configuring the PPP Properties for a Profile” on page 671.

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

primary-wins

Syntax primary-wins *primary-wins*;

Hierarchy Level [edit access group-profile *profile-name* client *client-name* ppp],
[edit access profile *profile-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the primary Windows Internet name server.

Options *primary-wins*—An IPv4 address.

Usage Guidelines See “Configuring the PPP Attributes for a Group Profile” on page 660 and “Configuring the PPP Properties for a Profile” on page 671.

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

profile

Syntax

```
profile profile-name {
    authentication-order [ authentication-methods ];
    client client-name {
        chap-secret chap-secret;
        group-profile profile-name;
        ike {
            allowed-proxy pair {
                remote remote-proxy-address local local-proxy-address;
            }
            pre-shared-key [ascii-text key-string] [hexadecimal key-string];
            interface-id <string-value>;
        }
        l2tp {
            interface-id interface-id;
            lcp-renegotiation;
            local-chap;
            maximum-sessions-per-tunnel number;
            multilink {
                drop-timeout milliseconds;
                fragmentation-threshold bytes;
            }
            ppp-authentication (chap | pap);
            ppp-profile profile-name;
            shared-secret shared-secret;
        }
        pap-password pap-password;
        ppp {
            framed-ip-address ip-address;
            framed-pool framed-pool;
            idle-timeout seconds;
            interface-id interface-id;
            keepalive seconds;
            primary-dns primary-dns;
            primary-wins primary-wins;
            secondary-dns secondary-dns;
            secondary-wins secondary-wins;
        }
        user-group-profile profile-name;
    }
    radius-server server-address {
        accounting-port port-number;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
    }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure PPP CHAP, or a profile and its L2TP or PPP properties.

Options *profile-name*—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the Challenge Handshake Authentication Protocol" on page 649, "Configuring the Profile" on page 662, "Configuring L2TP Properties for a Profile" on page 667, and "Configuring the PPP Properties for a Profile" on page 671.

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

radius-disconnect

Syntax

```
radius-disconnect {
    client-address {
      secret password;
   }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a disconnect server that listens on a configured User Datagram Protocol (UDP) port for disconnect messages from a configured client and processes these disconnect messages.

Options *client-address*—A valid IP address configured on one of the router interfaces.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the RADIUS Disconnect Server for L2TP" on page 679.

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

radius-disconnect-port

Syntax

```
radius-disconnect-port port-number;
```

Hierarchy Level [edit access]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify a port number on which to contact the RADIUS disconnect server. Most RADIUS servers use port number 1700.

Options *port-number*—The server port to which disconnect requests from the RADIUS client are sent. The LNS, which accepts these disconnect requests, is the server.

NOTE: The JUNOS software accepts only disconnect requests from the client address configured at the [edit access radius-disconnect client *client-address*] hierarchy level.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the RADIUS Disconnect Server for L2TP" on page 679.

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

radius-server

Syntax

```
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level [edit access],
[edit access profile *profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure RADIUS for L2TP or PPP

To configure multiple RADIUS servers, include multiple radius-server statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options *server-address*—Address of the RADIUS authentication server.

The remaining statements are explained separately.

Usage Guidelines See “Configuring RADIUS Authentication for L2TP” on page 677, “Configuring the Challenge Handshake Authentication Protocol” on page 649, and “Configuring RADIUS Authentication” on page 386.

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

retry

Syntax retry *attempts*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Number of times that the router attempts to contact a RADIUS authentication server.

Options *attempts*—Number of times to retry contacting a RADIUS server.
Range: 1 through 10
Default: 3

Usage Guidelines See “Configuring RADIUS Authentication for L2TP” on page 677 or “Example: CHAP Authentication with RADIUS” on page 651.

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

See Also timeout on page 706

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the routing instance used to send RADIUS packets to the RADIUS server.

Options *routing-instance-name*—Routing instance name.

Usage Guidelines See “Configuring the Challenge Handshake Authentication Protocol” on page 649 and “Configuring RADIUS Authentication” on page 386.

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

See also *JUNOS Network Interfaces Configuration Guide*

secondary-dns

Syntax secondary-dns *secondary-dns*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the secondary DNS server.

Options *secondary-dns*—An IPv4 address.

Usage Guidelines See "Configuring the PPP Attributes for a Group Profile" on page 660 and "Configuring the PPP Properties for a Profile" on page 671.

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

secondary-wins

Syntax secondary-wins *secondary-wins*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the secondary Windows Internet name server.

Options *secondary-wins*—An IPv4 address.

Usage Guidelines See "Configuring the PPP Attributes for a Group Profile" on page 660 and "Configuring the PPP Properties for a Profile" on page 671.

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

secret

Syntax secret *password* ;

Hierarchy Level [edit access radius-disconnect *client-address*],
[edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the password to use with the RADIUS server. The secret password used by the local router must match that used by the server.

Options *password*—Password to use; can include spaces.

Usage Guidelines See “Configuring RADIUS Authentication for L2TP” on page 677, “Configuring the RADIUS Disconnect Server for L2TP” on page 679, and “Example: CHAP Authentication with RADIUS” on page 651.

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

shared-secret

Syntax shared-secret *shared-secret*;

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the shared secret.

Options *shared-secret*—The shared secret key for authenticating the peer.

Usage Guidelines See “Configuring L2TP Properties for a Profile” on page 667.

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

source-address

Syntax source-address *source-address*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options *source-address*—A valid IPv4 address configured on one of the router interfaces.

Usage Guidelines See "Configuring RADIUS Authentication for L2TP" on page 677 or "Example: CHAP Authentication with RADIUS" on page 651.

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

timeout

Syntax timeout *seconds*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the amount of time that the local router waits to receive a response from a RADIUS server.

Options *seconds*—Amount of time to wait.
Range: 1 through 90 seconds
Default: 3 seconds

Usage Guidelines See "Configuring RADIUS Authentication for L2TP" on page 677 or "Example: CHAP Authentication with RADIUS" on page 651.

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

traceoptions

Syntax

```
traceoptions {
    file filename {
        files number;
        size maximum-file-size:
    }
    flag all;
    flag authentication;
    flag chap;
    flag configuration;
    flag kernel;
    flag radius;
    flag world-readable;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure access tracing options.

Options file *filename*—Save tracing information to a specified file.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

- all—All tracing operations
- authentication—All authentication module handling
- chap—All CHAP messages and handling
- configuration—Reading of configuration
- kernel–Send all configuration messages to the kernel
- radius—All RADIUS messages and handling
- world-readable–Allow any user to read the file

Usage Guidelines See “Tracing Access Processes” on page 654.

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

user-group-profile

Syntax user-group-profile *profile-name*;

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Apply a configured PPP group profile to PPP users.

Options *profile-name*—Name of a PPP group profile configured at the [edit access group-profile *profile-name*] hierarchy level.

Usage Guidelines See “Applying a Configured PPP Group Profile to a Tunnel” on page 672.

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

Part 7
Security Services

Chapter 31

Security Services Overview

The JUNOS software supports Internet Protocol Security (IPSec). This chapter discusses the following topics, which provide background information related to configuring IPSec:

For a list of IPSec- and IKE-supported standards, see "IPSec and IKE" on page 26.

IPSec Overview

IPSec architecture provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. In addition to IPSec, the JUNOS software also supports the Internet Key Exchange (IKE), which defines mechanisms for key generation and exchange, and manages security associations (SAs).

IPSec also defines a security association and key management framework that can be used with any network layer protocol. The SA specifies what protection policy to apply to traffic between two IP-layer entities. IPSec provides secure tunnels between two peers.

Security Associations

To use IPSec security services, you create SAs between hosts. An SA is a simplex connection that allows two hosts to communicate with each other securely by means of IPSec. There are two types of SAs: manual and dynamic.

- Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. Manual SAs statically define the Security Parameter Index (SPI) values, algorithms, and keys to be used, and require matching configurations on both ends of the tunnel. Each peer must have the same configured options for communication to take place.

- Dynamic SAs require additional configuration. With dynamic SAs, you configure IKE first and then the SA. IKE creates dynamic security associations; it negotiates SAs for IPSec. The IKE configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway. This connection is then used to dynamically agree upon keys and other data used by the dynamic IPSec SA. The IKE SA is negotiated first and then used to protect the negotiations that determine the dynamic IPSec SAs.

The JUNOS software implementation of IPSec supports two modes of security (transport and tunnel). For more information about transport and tunnel mode, see "Configuring IPSec Mode" on page 720.

IKE

IKE is a key management protocol that creates dynamic SAs; it negotiates SAs for IPSec. An IKE configuration defines the algorithms and keys used to establish a secure connection with a peer security gateway.

IKE does the following:

- Negotiates and manages IKE and IPSec parameters

- Authenticates secure key exchange

- Provides mutual peer authentication by means of shared secrets (not passwords) and public keys

- Provides identity protection (in main mode)

IKE occurs over two phases. In the first phase, it negotiates security attributes and establishes shared secrets to form the bidirectional IKE SA. In the second phase, inbound and outbound IPSec SAs are established. The IKE SA secures the exchanges in the second phase. IKE also generates keying material, provides Perfect Forward Secrecy, and exchanges identities

IPSec Requirements for JUNOS-FIPS

In a JUNOS-FIPS environment, hardware configurations with two Routing Engines must be configured to use IPSec and a private routing instance for all communications between the Routing Engines. IPSec communication between the Routing Engines and AS II FIPS PICs is also required. For more information about JUNOS-FIPS security, see the *JUNOS-FIPS Configuration Guide*

Chapter 32

Security Services Configuration Guidelines

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

```
[edit security]
certificates {
    cache-size bytes;
    cache-timeout-negative seconds;
    certification-authority ca-profile-name {
      ca-name ca-identity;
      crl file-name;
      encoding (binary | pem);
      enrollment-url url-name;
      file certificate-filename;
      ldap-url url-name;
    }
    enrollment-retry attempts;
    local certificate-filename {
      certificate-key-string;
      load-key-file key-filename;
    }
    maximum-certificates number;
    path-length certificate-path-length;
}
ike {
    proposal ike-proposal-name {
      authentication-algorithm (md5 | sha1);
      authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
      description description;
      dh-group (group1 | group2);
      encryption-algorithm (3des-cbc | des-cbc);
      lifetime-seconds seconds;
    }
    policy ike-peer-address {
      description description;
      encoding (binary | pem);
      identity identity-name;
      local-certificate certificate-filename;
      local-key-pair private-public-key-file;
      mode (aggressive | main);
      pre-shared-key (ascii-text key | hexadecimal key);
      proposals [ proposal-names ];
    }
}
ipsec {
    internal {
      security-association {
        manual {
          direction (bidirectional | inbound | outbound) {
            protocol esp;
            spi spi-value;
            encryption {
              algorithm 3des-cbc;
              key ascii-text ascii-text-string;
            }
          }
        }
      }
    }
    proposal ipsec-proposal-name {
      authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
      description description;
      encryption-algorithm (3des-cbc | des-cbc);
      lifetime-seconds seconds;
      protocol (ah | esp | bundle);
    }
    policy ipsec-policy-name {
      description description;
      perfect-forward-secrecy {
        keys (group1 | group2);
      }
      proposals [ proposal-names ];
    }
    security-association sa-name {
      description description;
      dynamic {
        ipsec-policy policy-name;
        replay-window-size (32 | 64);
      }
      manual {
        direction (JUNOS) (inbound | outbound | bi-directional) {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
          }
          auxiliary-spi auxiliary-spi;
          encryption {
            algorithm (des-cbc | 3des-cbc);
            key (ascii-text key | hexadecimal key);
          }
          protocol (ah | esp | bundle);
          spi spi-value;
        }
      }
      mode (tunnel | transport);
    }
}
traceoptions {
    file filename <files number> < size size>;
    flag all;
    flag database;
    flag general;
    flag ike;
    flag parse;
    flag policy-manager;
    flag routing-socket;
    flag timer;
}
```

NOTE: Most of the configuration statements do not have default values. If you do not specify an identifier for a statement that does not have a default value, you cannot commit the configuration.

For information about IP Security (IPSec) monitoring and troubleshooting, see the *JUNOS System Basics and Services Command Reference.*

This chapter describes the following tasks for configuring IPSec and Internet Key Exchange (IKE):

- Minimum Manual SA Configuration on page 718
- Minimum IKE Configuration on page 718
- Minimum Digital Certificates Configuration for IKE on page 719
- Configuring Security Associations on page 719
- Configuring an IKE Proposal (Dynamic SAs Only) on page 728
- Configuring an IKE Policy for Preshared Keys on page 731
- Configuring an IPSec Proposal on page 734
- Configuring the IPSec Policy on page 737
- Configuring Digital Certificates on page 739
- Configuring Trace Options on page 750
- Configuring the ES PIC on page 750
- Configuring Traffic on page 751
- Configuring an ES Tunnel Interface for a Layer 3 VPN on page 756

- Using JUNOScript SSL Service on page 757
- Configuring Internal IPSec for JUNOS-FIPS on page 758

Minimum Manual SA Configuration

To define a manual security association (SA) configuration, you must include at least the following statements at the [edit security ipsec] hierarchy level:

```
[edit security ipsec]
security-association sa-name {
      manual {
         direction (JUNOS) (inbound | outbound | bi-directional) {
            authentication {
               algorithm (hmac-md5-96 | hmac-sha1-96);
               key (ascii-text key | hexadecimal key);
            }
            encryption {
               algorithm (des-cbc | 3des-cbc);
               key (ascii-text key | hexadecimal key);
            }
            protocol (ah | esp | bundle);
            spi spi-value;
         }
      }
}
```

Minimum IKE Configuration

To define an IKE configuration, include at least the following statements at the [edit security] hierarchy level:

```
[edit security ike]
policy ike-peer-address {
      proposal [ ike-proposal-names ];
      pre-shared-key (ascii-text key | hexadecimal key);
}
```

Minimum Digital Certificates Configuration for IKE

To define a digital certificates configuration for IKE, include at least the following statements at the [edit security certificates] and [edit security ike] hierarchy levels:

```
[edit security]
certificates {
      certification-authority ca-profile-name {
         ca-name ca-identity;
         crl file-name;
         enrollment-url url-name;
         file certificate-filename;
         ldap-url url-name;
      }
}
ike {
      policy ike-peer-address {
         local-certificate certificate-filename;
         local-key-pair private-public-key-file;
         proposal [ ike-proposal-names ];
      }
      proposal ike-proposal-name {
         authentication-method rsa-signatures;
      }
}
```

Configuring Security Associations

To use IPSec security services, you create an SA between hosts. An SA is a simplex connection that allows two hosts to communicate with each other securely by means of IPSec. You can configure two types of SAs:

- Manual—Requires no negotiation; all values, including the keys, are static and specified in the configuration. As a result, each peer must have the same configured options for communication to take place. For information about how to configure a manual SA, see "Configuring Manual Security Associations" on page 722.
- Dynamic—Specifies proposals to be negotiated with the tunnel peer. The keys are generated as part of the negotiation and therefore do not need to be specified in the configuration. The dynamic SA includes one or more proposal statements, which allow you to prioritize a list of protocols and algorithms to be negotiated with the peer. For information about how to configure a dynamic SA, see "Configuring the ES PIC" on page 750.

NOTE: The JUNOS software does not perform a commit check when an SA name referenced in the Border Gateway Protocol (BGP) protocol section is not configured at the [edit security ipsec] hierarchy level.

We recommend that you configure no more than 512 dynamic security associations per ES Physical Interface Card (PIC).

To configure an SA for IPSec, include the security-association statement at the [edit security ipsec] hierarchy level:

```
[edit security ipsec]
security-association sa-name;
```

This section describes the following topics related to configuring security associations:

- Configuring the Description for an SA on page 720
- Configuring IPSec Mode on page 720
- Configuring Manual Security Associations on page 722
- Configuring Dynamic Security Associations on page 727

Configuring the Description for an SA

To specify a description for an IPSec SA, include the description statement at the [edit security ipsec security-association *sa-name*] hierarchy level:

```
[edit security ipsec security-association sa-name]
description description;
```

Configuring IPSec Mode

The JUNOS software implementation of IPSec supports two modes of security: transport and tunnel mode. By default, tunnel mode is enabled.

This section discusses the following topics:

- Configuring Transport Mode on page 720
- Configuring Tunnel Mode on page 721

Configuring Transport Mode

In transport mode, the data portion of the IP packet is encrypted, but the IP header is not. Transport mode can be used only when the communication endpoint and cryptographic endpoint are the same. Virtual private network (VPN) gateways that provide encryption and decryption services for protected hosts cannot use transport mode for protected VPN communications. SAs are manually configured and static values must be configured on both ends of the SA.

NOTE: When you use transport mode, the JUNOS software supports only BGP for manual SAs.

To configure IPSec security for transport mode, include the mode statement with the transport option at the [edit security ipsec security-association *sa-name*] hierarchy level:

```
[edit security ipsec security-association sa-name]
mode transport;
```

To apply tunnel mode, you configure manual SAs in transport mode and then reference the SA by name at the [edit protocols bgp] hierarchy level to protect a session with a given peer. For more information about how to reference the configured SA, see the *JUNOS Routing Protocols Configuration Guide*.

NOTE: You can configure BGP to peer over encrypted tunnels.

Configuring Tunnel Mode

You use tunnel mode when you use preshared keys with IKE to authenticate peers, or digital certificates with IKE to authenticate peers. In tunnel mode, encryption services are performed on an ES PIC.

When you use preshared keys, you manually configure a preshared key, which must match that of its peer. With digital certificates, each router is dynamically or manually enrolled with a certificate authority (CA). When a tunnel is established, the public keys used for IPsec are dynamically obtained through IKE and validated against the CA certificate. This avoids the manual configuration of keys on routers within the topology. Adding a new router to the topology does not require any security configuration changes to existing routers.

To configure the IPSec in tunnel mode, include the mode statement with the tunnel option at the [edit security ipsec security-association *sa-name*] hierarchy level:

```
[edit security ipsec security-association sa-name]
mode tunnel;
```

NOTE: Tunnel mode requires the ES PIC.

The JUNOS software supports only BGP in transport mode.

To enable tunnel mode, follow the steps in these sections:

1. Configuring an IKE Proposal (Dynamic SAs Only) on page 728

2. Configuring Security Associations on page 719

3. Configuring the ES PIC on page 750

4. Configuring Traffic on page 751

For more information about the ES PIC, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring Manual Security Associations

Manual SAs require no negotiation; all values, including the keys, are static and specified in the configuration. As a result, peers can communicate only when they all share the same configured options.

To configure the manual IPSec SA, include the manual statement at the [edit security ipsec security-association *sa-name*] hierarchy level:

```
[edit security ipsec security-association sa-name]
manual {
    direction (JUNOS) (inbound | outbound | bi-directional) {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      auxiliary-spi auxiliary-spi-value;
        encryption {
          algorithm (des-cbc | 3des-cbc);
          key (ascii-text key | hexadecimal key);
      }
      protocol (ah | esp | bundle);
      spi spi-value;
    }
}
```

The following sections describe how to configure a manual SA:

- Configuring the Processing Direction on page 722
- Configuring the Protocol for a Manual SA on page 723
- Configuring the Security Parameter Index on page 724
- Configuring the Auxiliary Security Parameter Index on page 724
- Configuring the Authentication Algorithm and Key on page 725
- Configuring the Encryption Algorithm and Key on page 726

Configuring the Processing Direction

The direction statement sets inbound and outbound IPSec processing. If you want to define different algorithms, keys, or security parameter index (SPI) values for each direction, you configure the inbound and outbound options. If you want the same attributes in both directions, use the bidirectional option.

To configure the direction of IPSec processing, include the direction statement and specify the direction at the [edit security ipsec security-association *sa-name* manual] hierarchy level:

```
[edit security ipsec security-association sa-name manual]
direction (JUNOS) (inbound | outbound | bidirectional);
```

For sample configurations, see the following sections:

- Example: Configuring Inbound and Outbound Processing on page 723
- Example: Configuring Bidirectional Processing on page 723

Example: Configuring Inbound and Outbound Processing

Define different algorithms, keys, and security parameter index values for each direction:

```
[edit security ipsec security-association sa-name]
manual {
    direction inbound {
        encryption {
            algorithm 3des-cbc;
            key ascii-text 2345678901234567890l234;
        }
        protocol esp;
        spi 16384;
    }
    direction outbound {
        encryption {
            algorithm 3des-cbc;
            key ascii-text 12345678901234567890abcd;
        }
        protocol esp;
        spi 24576;
    }
}
```

Example: Configuring Bidirectional Processing

Define the same algorithms, keys, and security parameter index values for each direction:

```
[edit security ipsec security-association sa-name manual]
direction bidirectional {
    authentication {
        algorithm hmac-md5-96;
        key ascii-text 123456789012abcd;
    }
    protocol ah;
    spi 20001;
}
```

Configuring the Protocol for a Manual SA

IPSec uses two protocols to protect IP traffic: Encapsulating Security Payload (ESP) and authentication header (AH). For transport mode SAs, both ESP and AH are supported. The AH protocol is used for strong authentication. The bundle option uses AH authentication and ESP encryption; it does not use ESP authentication because AH provides stronger authentication of IP packets.

NOTE: The AH protocol is supported only on M-series platforms.

To configure the IPSec protocol, include the protocol statement at the [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bidirectional) hierarchy level and specify the ah, bundle, or esp option:

```
[edit security ipsec security-association sa-name manual direction (inbound |
    outbound | bi-directional)]
protocol (ah | bundle | esp);
```

Configuring the Security Parameter Index

An SPI is an arbitrary value that uniquely identifies which SA to use at the receiving host. The sending host uses the SPI to identify and select which SA to use to secure every packet. The receiving host uses the SPI to identify and select the encryption algorithm and key used to decrypt packets.

NOTE: Each manual SA must have a unique SPI and protocol combination.

Use the auxiliary SPI when you configure the protocol statement to use the bundle option. For more information, see “Configuring the Auxiliary Security Parameter Index” on page 724.

To configure the SPI, include the spi statement and specify a value (256 through 16,639) at the [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bi-directional] hierarchy level:

```
[edit security ipsec security-association sa-name manual direction (inbound |
    outbound | bidirectional)]
spi spi-value;
```

Configuring the Auxiliary Security Parameter Index

When you configure the protocol statement to use the bundle option, the JUNOS software uses the auxiliary SPI for the ESP and the SPI for the AH.

NOTE: Each manual SA must have a unique SPI and protocol combination.

To configure the auxiliary SPI, include the auxiliary-spi statement at the [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bi-directional)] hierarchy level and set the value to an integer between 256 and 16,639:

```
[edit security ipsec security-association sa-name manual direction (inbound |
    outbound | bidirectional)]
auxiliary-spi auxiliary-spi-value;
```

Configuring the Authentication Algorithm and Key

To configure an authentication algorithm and key, include the authentication statement at the [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bi-directional)] hierarchy level:

```
[edit security ipsec security-association sa-name manual direction (inbound |
      outbound | bidirectional)]
authentication {
      algorithm (hmac-md5-96 | hmac-sha1-96);
      key (ascii-text key | hexadecimal key);
}
```

The algorithm can be one of the following:

- hmac-md5-96—Hash algorithm that authenticates packet data. It produces a 128-bit authenticator value and 96-bit digest.
- hmac-sha1-96—Hash algorithm that authenticates packet data. It produces a 160-bit authenticator value and a 96-bit digest.

The key can be one of the following:

- ascii-text *key*—ASCII text key. With the hmac-md5-96 option, the key contains 16 ASCII characters. With the hmac-sha1-96 option, the key contains 20 ASCII characters.
- hexadecimal *key*—Hexadecimal key. With the hmac-md5-96 option, the key contains 32 hexadecimal characters. With the hmac-sha1-96 option, the key contains 40 hexadecimal characters.

Configuring the Encryption Algorithm and Key

To configure IPSec encryption, include the encryption statement and specify an algorithm and key at the [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bi-directional)] hierarchy level:

```
[edit security ipsec security-association sa-name manual direction (inbound |
     outbound | bi-directional)]
encryption {
     algorithm (des-cbc | 3des-cbc);
     key (ascii-text key | hexadecimal key);
}
```

The algorithm can be one of the following:

- des-cbc—Encryption algorithm that has a block size of 8 bytes; its key size is 64 bits long.
- 3des-cbc—Encryption algorithm that has a block size of 24 bytes; its key size is 192 bits long.

NOTE: For a list of Data Encryption Standard (DES) encryption algorithm weak and semi-weak keys, see RFC 2409.

For 3des-cbc, we recommend that the first 8 bytes not be the same as the second 8 bytes, and that the second 8 bytes be the same as the third 8 bytes.

The key can be one of the following:

- ascii-text—ASCII text key. With the des-cbc option, the key contains 8 ASCII characters. With the 3des-cbc option, the key contains 24 ASCII characters.
- hexadecimal—Hexadecimal key. With the des-cbc option, the key contains 16 hexadecimal characters. With the 3des-cbc option, the key contains 48 hexadecimal characters.

NOTE: You cannot configure encryption when you use the AH protocol.

Configuring Dynamic Security Associations

You configure dynamic SAs with a set of proposals that are negotiated by the security gateways. The keys are generated as part of the negotiation and do not need to be specified in the configuration. The dynamic SA includes one or more proposals, which allow you to prioritize a list of protocols and algorithms to be negotiated with the peer.

To enable a dynamic SA, follow these steps:

1. Configure IKE proposals and IKE policies associated with these proposals.

2. Configure IPSec proposals and an IPSec policy associated with these proposals.

3. Associate an SA with an IPSec policy.

For more information about IKE policies and proposals, see "Configuring an IKE Policy for Preshared Keys" on page 731 and "Configuring an IKE Proposal (Dynamic SAs Only)" on page 728. For more information about IPSec policies and proposals, see "Configuring the IPSec Policy" on page 737.

NOTE: Dynamic tunnel SAs require an ES PIC.

To configure a dynamic SA, include the dynamic statement at the [edit security ipsec security-association *sa-name*] hierarchy level. Specify an IPSec policy name, and optionally, a 32- or 64-packet replay window size.

```
[edit security ipsec security-association sa-name]
dynamic {
    ipsec-policy policy-name;
    replay-window-size (32 | 64);
}
```

NOTE: If you want to establish a dynamic SA, the attributes in at least one configured IPSec and IKE proposal must match those of its peer.

The replay window is not used with manual SAs.

Configuring an IKE Proposal (Dynamic SAs Only)

Dynamic SAs require IKE configuration. The IKE configuration defines the algorithms and keys used to establish the secure IKE connection with the peer security gateway.

You can configure one or more IKE proposals. Each proposal is a list of IKE attributes to protect the IKE connection between the IKE host and its peer.

To configure an IKE proposal and define its properties, include the following statements at the [edit security ike] hierarchy level:

```
[edit security ike]
proposal ike-proposal-name {
    authentication-algorithm (md5 | sha1);
    authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2);
    encryption-algorithm (3des-cbc | des-cbc);
    lifetime-seconds seconds;
}
```

For information about associating an IKE proposal with an IKE policy, see "Associating Proposals with an IKE Policy" on page 732.

This section discusses the following topics:

- Configuring the Authentication Algorithm for an IKE Proposal on page 728
- Configuring the Authentication Method for an IKE Proposal on page 729
- Configuring the Description for an IKE Proposal on page 729
- Configuring the Diffie-Hellman Group for an IKE Proposal on page 729
- Configuring the Encryption Algorithm for an IKE Proposal on page 730
- Configuring the Lifetime for an IKE SA on page 730
- Example: Configuring an IKE Proposal on page 730

Configuring the Authentication Algorithm for an IKE Proposal

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

```
[edit security ike proposal ike-proposal-name]
authentication-algorithm (md5 | sha1);
```

The authentication algorithm can be one of the following:

- md5—Produces a 128-bit digest.
- sha1—Produces a 160-bit digest.

Configuring the Authentication Method for an IKE Proposal

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

```
[edit security ike proposal ike-proposal-name]
authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
```

The authentication method can be one of the following:

- dsa-signatures—Digital Signature Algorithm (DSA)
- pre-shared-keys—Preshared keys; a key derived from an out-of-band mechanism is used to authenticate an exchange
- rsa-signatures—A public key algorithm, which supports encryption and digital signatures

Configuring the Description for an IKE Proposal

To specify a description for an IKE proposal, include the description statement at the [edit security ike proposal *ike-proposal-name*] hierarchy level:

```
[edit security ike proposal ike-proposal-name]
description description;
```

Configuring the Diffie-Hellman Group for an IKE Proposal

Diffie-Hellman is a public-key cryptography scheme that allows two parties to establish a shared secret over an insecure communications channel. It is also used within IKE to establish session keys.

To configure an IKE Diffie-Hellman group, include the dh-group statement at the [edit security ike proposal *ike-proposal-name*] hierarchy level:

```
[edit security ike proposal ike-proposal-name]
dh-group (group1 | group2);
```

The group can be one of the following:

- group1—Specifies that IKE use the 768-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- group2—Specifies that IKE use the 1024-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.

group2 provides more security but requires more processing time.

Configuring the Encryption Algorithm for an IKE Proposal

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

```
[edit security ike proposal ike-proposal-name]
encryption-algorithm (3des-cbc | des-cbc);
```

The encryption algorithm can be one of the following:

- 3des-cbc—Encryption algorithm that has a key size of 24 bytes; its key size is 192 bits long.
- des-cbc—Encryption algorithm that has a key size of 8 bytes; its key size is 56 bits long.

Configuring the Lifetime for an IKE SA

The IKE lifetime sets the lifetime of an IKE SA. When the IKE SA expires, it is replaced by a new SA (and SPI) or is terminated. The default value IKE lifetime is 3600 seconds.

To configure the IKE lifetime, include the lifetime-seconds statement and specify the number of seconds (180 through 86,400) at the [edit security ike proposal *ike-proposal-name*] hierarchy level:

```
[edit security ike proposal ike-proposal-name]
lifetime-seconds seconds;
```

Example: Configuring an IKE Proposal

Configure an IKE proposal:

```
[edit security ike]
proposal ike-proposal {
    authentication-method pre-shared-keys;
    dh-group group1;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
}
```

Configuring an IKE Policy for Preshared Keys

An IKE policy defines a combination of security parameters (IKE proposals) to be used during IKE negotiation. It defines a peer address, the preshared key for the given peer, and the proposals needed for that connection. During the IKE negotiation, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

A match is made when both policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used. The configured preshared key must also match its peer.

You can create multiple, prioritized proposals at each peer to ensure that at least one proposal will match a remote peer's proposal.

First, you configure one or more IKE proposals; then you associate these proposals with an IKE policy. You can also prioritize a list of proposals used by IKE in the policy statement by listing the proposals you want to use, from first to last.

To configure an IKE policy, include the policy statement at the [edit security ike] hierarchy level and specify a peer address:

```
[edit security ike]
policy ike-peer-address;
```

NOTE: The IKE policy peer address must be an IPSec tunnel destination address.

This section discusses the following topics:

- Configuring the Description for an IKE Policy on page 732
- Configuring the Mode for an IKE Policy on page 732
- Configuring the Preshared Key for an IKE Policy on page 732
- Associating Proposals with an IKE Policy on page 732
- Example: Configuring an IKE Policy on page 733

Configuring the Description for an IKE Policy

To specify a description for an IKE policy, include the description statement at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
description description;
```

Configuring the Mode for an IKE Policy

IKE policy has two modes: aggressive and main. By default, main mode is enabled. Main mode uses six messages, in three exchanges, to establish the IKE SA. (These three steps are IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer.) Main mode also allows a peer to hide its identity.

Aggressive mode also establishes an authenticated IKE SA and keys. However, aggressive mode uses half the number of messages, has less negotiation power, and does not provide identity protection. The peer can use the aggressive or main mode to start IKE negotiation; the remote peer accepts the mode sent by the peer.

To configure IKE policy mode, include the mode statement and specify aggressive or main at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
mode (aggressive | main);
```

Configuring the Preshared Key for an IKE Policy

IKE policy preshared keys authenticate peers. You must manually configure a preshared key, which must match that of its peer. The preshared key can be an ASCII text (alphanumeric) key or a hexadecimal key.

To configure an IKE policy preshared key, include the pre-shared-key statement at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
pre-shared-key (ascii-text key | hexadecimal key);
```

Associating Proposals with an IKE Policy

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

To configure an IKE policy proposal, include the proposals statement at the [edit security ike policy *ike-peer-address*] hierarchy level and specify one or more proposal names:

```
[edit security ike policy ike-peer-address]
proposals [ proposal-names ];
```

For more information about configuring an individual proposal, see “Configuring an IKE Proposal (Dynamic SAs Only)” on page 728.

Example: Configuring an IKE Policy

Define two IKE policies: policy 10.1.1.2 and policy 10.1.1.1. Each policy is associated with proposal-1 and proposal-2.

```
[edit security]
ike {
    proposal proposal-1 {
      authentication-method pre-shared-keys;
      dh-group group1;
      authentication-algorithm sha1;
      encryption-algorithm 3des-cbc;
      lifetime-seconds 1000;
    }
    proposal proposal-2 {
      authentication-method pre-shared-keys;
      dh-group group2;
      authentication-algorithm md5;
      encryption-algorithm des-cbc;
      lifetime-seconds 10000;
    }
    proposal proposal-3 {
      authentication-method rsa-signatures;
      dh-group group2;
      authentication-algorithm md5;
      encryption-algorithm des-cbc;
      lifetime-seconds 10000;
    }
    policy 10.1.1.2 {
      mode main;
      proposals [ proposal-1 proposal-2 ];
      pre-shared-key ascii-text example-pre-shared-key;
    }
    policy 10.1.1.1 {
      local-certificate certificate-filename;
      local-key-pair private-public-key-file;
      mode aggressive;
      proposals [ proposal-2 proposal-3 ]
      pre-shared-key hexadecimal 0102030abbcd;
    }
}
```

NOTE: Updates to the current IKE proposal and policy configuration are not applied to the current IKE SA; updates are applied to new IKE SAs.

If you want the new updates to take immediate effect, you must clear the existing IKE security associations so that they will be reestablished with the changed configuration. For information about how to clear the current IKE security association, see the *JUNOS System Basics and Services Command Referencee.*

Configuring an IPSec Proposal

An IPSec proposal lists protocols and algorithms (security services) to be negotiated with the remote IPSec peer.

To configure an IPSec proposal and define its properties, include the following statements at the [edit security ipsec] hierarchy level:

```
[edit security ipsec]
proposal ipsec-proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    description description;
    encryption-algorithm (3des-cbc | des-cbc);
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
}
```

This section discusses the following topics:

- Configuring the Authentication Algorithm for an IPSec Proposal on page 735
- Configuring the Description for an IPSec Proposal on page 735
- Configuring the Encryption Algorithm for an IPSec Proposal on page 735
- Configuring the Lifetime for an IPSec SA on page 736
- Configuring the Protocol for a Dynamic IPSec SA on page 736

Configuring the Authentication Algorithm for an IPSec Proposal

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

```
[edit security ipsec proposal ipsec-proposal-name]
authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
```

The authentication algorithm can be one of the following:

- hmac-md5-96—Hash algorithm that authenticates packet data. It produces a 128-bit digest. Only 96 bits are used for authentication.

- hmac-sha1-96—Hash algorithm that authenticates packet data. It produces a 160-bit digest. Only 96 bits are used for authentication.

Configuring the Description for an IPSec Proposal

To specify a description for an IPSec proposal, include the description statement at the [edit security ipsec proposal *ipsec-proposal-name*] hierarchy level*:*

```
[edit security ike policy ipsec-proposal-name]
description description;
```

Configuring the Encryption Algorithm for an IPSec Proposal

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

```
[edit security ipsec proposal ipsec-proposal-name]
encryption-algorithm (3des-cbc | des-cbc);
```

The encryption algorithm can be one of the following:

- 3des-cbc—Encryption algorithm that has a block size of 24 bytes; its key size is 192 bits long.

- des-cbc—Encryption algorithm that has a block size of 8 bytes; its key size is 48 bits long.

NOTE: We recommend that you use the triple DES cipher block chaining (3DES-CBC) encryption algorithm.

Configuring the Lifetime for an IPSec SA

The IPSec lifetime option sets the lifetime of an IPSec SA. When the IPSec SA expires, it is replaced by a new SA (and SPI) or is terminated. A new SA has new authentication and encryption keys, and SPI; however, the algorithms may remain the same if the proposal is not changed. If you do not configure a lifetime and a lifetime is not sent by a responder, the lifetime is 28,800 seconds.

To configure the IPSec lifetime, include the lifetime-seconds statement and specify the number of seconds (180 through 86,400) at the [edit security ipsec proposal *ipsec-proposal-name*] hierarchy level:

```
[edit security ipsec proposal ipsec-proposal-name]
lifetime-seconds seconds;
```

NOTE: When a dynamic SA is created, two types of lifetimes are used: hard and soft. The hard lifetime specifies the lifetime of the SA. The soft lifetime, which is derived from the hard lifetime, informs the IPSec key management system that that the SA is about to expire. This allows the key management system to negotiate a new SA before the hard lifetime expires. When you specify the lifetime, you specify a hard lifetime.

Configuring the Protocol for a Dynamic IPSec SA

The protocol statement sets the protocol for a dynamic SA. The ESP protocol can support authentication, encryption, or both. The AH protocol is used for strong authentication. AH also authenticates the IP packet. The bundle option uses AH authentication and ESP encryption; it does not use ESP authentication because AH provides stronger authentication of IP packets.

To configure the protocol for a dynamic SA, include the protocol statement at the [edit security ipsec proposal *ipsec-proposal-name*] hierarchy level:

```
[edit security ipsec proposal ipsec-proposal-name]
protocol (ah | esp | bundle);
```

Configuring the IPSec Policy

An IPSec policy defines a combination of security parameters (IPSec proposals) used during IPSec negotiation. It defines Perfect Forward Secrecy (PFS) and the proposals needed for the connection. During the IPSec negotiation, IPSec looks for an IPSec proposal that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

A match is made when both policies from the two peers have a proposal that contains the same configured attributes. If the lifetimes are not identical, the shorter lifetime between the two policies (from the host and peer) is used.

You can create multiple, prioritized IPSec proposals at each peer to ensure that at least one proposal will match a remote peer's proposal.

First, you configure one or more IPSec proposals; then you associate these proposals with an IPSec policy. You can prioritize the proposals in the list by listing them in the order in which the IPSec policy uses them (first to last).

To configure an IPSec policy, include the policy statement at the [edit security ipsec] hierarchy level, specifying the policy name and one or more proposals you want to associate with this policy:

```
[edit security ipsec]
policy ipsec-policy-name {
    proposals [ proposal-names ];
}
```

This section discusses the following topics related to configuring an IPSec policy:

- Configuring Perfect Forward Secrecy on page 738
- Example: IPSec Policy Configuration on page 738

Configuring Perfect Forward Secrecy

PFS provides additional security by means of a Diffie-Hellman shared secret value. With PFS, if one key is compromised, previous and subsequent keys are secure because they are not derived from previous keys. This statement is optional.

To configure PFS, include the perfect-forward-secrecy statement and specify a Diffie-Hellman group at the [edit security ipsec policy *ipsec-policy-name*] hierarchy level:

```
[edit security ipsec policy ipsec-policy-name]
perfect-forward-secrecy {
    keys (group1 | group2);
}
```

The key can be one of the following:

- group1—Specifies that IKE use the 768-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.
- group2—Specifies that IKE use the 1024-bit Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange.

group2 provides more security than group1, but requires more processing time.

Example: IPSec Policy Configuration

Define an IPSec policy, dynamic policy-1, that is associated with two proposals (dynamic-1 and dynamic-2):

```
[edit security ipsec]
proposal dynamic-1 {
    protocol esp;
    authentication-algorithm hmac-md5-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 6000;
}
proposal dynamic-2 {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 6000;
}
policy dynamic-policy-1 {
    perfect-forward-secrecy {
      keys group1;
    }
    proposals [ dynamic-1 dynamic-2 ];
}
security-association dynamic-sa1 {
    dynamic {
      replay-window-size 64;
      ipsec-policy dynamic-policy-1;
    }
}
```

NOTE: Updates to the current IPSec proposal and policy configuration are not applied to the current IPSec SA; updates are applied to new IPSec SAs.

If you want the new updates to take immediate effect, you must clear the existing IPSec security associations so that they will be reestablished with the changed configuration. For information about how to clear the current IPSec security association, see the *JUNOS System Basics and Services Command Reference*.

Configuring Digital Certificates

To define the digital certificate configuration, include the following statements at the [edit security certificates] and [edit security ike] hierarchy levels:

```
[edit security]
certificates {
    cache-size bytes;
    cache-timeout-negative seconds;
    certification-authority ca-profile-name {
        ca-name ca-identity;
        crl file-name;
        encoding (binary | pem);
        enrollment-url url-name;
        file certificate-filename;
        ldap-url url-name;
    }
    enrollment-retry attempts;
    local certificate-filename {
        certificate-key-string;
        load-key-file key-filename;
    }
    maximum-certificates number;
    path-length certificate-path-length;
}
ike {
    policy ike-peer-address {
        description policy;
        encoding (binary | pem);
        identity identity-name;
        local-certificate certificate-filename;
        local-key-pair private-public-key-file;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals [ proposal-names ];
    }
}
```

For information about how to configure the description and mode statements, see “Configuring the Description for an IKE Policy” on page 732 and “Configuring the Mode for an IKE Policy” on page 732. For information about how to configure the IKE proposal, see “Associating Proposals with an IKE Policy” on page 732.

NOTE: For digital certificates, the JUNOS software supports only VeriSign.

Overview

Digital certificates provide a way of authenticating users through a trusted third party called a certificate authority (CA). The CA validates the identity of a certificate holder and “signs” the certificate to attest that it has not been forged or altered.

A certificate includes the following information:

- The distinguished name (DN) of the owner. A DN is a unique identifier and consists of a fully qualified name including not only the common name (CN) of the owner, the owner's organization, and other distinguishing information.
- The public key of the owner.
- The date on which the certificate was issued.
- The date on which the certificate expires.
- The distinguished name of the issuing CA.
- The digital signature of the issuing CA.

The additional information in a certificate allows recipients to decide whether to accept the certificate. The recipient can determine if the certificate is still valid based on the expiration date. The recipient can check whether the CA is trusted by the site based on the issuing CA.

With a certificate, a CA takes the owner's public key, signs that public key with the its own private key, and returns this to the owner as a certificate. The recipient can extract the certificate (containing the CA's signature) with the owner's public key. By using the CA's public key and the CA's signature on the extracted certificate, the recipient can validate the CA's signature and owner of the certificate.

When you use digital certificates, your first send in a request to obtain a certificate from your CA. You then configure digital certificates and a digital certificate IKE policy. Finally, you obtain a digitally signed certificate from a CA.

To use digital certificates for dynamic SAs, perform the tasks described in the following sections:

1. Obtaining a Certificate from a Certificate Authority on page 741
2. Configuring Digital Certificates on page 742
3. Configuring an IKE Policy for Digital Certificates on page 747
4. Obtaining a Signed Certificate from the CA on page 748

Obtaining a Certificate from a Certificate Authority

Certificate authorities manage certificate requests and issue certificates to participating IPSec network devices. When you create a certificate request, you need to provide the information about the owner of the certificate. The required information and its format vary across certificate authorities.

Certificates use names in the X.500 format, a directory access protocol that provides both read and update access**.** The entire name is called a DN (distinguished name). It consists of a set of components, which often includes a CN (common name), an organization (O), an organization unit (OU), a country (C), a locality (L), and so on.

NOTE: For the dynamic registration of digital certificates, the JUNOS software supports only the Simple Certificate Enrollment Protocols (SCEP).

Issue the following command to obtain a public key certificate from a CA. The results are saved in the specified file in the /var/etc/ikecert directory. The CA public key verifies certificates from remote peers.

```
user@host> request security certificate enroll filename filename ca-name
ca-name parameters parameters
```

For more information, see the following sections:

- Example: Obtaining a Certificate from a Certificate Authority on page 741
- Generating a Private and Public Key on page 742

Example: Obtaining a Certificate from a Certificate Authority

Specify a URL to the SCEP server and the name of the certification authority whose certificate you want: mycompany.com. filename 1 is name of the file that stores the result. The output, "Received CA certificate:" provides the signature for the certificate, which allows you to verify (offline) that the certificate is genuine.

```
user@host> request security certificate enroll filename ca_verisign ca-file
verisign ca-name xyzcompany url
http://pilotonsiteipsec.verisign.com/cgi-bin/pkiclient.exe
URL: http://pilotonsiteipsec.verisign.com/cgi-bin/pkiclient.exe CA name:
juniper.net CA file: verisign Encoding: binary
Certificate enrollment has started. To see the certificate enrollment status, check
the key management process (kmd) log file at /var/log/kmd. <––––––
```

NOTE: Each router is initially manually enrolled with a certificate authority.

Generating a Private and Public Key

To generate a private and public key, issue the following command:

```
user@host> request security key-pair name size key-size type ( rsa | dsa )
```

name specifies the filename in which to store the public and private keys.

key-size can be 512, 1024, 1596, or 2048 bytes. The default key size is 1024 bytes.

type can be rsa or dsa. The default is RSA.

NOTE: When you use SCEP, the JUNOS software only supports RSA.

Example: Generating a Key Pair

Generate a private and public key:

```
user@host> request security key-pair batt
Generated key pair, key size 1024, file batt Algorithm RSA
```

Configuring Digital Certificates

This section includes the following topics:

- Configuring the Certificate Authority Properties on page 743
- Configuring the Cache Size on page 745
- Configuring the Negative Cache on page 745
- Configuring the Number of Enrollment Retries on page 746
- Configuring the Maximum Number of Peer Certificates on page 746
- Configuring the Path Length for the Certificate Hierarchy on page 746

For information about the minimum digital certificate configuration for IKE, see “Minimum Digital Certificates Configuration for IKE” on page 719.

Configuring the Certificate Authority Properties

A CA is a trusted third-party organization that creates, enrolls, validates, and revokes digital certificates. The certificate authority guarantees a user's identity and issues public and private “keys” for message encryption and decryption (coding and decoding).

To configure a certificate authority and its properties, include the following statements at the [edit security certificates] hierarchy level:

```
[edit security certificates]
certification-authority ca-profile-name {
    ca-name ca-identity;
    crl file-name;
    encoding (binary | pem);
    enrollment-url url-name;
    file certificate-filename;
    ldap-url url-name;
```

ca-profile-name is the CA profile name.

This section discusses the following topics:

- Specifying the Certificate Authority Name on page 743
- Configuring the Certificate Revocation List on page 744
- Configuring the Type of Encoding Your CA Supports on page 744
- Specifying an Enrollment URL on page 744
- Specifying a File to Read the Digital Certificate on page 744
- Specifying an LDAP URL on page 745

Specifying the Certificate Authority Name

If you are enrolling with a CA using simple certificate enrollment protocols (SCEP), you need to specify the CA name (CA identity) that is used in the certificate request, in addition to the URL for the SCEP server.

To specify the name of the CA identity, include the ca-name statement at the [edit security certificates certification-authority *ca-profile-name*] hierarchy level:

```
[edit security certificates certification-authority ca-profile-name]
ca-name ca-identity;
```

ca-identity specifies the CA identity to use in the certificate request. It is typically the CA domain name.

Configuring the Certificate Revocation List

A certificate revocation list (CRL) contains a list of digital certificates that have been cancelled before their expiration date. When a participating peer uses a digital certificate, it checks the certificate signature and validity. It also acquires the most recently issued CRL and checks that the certificate serial number is not on that CRL.

To configure the CA certificate revocation list, include the crl statement and specify the file from which to read the CRL at the [edit security certificates certification-authority *ca-profile-name*] hierarchy level:

```
[edit security certificates certification-authority ca-profile-name]
crl file-name;
```

Configuring the Type of Encoding Your CA Supports

By default, encoding is set to binary. Encoding specifies the file format used for the local-certificate and local-key-pair statements. By default, the binary (distinguished encoding rules) format is enabled. Privacy-enhanced mail (PEM) is an ASCII base 64 encoded format. Check with your CA to determine which file formats it supports.

To configure the file format that your CA supports, include the encoding statement and specify a binary or PEM format at the [edit security certificates certification-authority *ca-profile-name*] hierarchy level*:*

```
[edit security certificates certification-authority ca-profile-name]
encoding (binary | pem);
```

Specifying an Enrollment URL

You specify the CA location where your router should send SCEP-based certificate enrollment requests. To specify the CA location by naming the CA URL, include the enrollment-url statement at the [edit security certificates certification-authority *ca-profile-name*] hierarchy level:

```
[edit security certificates certification-authority ca-profile-name]
enrollment-url url-name;
```

url-name is the CA location. The format is http://*CA_name,* where *CA_name* is the CA host DNS name or IP address.

Specifying a File to Read the Digital Certificate

To specify the file from which to read the digital certificate, include the file statement and specify the certificate filename at the [edit security certificates certification-authority *ca-profile-name*] hierarchy level:

```
[edit security certificates certification-authority ca-profile-name]
file certificate-filename;
```

Specifying an LDAP URL

If your CA stores its current CRL at its Lightweight Directory Access Protocol (LDAP) server, you can optionally check your CA CRL list before using a digital certificate. If the digital certificate appears on the CA CRL, your router cannot use it. To access your CA CRL, include the ldap-url statement at the [edit security certificates certification-authority *ca-profile-name*] hierarchy level:

```
[edit security certificates certification-authority ca-profile-name]
ldap-url url-name;
```

url-name is the certification authority LDAP server name. The format is ldap://*server_name,* where *server_name* is the CA host DNS name or IP address.

Configuring the Cache Size

By default, the cache size is 2 megabytes (MB). To configure total cache size for digital certificates, include the cache-size statement at the [edit security certificates] hierarchy level:

```
[edit security certificates]
cache-size bytes;
```

bytes is the cache size for digital certificates. The range can be from 64 through 4,294,967,295 bytes.

NOTE: We recommend that you limit your cache size to 4 MB.

Configuring the Negative Cache

Negative caching stores negative results and reduces the response time for negative answers. It also reduces the number of messages that are sent to the remote server. Maintaining a negative cache state allows the system to quickly return a failure condition when a lookup attempt is retried. Without a negative cache state, a retry would require waiting for the remote server to fail to respond, even though the system already “knows” that remote server is not responding.

By default, the negative cache is 20 seconds. To configure the negative cache, include the cache-timeout-negative statement at the [edit security certificates] hierarchy level:

```
[edit security certificates]
cache-timeout-negative seconds;
```

seconds is the amount of time for which a failed CA or router certificate is present in the negative cache. While searching for certificates with a matching CA identity (domain name for certificates or CA domain name and serial for CRLs), the negative cache is searched first. If an entry is found in the negative cache, the search fails immediately.

NOTE: Configuring a large negative cache value can make you susceptible to a denial-of-service (DoS) attack.

Configuring the Number of Enrollment Retries

By default, the number of enrollment retries is set to 0, an infinite number of retries**.** To specify how many times a router will resend a certificate request, include the enrollment-retry statement at the [edit security certificates] hierarchy level:

```
[edit security certificates]
enrollment-retry attempts;
```

attempts is the number of enrollment retries (0 through 100).

Configuring the Maximum Number of Peer Certificates

By default, the maximum number of peer certificates to be cached is 1024. To configure the maximum number of peer certificates to be cached, include the maximum-certificates statement at the [edit security certificates] hierarchy statement level:

```
[edit security certificates]
maximum-certificates number;
```

number is the maximum number of peer certificates to be cached. The range is from 64 through 4,294,967,295 peer certificates.

Configuring the Path Length for the Certificate Hierarchy

Certification authorities can issue certificates to other CAs. This creates a tree-like certification hierarchy. The highest trusted CA in the hierarchy is called the *trust anchor*. Sometimes the trust anchor is the root CA, which is usually signed by itself. In the hierarchy, every certificate is signed by the CA immediately above it. An exception is the root CA certificate, which is usually signed by the root CA itself. In general, a chain of multiple certificates may be needed, comprising a certificate of the public key owner (the end entity) signed by one CA, and zero or more additional certificates of CAs signed by other CAs. Such chains, called certification paths, are required because a public key user is only initialized with a limited number of assured CA public keys.

Path length refers to a path of certificates from one certificate to another certificate, based on the relationship of a CA and its "children". When you configure the path length statement, you specify the maximum depth of the hierarchy to validate a certificate from the trusted root CA certificate to the certificate in question. For more information about the certificate hierarchy, see RFC 3280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile.*

By default, the maximum certificate path length is set to 15. The root anchor is 1. To configure path length, include the path-length statement at the [edit security certificates] hierarchy level:

```
[edit security certificates]
path-length certificate-path-length;
```

certificate-path-length is the maximum number certificates for the certificate path length. The range is from 2 through 15 certificates.

Configuring an IKE Policy for Digital Certificates

An IKE policy for digital certificates defines a combination of security parameters (IKE proposals) to be used during IKE negotiation. It defines a peer address and the proposals needed for that connection. During the IKE negotiation, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation sends all its policies to the remote peer, and the remote peer tries to find a match.

To configure an IKE policy for digital certificates, include the following statements at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike]
policy ike-peer-address {
      encoding (binary | pem);
      identity identity-name;
      local-certificate certificate-filename;
      local-key-pair private-public-key-file ;
}
```

This section contains the following topics:

- Configuring the Type of Encoding Your CA Supports on page 747
- Configuring the Identity to Define the Remote Certificate Name on page 747
- Specifying the Certificate Filename on page 748
- Specifying the Private and Public Key File on page 748

Configuring the Type of Encoding Your CA Supports

By default, the encoding is set to binary. Encoding specifies the file format used for the local-certificate and local-key-pair statements. By default, the binary (distinguished encoding rules) format is enabled. PEM is an ASCII base 64 encoded format. Check with your CA to determine which file formats it supports.

To configure the file format that your CA supports, include the encoding statement and specify a binary or PEM format at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
encoding (binary | pem);
```

Configuring the Identity to Define the Remote Certificate Name

To define the remote certificate name, include the identity statement at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
identity identity-name;
```

identity-name defines the identity of the remote certificate name if the identity cannot be learned through IKE (ID payload or IP address).

Specifying the Certificate Filename

To configure the certificate filename from which to read the local certificate, include the local-certificate statement at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
local-certificate certificate-filename;
```

certificate-filename specifies the file from which to read the local certificate.

Specifying the Private and Public Key File

To specify the filename from which to read the public and private key, include the local key-pair statement at the [edit security ike policy *ike-peer-address*] hierarchy level:

```
[edit security ike policy ike-peer-address]
local-key-pair private-public-key-file;
```

private-public-key-file specifies the file from which to read the pair key.

Obtaining a Signed Certificate from the CA

To obtain signed certificate from the CA, issue the following command:

```
user@host> request security certificate enroll filename filename subject
c=us,o=x alternative-subject certificate-ip-address certification-authority
certificate-authority key-file key-file-name domain-name domain-name
```

The results are saved in a specified file to the /var/etc/ikecert directory.

Example: Obtaining a Signed Certificate

Obtain a CA signed certificate by referencing the configured certification-authority statement “local”. This statement is referenced by the request security certificate enroll filename m subject c=us,0=x alternative subject 1.1.1.1 certification-authority command.

```
[edit]
security {
    certificates {
        certification-authority local {
            ca-name xyz.company.com;
            file l;
            enrollment-url "http://www.xyzcompany.com";
        }
    }
}
```

To obtain a signed certificate from the CA, issue the following command:

```
user@host> request security certificate enroll filename l subject c=uk,o=london
alternative-subject 10.50.1.4 certification-authority verisign key-file host-1.prv
domain-name host.xyzcompany.com
CA name: xyz.company.com CA file: ca_verisign
local pub/private key pair: host.prv
subject: c=uk,o=london domain name: host.juniper.net
alternative subject: 10.50.1.4
Encoding: binary
Certificate enrollment has started. To see the certificate enrollment status, check
the key management process (kmd) log file at /var/log/kmd. <––––––
```

For information about how to use the operational mode commands to obtain a signed certificate, see the *JUNOS System Basics and Services Command Reference*.

Another way to obtain a signed certificate from the CA is to reference the configured statements such as the URL, CA name, and CA certificate file by means of the certification-authority statement:

```
user@host> request security certificate enroll filename m subject c=us,o=x
alternative-subject 1.1.1.1 certification-authority local key-file y domain-name
abc.company.com
```

Configuring Trace Options

To configure security traceoptions, you can specify options in the traceoptions statement at the [edit security] hierarchy level:

```
[edit security]
traceoptions {
    file filename <files number> <size size>;
    flag all;
    flag database;
    flag general;
    flag ike;
    flag parse;
    flag policy-manager;
    flag routing-socket;
    flag timer;
}
```

The output of the security tracing options is placed in the /var/log/kmd file.

You can specify one or more of the following security tracing flags:

- all—Trace all security events
- database—Trace database events
- general—Trace general events
- ike—Trace IKE module processing
- parse—Trace configuration processing
- policy-manager—Trace policy manager processing
- routing-socket—Trace routing socket messages
- timer—Trace internal timer events

Configuring the ES PIC

Configuring the ES PIC associates the configured SA with a logical interface. This configuration defines the tunnel itself (logical subunit, tunnel addresses, maximum transmission unit [MTU], optional interface addresses, and the name of the SA to apply to traffic).

The addresses configured as the tunnel source and destination are the addresses in the outer IP header of the tunnel.

NOTE: The tunnel source address must be configured locally on the router, and the tunnel destination address must be a valid address for the security gateway terminating the tunnel.

The M5, M10, M20, and M40 routers support the ES PIC.

You can also configure IPSec on the Adaptive Services PIC. For information about how to configure IPSec on a Adaptive Services PIC, see the *JUNOS Services Interfaces Configuration Guide*.

The SA must be a valid tunnel-mode SA. The interface address and destination address listed are optional. The destination address allows the user to configure a static route to encrypt traffic. If a static route uses that destination address as the next hop, traffic is forwarded through the portion of the tunnel in which encryption occurs. For more information about the ES PIC, see the *JUNOS Services Interfaces Configuration Guide*.

Example: Configuring the ES PIC

Configure an IPSec tunnel as a logical interface on the ES PIC. The logical interface specifies the tunnel through which the encrypted traffic travels. The ipsec-sa statement associates the security profile with the interface.

```
[edit interfaces]
es-0/0/0 {
    unit 0 {
      tunnel {
        source tunnel 10.5.5.5;                # tunnel source address
        destination 10.6.6.6;                  # tunnel destination address
      }
      family inet {
        ipsec-sa ipsec-sa; # name of security association to apply to packet
        address 10.1.1.8/32 { # local interface address inside local VPN
          destination 10.2.2.254; # destination address inside remote VPN
        }
      }
    }
}
```

Configuring Traffic

This section contains the following topics:

- Traffic Overview on page 752
- Example: Configuring Outbound Traffic Filter on page 754
- Example: Applying Outbound Traffic Filter on page 755
- Example: Configuring Inbound Traffic Filter for Policy Check on page 755
- Example: Applying Inbound Traffic Filter to ES PIC for Policy Check on page 756

Traffic Overview

Traffic configuration defines the traffic that must flow through the tunnel. You configure outbound and inbound firewall filters, which identify and direct traffic to be encrypted and confirm that decrypted traffic parameters match those defined for the given tunnel. The outbound filter is applied to the LAN or WAN interface for the incoming traffic you want to encrypt off of that LAN or WAN. The inbound filter is applied to the ES PIC to check the policy for traffic coming in from the remote host. Because of the complexity of configuring a router to forward packets, no automatic checking is done to ensure that the configuration is correct. Make sure that you configure the router very carefully.

NOTE: The valid firewall filters statements for IPSec are destination-port, source-port, protocol, destination-address, and source-address.

In Figure 11, Gateway A protects the network 10.1.1.0/24, and Gateway B protects the network 10.2.2.0/24. The gateways are connected by an IPSec tunnel. For more information about firewalls, see the *JUNOS Policy Framework Configuration Guide.*

Figure 11: Example: IPSec Tunnel Connecting Security Gateways

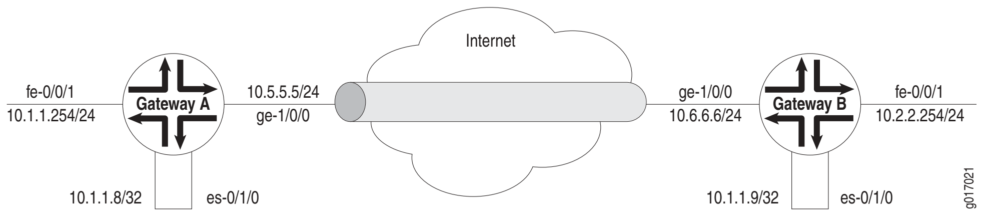

The SA and ES interface for security Gateway A are configured as follows:

```
[edit security ipsec]
security-association manual-sa1 {
    manual {
      direction bidirectional {
        protocol esp;
        spi 2312;
        authentication {
          algorithm hmac-md5-96;
          key ascii-text 1234123412341234;
        }
        encryption {
          algorithm 3des-cbc;
          key ascii-text 123456789009876543211234;
        }
      }
    }
}
```

```
[edit interfaces es-0/1/0]
    unit 0 {
      tunnel {
        source 10.5.5.5;
        destination 10.6.6.6;
      }
      family inet {
        ipsec-sa manual-sa1;
        address 10.1.1.8/32 {
          destination 10.1.1.9;
        }
      }
    }
}
```

The SA and ES interface for security Gateway B are configured as follows:

```
[edit security ipsec]
security-association manual-sa1 {
    manual {
      direction bidirectional {
        protocol esp;
        spi 2312;
        authentication {
          algorithm hmac-md5-96;
          key ascii-text 1234123412341234;
        }
        encryption {
          algorithm 3des-cbc;
          key ascii-text 123456789009876543211234;
        }
      }
    }
}

[edit interfaces es-0/1/0]
    unit 0 {
      tunnel {
        source 10.6.6.6;
        destination 10.5.5.5;
      }
      family inet {
        ipsec-sa manual-sa1;
        address 10.1.1.9/32; {
          destination 10.1.1.8;
        }
      }
    }
}
```

Example: Configuring Outbound Traffic Filter

Firewall filters for outbound traffic direct the traffic through the desired IPSec tunnel and ensure that the tunneled traffic goes out the appropriate interface (see Figure 11 on page 752). Here, an outbound firewall filter is created on security Gateway A; it identifies the traffic to be encrypted and adds it to the input side of the interface that carries the internal VPN traffic:

```
[edit firewall]
filter ipsec-encrypt-policy-filter {
    term term1 {
        from {
            source-address {          # local network
                10.1.1.0/24;
            }
            destination-address {    # remote network
                10.2.2.0/24;
            }
        }
        then ipsec-sa manual-sa1;     # apply SA name to packet
    }
    term default {
        then accept;
    }
}
```

NOTE: The source address, port, and protocol on the outbound traffic filter must match the destination address, port, and protocol on the inbound traffic filter. The destination address, port, and protocol on the outbound traffic filter must match the source address, port, and protocol on the inbound traffic filter.

Example: Applying Outbound Traffic Filter

After you have configured the outbound firewall filter, you apply it:

```
[edit interfaces]
fe-0/0/1 {
    unit 0 {
      family inet {
        filter {
          input ipsec-encrypt-policy-filter;
        }
        address 10.1.1.254/24;
      }
    }
}
```

The outbound filter is applied on the Fast Ethernet interface at the [edit interfaces fe-0/0/1 unit 0 family inet] hierarchy level. Any packet matching the IPSec action term (term 1) on the input filter (ipsec-encrypt-policy-filter), configured on the Fast Ethernet interface, is directed to the ES PIC interface at the [edit interfaces es-0/1/0 unit 0 family inet] hierarchy level. If a packet arrives from the source address 10.1.1.0/24 and goes to the destination address 10.2.2.0/24, the Packet Forwarding Engine directs the packet to the ES PIC interface, which is configured with the manual-sa1 SA. The ES PIC receives the packet, applies the manual-sa1 SA, and sends the packet through the tunnel.

The router must have a route to the tunnel endpoint; add a static route if necessary.

Example: Configuring Inbound Traffic Filter for Policy Check

Here, an inbound firewall filter, which performs the final IPSec policy check, is created on security gateway A. This check ensures that only packets that match the traffic configured for this tunnel are accepted.

```
filter ipsec-decrypt-policy-filter {
    term term1 {                          # perform policy check
      from {
        source-address {                  # remote network
          10.2.2.0/24;
        }
        destination-address {              # local network
          10.1.1.0/24;
        }
      }
      then accept;
    }
}
```

Example: Applying Inbound Traffic Filter to ES PIC for Policy Check

After you create the inbound firewall filter, apply it to the ES PIC. Here, the inbound firewall filter (ipsec-decrypt-policy-filter) is applied on the decrypted packet to perform the final policy check. The IPSec manual-sa1 SA is referenced at the [edit interfaces es-1/2/0 unit 0 family inet] hierarchy level and decrypts the incoming packet.

```
[edit interfaces]
es-1/2/0 {
    unit 0 {
      tunnel {
        source 10.5.5.5;                    # tunnel source address
        destination 10.6.6.6;               # tunnel destination address
      }
      family inet {
        filter {
          input ipsec-decrypt-policy-filter;
        }
        ipsec-sa manual-sa1; # SA name applied to packet
        address 10.1.1.8/32 { # local interface address inside local VPN
          destination 10.2.2.254; # destination address inside remote VPN
        }
      }
    }
}
```

The Packet Forwarding Engine directs IPSec packets to the ES PIC. It uses the packet's SPI, protocol, and destination address to look up the SA configured on one of the ES interfaces. The IPSec manual-sa1 SA is referenced at the [edit interfaces es-1/2/0 unit 0 family inet] hierarchy level and is used to decrypt the incoming packet. When the packets are processed (decrypted, authenticated, or both), the input firewall filter (ipsec-decrypt-policy-filter) is applied on the decrypted packet to perform the final policy check. Term1 defines the decrypted (and verified) traffic and performs the required policy check.

NOTE: The inbound traffic filter is applied after the ES PIC has processed the packet, so the decrypted traffic is defined as any traffic that the remote gateway is encrypting and sending to this router. IKE uses this filter to determine the policy required for a tunnel. This policy is used during the negotiation with the remote gateway to find the matching SA configuration.

Configuring an ES Tunnel Interface for a Layer 3 VPN

To configure an ES tunnel interface for a Layer 3 VPN, you need to configure an ES tunnel interface on the provider edge (PE) router and on the customer edge (CE) router. You also need to configure IPSec on the PE and CE routers. For more information about configuring an ES tunnel for a Layer 3 VPN, see the *JUNOS VPNs Configuration Guide.*

Using JUNOScript SSL Service

This section contains the following topics:

- Configuring the JUNOScript SSL Service on page 757
- Loading the SSL Certificate from a File or URL on page 758

Configuring the JUNOScript SSL Service

The Secure Sockets Layer (SSL) protocol uses public-private key technology, which requires a paired private key and authentication certificate SSL service. This section describes how to import the SSL certificate in to the JUNOScript server. For a complete example on how to configure the xnm-ssl service, see the *JUNOScript API Guide*.

NOTE: Configuring xnm-ssl service does not apply to IPSec.

To import an SSL certificate into the router, include the local statement at the [edit security certificates] hierarchy level:

```
[edit security certificates]
local certificate-filename;
```

To import the certificate, enter CLI configuration mode on the JUNOScript server machine and issue the following commands at the [edit security certificates] and [edit security certificates local *certificate-name*] hierarchy level:

```
[edit security certificates]
user@host# edit local certificate-name
```

certificate-filename is the name of the certificate.

```
[edit security certificates local certificate-name]
user@host# set load-key-file URL-or-path
```

URL-or-path is the URL or pathname on the local disk.

NOTE: The CLI expects the private key in the specified file (*URL-or-path*) to be unencrypted. If the key is encrypted, the CLI prompts for the passphrase associated with it, decrypts it, and stores the unencrypted version.

Loading the SSL Certificate from a File or URL

To load an SSL certificate from a file or URL, include the local statement at the [edit security certificates] hierarchy level:

```
[edit security certificates]
local local-certificate-filename {
     load-key-file (file-name | url);
}
```

local-certificate-filename is the name of the SSL certificate name that you want to load.

file-name is the name of the file that contains an SSL certificate and private key in PEM format.

url is the SSL certificate and private key PEM location.

Configuring Internal IPSec for JUNOS-FIPS

In a JUNOS-FIPS environment, routers with two Routing Engines must use IPSec for internal communication between the Routing Engines. You configure internal IPSec after you install JUNOS-FIPS. You must be a Crypto Officer to configure internal IPSec.

To configure internal IPSec, include the security-association statement at the [edit security] hierarchy level:

```
[edit security]
ipsec {
    internal {
      security-association {
        manual {
          direction (bidirectional | inbound | outbound) {
            protocol esp;
            spi spi-value;
            encryption {
              algorithm 3des-cbc;
              key ascii-text ascii-text-string;
            }
          }
        }
      }
    }
}
```

This section describes the following tasks for configuring internal IPSec:

- Configuring the SA Direction on page 759
- Configuring the IPSec SPI on page 760
- Configuring the IPSec Key on page 760
- Example: Configuring Internal IPSec on page 760

Configuring the SA Direction

To configure the IPSec SA direction, include the direction statement at the [edit security ipsec internal security-association manual] hierarchy level:

```
direction (bidirectional | inbound | outbound);
```

The value can be one of the following:

- bidirectional—Apply the same SA values in both directions between Routing Engines.
- inbound—Apply these SA properties only to the inbound IPSec tunnel.
- outbound—Apply these SA properties only to the outbound IPSec tunnel.

If you do not configure the SA to be bidirectional, you must configure SA parameters for IPSec tunnels in both directions. The following example uses an inbound and outbound IPSec tunnel:

```
[edit security]
ipsec {
    internal {
      security-association {
        manual {
          direction inbound {
            protocol esp;
            spi 512;
            encryption {
              algorithm 3des-cbc;
              key ascii-text “$.KL3rngIH7,theOPcn87lxfpe9GJKdme”;
            }
          }
          direction outbound {
            protocol esp;
            spi 513;
            encryption {
              algorithm 3des-cbc;
              key ascii-text “.n87IngIH7,thxefpe9GJKdme.KL3rOPc”;
            }
          }
        }
      }
    }
```

Configuring the IPSec SPI

A security parameter index (SPI) is a 32-bit index identifying a security context between a pair of Routing Engines. To configure the IPSec Security Parameter Index (SPI) value, include the spi statement at the [edit security ipsec internal security-association manual direction] hierarchy level:

```
spi value;
```

The value must be from 256 through 16639.

Configuring the IPSec Key

To configure the ASCII text key, include the key statement at the [edit security ipsec internal security-association manual direction encryption] hierarchy level:

```
key ascii-text ascii-text-string;
```

The value must be from 256 through 16639. You must enter the key ASCII value twice and the strings entered must match or the key will not be set. The ASCII text key is never displayed in plain text.

Example: Configuring Internal IPSec

Configure a bidirectional IPSec SA with an SPI value of 512 and a key value conforming to the FIPS 140-2 rules:

```
[edit security]
ipsec {
    internal {
      security-association {
        manual {
          direction bidirectional {
            protocol esp;
            spi 512;
            encryption {
              algorithm 3des-cbc;
              key ascii-text "$9$90j.COIek8X7VevbYgoji1rh";
            }
          }
        }
      }
    }
}
```

Chapter 33

Summary of Security Services Configuration Statements

The following sections explain each of the security services configuration statements. The statements are organized alphabetically.

algorithm

Syntax

```
algorithm 3des-cbc;
```

Hierarchy Level [edit security ipsec internal security-association manual direction encryption]

Release Information Statement introduced before JUNOS Release 7.4.

Description Select the encryption algorithm for the internal Routing-Engine-to-Routing-Engine IPSec security association (SA) configuration.

Options Only 3des-cbc is supported.

Usage Guidelines See "Configuring Internal IPSec for JUNOS-FIPS" on page 758.

Required Privilege Level Crypto Officer—To add and view this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

authentication

Syntax

```
authentication {
    algorithm (hmac-md5-96 | hmac-sha1-96);
    key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bi-directional)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure IP Security (IPSec) authentication parameters for manual security association (SA).

Options algorithm—Hash algorithm that authenticates packet data. It can be one of the following:

- hmac-md5-96—Produces a 128-bit digest.
- hmac-sha1-96—Produces a 160-bit digest.

key—Type of authentication key. It can be one of the following:

- ascii-text key—ASCII text key. For hmac-md5-96, the key is 16 ASCII characters; for hmac-sha1-96, the key is 20 ASCII characters.
- hexadecimal key—Hexadecimal key. For hmac-md5-96, the key is 32 hexadecimal characters; for hmac-sha1-96, the key is 40 hexadecimal characters.

Usage Guidelines See “Configuring the Authentication Algorithm and Key” on page 725.

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

authentication-algorithm

See the following sections:

- authentication-algorithm (IKE) on page 762
- authentication-algorithm (IPSec) on page 763

authentication-algorithm (IKE)

Syntax authentication-algorithm (md5 | sha1);

Hierarchy Level [edit security ike proposal *ike-proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the Internet Key Exchange (IKE) authentication algorithm.

Options authentication-algorithm—Hash algorithm that authenticates packet data. It can be one of two algorithms:

- md5—Produces a 128-bit digest.
- sha1—Produces a 160-bit digest.

Usage Guidelines See “Configuring the Authentication Algorithm for an IKE Proposal” on page 728.

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

authentication-algorithm (IPSec)

Syntax authentication-algorithm (hmac-md5-96 | hmac-sha1-96);

Hierarchy Level [edit security ipsec proposal *ipsec-proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the IPSec authentication algorithm.

Options authentication-algorithm—Hash algorithm that authenticates packet data. It can be one of two algorithms:

- hmac-md5-96—Produces a 128-bit digest.
- hmac-sha1-96—Produces a 160-bit digest.

Usage Guidelines See "Configuring the Authentication Algorithm for an IPSec Proposal" on page 735.

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

authentication-method

Syntax authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);

Hierarchy Level [edit security ike proposal *ike-proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the IKE authentication method.

Options dsa-signatures—Digital Signature Algorithm (DSA)

rsa-signatures—A public key algorithm, which supports encryption and digital signatures

pre-shared-keys—A key derived from an out-of-band mechanism; the key authenticates the exchange

Usage Guidelines See "Configuring the Authentication Method for an IKE Proposal" on page 729.

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

auxiliary-spi

Syntax auxiliary-spi *auxiliary-spi-value*;

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction (JUNOS) (inbound | outbound | bi-directional)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the auxiliary Security Parameter Index (SPI) for a manual SA. Use the auxiliary SPI when you configure the protocol statement to use the bundle option.

Options *auxiliary-spi-value*—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).
Range: 256 through 16,639

Usage Guidelines See "Configuring the Auxiliary Security Parameter Index" on page 724. For information about SPI, see "Configuring the Security Parameter Index" on page 724 and spi on page 793.

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

ca-name

Syntax ca-name *ca-identity*;

Hierarchy Level [edit security certificates certification-authority]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the certificate authority (CA) identity to use in the certificate request.

Usage Guidelines See "Specifying the Certificate Authority Name" on page 743.

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

cache-size

Syntax cache-size *bytes*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the cache size for digital certificates.

Options *bytes*—Cache size for digital certificates.
Range: 64 through 4,294,967,295
Default: 2 megabytes (MB)

NOTE: We recommend that you limit your cache size to 4 MB.

Usage Guidelines See “Configuring the Cache Size” on page 745.

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

cache-timeout-negative

Syntax cache-timeout-negative *seconds*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a negative cache for digital certificates.

Options *seconds*—Negative time to cache digital certificates, in seconds.
Range: 10 through 4,294,967,295
Default: 20

CAUTION: Configuring a large negative cache value can lead to a denial-of-service attack.

Usage Guidelines See “Configuring the Negative Cache” on page 745.

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

certificates

Syntax

```
certificates {
    cache-size bytes;
    cache-timeout-negative seconds;
    certification-authority ca-profile-name {
        ca-name ca-identity;
        crl file-name;
        encoding (binary | pem);
        enrollment-url url-name;
        file certificate-filename;
        ldap-url url-name;
    }
    enrollment-retry attempts;
    local certificate-filename {
        certificate-key-string;
        load-key-file key-filename;
    }
    maximum-certificates number;
    path-length certificate-path-length;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the digital certificates for IPSec.

Usage Guidelines See “Configuring Digital Certificates” on page 739.

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

certification-authority

Syntax

```
certification-authority ca-profile-name {
    ca-name ca-identity;
    crl file-name;
    encoding (binary | pem);
    enrollment-url url-name;
    file certificate-filename;
    ldap-url url-name;
}
```

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a certificate authority profile name. The remaining statements are explained separately.

Usage Guidelines See “Configuring the Certificate Authority Properties” on page 743.

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

crl

Syntax crl *file-name*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPSec peers on a regular periodic basis.

Options *file-name*—Specifies the file from which to read the CRL.

Usage Guidelines See “Configuring the Certificate Authority Properties” on page 743.

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

description

Syntax description *description*;

Hierarchy Level [edit security ike policy *ike-peer-address*],
[edit security ike proposal *ike-proposal-name*],
[edit security ipsec policy *ipsec-policy-name*],
[edit security ipsec proposal *ipsec-proposal-name*],
[edit security ipsec security-association *sa-name*]

Description Specify a text description for an IKE proposal or policy, or an IPSec proposal, policy, or SA.

Usage Guidelines See “Configuring Security Associations” on page 719, “Configuring the Description for an IKE Proposal” on page 729, “Configuring the Description for an IKE Policy” on page 732, “Configuring an IPSec Proposal” on page 734, and “Configuring the IPSec Policy” on page 737.

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

dh-group

Syntax

```
dh-group (group1 | group2);
```

Hierarchy Level [edit security ike proposal *ike-proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the IKE Diffie-Hellman group.

Options dh-group—Type of Diffie-Hellman prime modulus group when performing the new Diffie-Hellman exchange. It can be one of the following:

- group1—768-bit.
- group2—1024-bit.

Usage Guidelines See “Configuring the Diffie-Hellman Group for an IKE Proposal” on page 729.

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

direction

See the following sections:

- direction (JUNOS) on page 768
- direction (JUNOS-FIPS) on page 769

direction (JUNOS)

Syntax

```
direction (inbound | outbound | bi-directional) {
    authentication {
      algorithm (hmac-md5-96 | hmac-sha1-96);
      key (ascii-text key | hexadecimal key);
    }
    auxiliary-spi auxiliary-spi-value;
    encryption {
      algorithm (des-cbc | 3des-cbc);
      key (ascii-text key | hexadecimal key);
    }
    protocol (ah | esp | bundle);
    spi spi-value;
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the direction of IPSec processing.

Options inbound—Inbound SA.

outbound—Outbound SA.

bidirectional—Bidirectional SA.

Usage Guidelines See “Configuring the Processing Direction” on page 722.

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

direction (JUNOS-FIPS)

Syntax

```
direction (bidirectional | inbound | outbound) {
    protocol esp;
    spi (JUNOS) spi-value;
    encryption {
        algorithm 3des-cbc;
        key ascii-text ascii-text-string;
    }
}
```

Hierarchy Level [edit security ipsec internal security-association manual]

Description Establish a manual security association (SA) for internal Routing-Engine-to-Routing-Engine communication.

Options bidirectional—Apply the same SA values in both directions between Routing Engines.

inbound—Apply these SA properties only to the inbound IPSec tunnel.

outbound—Apply these SA properties only to the outbound IPSec tunnel.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Internal IPSec for JUNOS-FIPS” on page 758.

Required Privilege Level Crypto Officer—To view and add this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

dynamic

Syntax

```
dynamic {
    ipsec-policy ipsec-policy-name;
    replay-window-size (32 | 64);
}
```

Hierarchy Level [edit security ipsec security-association *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define a dynamic IPSec SA.

Options ipsec-policy *ipsec-policy-name*—Name of the IPSec policy.

replay-window-size—(Optional) Antireplay window size. It can be one of the following values:

- 32—32-packet window size.
- 64—64-packet window size.

Usage Guidelines See “Configuring Dynamic Security Associations” on page 727 and “Configuring the ES PIC” on page 750.

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

encoding

Syntax encoding (binary | pem);

Hierarchy Level [edit security ike policy *ike-peer-address*],
[edit security certificates certification-authority *ca-profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the file format used for the local-certificate and local-key-pair statements.

Options binary—Binary file format.

pem—Privacy-enhanced mail (PEM), an ASCII base 64 encoded format.
Default: binary

Usage Guidelines See “Configuring the Type of Encoding Your CA Supports” on page 744 and “Configuring the Type of Encoding Your CA Supports” on page 747.

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

encryption

See the following sections:

- encryption (JUNOS) on page 771
- encryption (JUNOS-FIPS) on page 772

encryption (JUNOS)

Syntax

```
encryption {
    algorithm (des-cbc | 3des-cbc);
    key (ascii-text key | hexadecimal key);
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bidirectional)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure an encryption algorithm and key for manual SA.

Options algorithm—Type of encryption algorithm. It can be one of the following:

- des-cbc—Has a block size of 8 bytes (64 bits); its key size is 48 bits long.
- 3des-cbc—Has block size of 8 bytes (64 bits); its key size is 192 bits long.

NOTE: For 3des-cbc, we recommend that the first 8 bytes be different from the second 8 bytes, and the second 8 bytes be the same as the third 8 bytes.

key—Type of encryption key. It can be one of the following:

- ascii-text—ASCII text key. For the des-cbc option, the key contains 8 ASCII characters; for 3des-cbc, the key contains 24 ASCII characters.
- hexadecimal—Hexadecimal key. For the des-cbc option, the key contains 16 hexadecimal characters; for the 3des-cbc option, the key contains 48 hexadecimal characters.

Usage Guidelines See “Configuring the Encryption Algorithm and Key” on page 726.

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

encryption (JUNOS-FIPS)

Syntax

```
encryption {
    algorithm 3des-cbc;
    key ascii-text ascii-text-string;
}
```

Hierarchy Level [edit security ipsec internal security-association manual direction]

Description Define the encryption parameters for internal Routing-Engine-to-Routing-Engine communication. The remaining statements are explained separately.

Usage Guidelines See “Configuring Internal IPSec for JUNOS-FIPS” on page 758.

Required Privilege Level Crypto Officer—To view and add this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

encryption-algorithm

Syntax encryption-algorithm (3des-cbc | des-cbc);

Hierarchy Level [edit security ike proposal *ike-proposal-name*],
[edit security ipsec proposal *ipsec-proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure an IKE or IPSec encryption algorithm.

Options 3des-cbc—Encryption algorithm with key size of 24 bytes; its key size is 192 bits long.

des-cbc—Encryption algorithm with key size of 8 bytes; its key size is 48 bits long.

Usage Guidelines See “Configuring the Encryption Algorithm for an IKE Proposal” on page 730 and “Configuring the Encryption Algorithm for an IPSec Proposal” on page 735.

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

enrollment-retry

Syntax enrollment-retry *attempts*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify how many times a router will resend a digital certificate request.

Options *number*—Number of enrollment retries.
Range: 0 through 100
Default: 0

Usage Guidelines See "Configuring the Number of Enrollment Retries" on page 746

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

enrollment-url

Syntax enrollment-url *url-name*;

Hierarchy Level [edit security certificates certification-authority *ca-profile-name*]

Options *url-name*—Certificate authority URL.

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify where your router should send Simple Certificate Enrollment Protocols-based (SCEP-based) certificate enrollment requests (certificate authority URL).

Usage Guidelines See "Specifying an Enrollment URL" on page 744.

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

file

Syntax file *certificate-filename*;

Hierarchy Level [edit security certificates certification-authority *ca-profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the file from which to read the digital certificate.

Options *certificate-filename*—File from which to read the digital certificate.

Usage Guidelines See “Specifying a File to Read the Digital Certificate” on page 744.

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

identity

Identity identity *identity-name*;

Hierarchy Level [edit security ike]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the identity of the remote certificate name if the identity cannot be learned through IKE (ID payload or IP address).

Usage Guidelines See “Configuring the Identity to Define the Remote Certificate Name” on page 747.

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

ike

Syntax

```
ike {
    policy ike-peer-address {
        description policy-description;
        encoding (binary | pem);
        identity identity-name;
        local-certificate certificate-filename;
        local-key-pair private-public-key-file;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals [ proposal-names ];
    }
    proposal ike-proposal-name {
        authentication-algorithm (md5 | sha1);
        authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
        dh-group (group1 | group2);
        encryption-algorithm (3des-cbc | des-cbc);
        lifetime-seconds seconds;
    }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure IKE.

The statements are explained separately.

Usage Guidelines See “Configuring an IKE Proposal (Dynamic SAs Only)” on page 728 and “Configuring an IKE Policy for Preshared Keys” on page 731.

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

internal

Syntax

```
internal {
    security-association {
        manual {
            direction (bidirectional | inbound | outbound) {
                protocol esp;
                spi spi-value;
                encryption {
                    algorithm 3des-cbc;
                    key ascii-text ascii-text-string;
                }
            }
        }
    }
}
```

Hierarchy Level [edit security ipsec]

Description Define an internal security association (SA) for internal Routing-Engine-to-Routing-Engine communication. The remaining statements are explained separately.

Usage Guidelines See "Configuring Internal IPSec for JUNOS-FIPS" on page 758.

Required Privilege Level Crypto Officer—To view and add this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

ipsec

Syntax

```
ipsec {
    internal {
        security-association {
            manual {
                direction (bidirectional | inbound | outbound) {
                    protocol esp;
                    spi spi-value;
                    encryption {
                        algorithm 3des-cbc;
                        key ascii-text ascii-text-string;
                    }
                }
            }
        }
    }
    policy ipsec-policy-name {
        perfect-forward-secrecy {
            keys (group1 | group2);
        }
        proposals [ proposal-names ];
    }
    proposal ipsec-proposal-name {
      authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
      encryption-algorithm (3des-cbc | des-cbc);
      lifetime-seconds seconds;
      protocol (ah | esp | bundle);
    }
    security-association name {
      dynamic {
        ipsec-policy policy-name;
        replay-window-size (32 | 64);
      }
      manual {
        direction (inbound | outbound | bi-directional) {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
          }
          auxiliary-spi auxiliary-spi-value;
          encryption {
            algorithm (des-cbc | 3des-cbc);
            key (ascii-text key | hexadecimal key);
          }
          protocol (ah | esp | bundle);
          spi spi-value;
        }
      }
      mode (tunnel | transport);
    }
    traceoptions {
      file <files number> < size size>;
      flag all;
      flag database;
      flag general;
      flag ike;
      flag parse;
      flag policy-manager;
      flag routing-socket;
      flag timer;
    }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure IPSec.

The statements are explained separately.

Usage Guidelines See “Configuring Security Associations” on page 719.

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

key

Syntax key ascii-text *ascii-text-string*;

Hierarchy Level [edit security ipsec internal security-association manual direction encryption]

Description The key used for the internal Routing-Engine-to-Routing-Engine IPSec security association (SA) configuration.

Options Only ascii-text is supported.

ascii-text-string—The encrypted ASCII text key.

Usage Guidelines See "Configuring Internal IPSec for JUNOS-FIPS" on page 758.

Required Privilege Level Crypto Officer—To add and view this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

ldap-url

Syntax ldp-url *url-name*;

Hierarchy Level [edit security certificates certification-authority *ca-profile-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description (Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.

Options *url*—Name of the LDAP URL.

Usage Guidelines See "Specifying an LDAP URL" on page 745.

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

lifetime-seconds

Syntax lifetime-seconds *seconds*;

Hierarchy Level [edit security ike proposal *ike-proposal-name*],
[edit security ipsec proposal *ipsec-proposal-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description (Optional) Configure the lifetime of IKE or IPSec SA. When the SA expires, it is replaced by a new SA (and SPI) or terminated.

Options *seconds*—Lifetime, in seconds.

Range: 180 through 86,400

Usage Guidelines See “Configuring the Lifetime for an IKE SA” on page 730 and “Configuring the Lifetime for an IPSec SA” on page 736.

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

local

Syntax

```
local certificate-filename {
    certificate-key-string;
    load-key-file key-filename;
}
```

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Import a Secure Sockets Layer (SSL) certificate into the router.

NOTE: Configuring xnm-ssl service does not apply to IPSec.

Options *certificate-filename*—SSL certificate name.

Usage Guidelines See “Using JUNOScript SSL Service” on page 757.

local-certificate

Syntax local-certificate *certificate-filename*;

Hierarchy Level [edit security ike policy *ike-peer-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the certificate filename from which to read the local certificate.

Options *certificate-filename*—File from which to read the local certificate.

Usage Guidelines See "Specifying the Certificate Filename" on page 748.

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

local-key-pair

Syntax local-key-pair *private-public-key-file*;

Hierarchy Level [edit security ike policy *ike-peer-address*]

Description Specify private and public keys.

Options *private-public-key-file*—Specifies the file from which to read the private and public key pair.

Usage Guidelines See "Specifying the Private and Public Key File" on page 748.

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

manual

See the following sections:

- manual (JUNOS) on page 781
- manual (JUNOS-FIPS) on page 782

manual (JUNOS)

Syntax

```
manual {
    direction (inbound | outbound | bi-directional) {
        authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
            algorithm (des-cbc | 3des-cbc);
            key (ascii-text key | hexadecimal key);
        }
        protocol (ah | esp | bundle);
        spi spi-value;
    }
}
```

Hierarchy Level [edit security ipsec security-association]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define a manual IPSec SA.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Manual Security Associations” on page 722.

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

manual (JUNOS-FIPS)

Syntax

```
manual {
    direction (bidirectional | inbound | outbound) {
        protocol esp;
        spi spi-value;
        encryption {
            algorithm 3des-cbc;
            key ascii-text ascii-text-string;
        }
    }
}
```

Hierarchy Level [edit security ipsec internal security-association]

Description Define a manual security association (SA) for internal Routing-Engine-to-Routing-Engine communication.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring Internal IPSec for JUNOS-FIPS” on page 758.

Required Privilege Level Crypto Officer—To view and add this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

maximum-certificates

Syntax maximum-certificates *number*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the maximum number of peer digital certificates to be cached.

Options *number*—Maximum number of peer digital certificates to be cached.
Range: 64 through 4,294,967,295 peer certificates
Default: 1024 peer certificates

Usage Guidelines See “Configuring the Maximum Number of Peer Certificates” on page 746.

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

mode

See the following sections:

- mode (IKE) on page 783
- mode (IPSec) on page 783

mode (IKE)

Syntax mode (aggressive | main);

Hierarchy Level [edit security ike policy *ike-peer-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the IKE policy mode.

Options mode—Type of IKE policy.

aggressive—Takes half the number of messages of main mode, has less negotiation power, and does not provide identity protection.

main—Uses six messages, in three peer-to-peer exchanges, to establish the IKE SA. These three steps include the IKE SA negotiation, a Diffie-Hellman exchange, and authentication of the peer. Also provides identity protection.
Default: main

Usage Guidelines See "Configuring the Mode for an IKE Policy" on page 732.

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

mode (IPSec)

Syntax mode (transport | tunnel);

Hierarchy Level [edit security ipsec security-association *name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the mode for the IPSec security association.

Options transport— Protects traffic when the communication endpoint and cryptographic endpoint are the same. The data portion of the IP packet is encrypted, but the IP header is not. Virtual Private Network (VPN) gateways that provide encryption and decryption services for protected hosts cannot use transport mode for protected VPN communications.

tunnel—Protects traffic using preshared keys with IKE to authenticate peers or digital certificates with IKE to authenticate peers.
Default: tunnel

NOTE: Tunnel mode requires the ES Physical Interface Card (PIC).

The JUNOS software supports only encapsulating security payload (ESP) when you use tunnel mode.

In transport mode, the JUNOS software does not support authentication header (AH) and ESP header bundles.

In transport mode, the JUNOS software supports only Border Gateway Protocol (BGP).

Usage Guidelines See "Configuring IPSec Mode" on page 720.

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

path-length

Syntax path-length *certificate-path-length*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the digital certificate path length.

Options *certificate-path-length*—Digital certificate path length.
Range: 2 through 15 certificates
Default: 15 certificates

Usage Guidelines See "Configuring the Path Length for the Certificate Hierarchy" on page 746.

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

perfect-forward-secrecy

Syntax

```
perfect-forward-secrecy {
      keys (group1 | group2);
}
```

Hierarchy Level [edit security ipsec policy *ipsec-policy-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description (Optional) Define the Perfect Forward Secrecy (PFS) protocol. Creates single use keys.

Options keys—Type of Diffie-Hellman prime modulus group that IKE uses when performing the new Diffie-Hellman exchange.

The key can be one of the following:

- group1—768-bit.
- group2—1024-bit.

Usage Guidelines See “Configuring Perfect Forward Secrecy” on page 738.

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

policy

See the following sections:

- policy (IKE) on page 786
- policy (IPSec) on page 787

policy (IKE)

Syntax

```
policy ike-peer-address {
    description policy-description;
    encoding (binary | pem);
    identity identity-name;
    local-certificate certificate-filename;
    local-key-pair private-public-key-file;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [ proposal-names ];
}
```

Hierarchy Level [edit security ike]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define an IKE policy.

Options *ike-peer-address*—A tunnel address configured at the [edit interfaces es] hierarchy level.

The remaining statements are explained separately.

Usage Guidelines See “Configuring an IKE Policy for Preshared Keys” on page 731 and “Configuring an IKE Policy for Digital Certificates” on page 747.

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

policy (IPSec)

Syntax

```
policy ipsec-policy-name {
      perfect-forward-secrecy {
         keys (group1 | group2);
      }
      proposals [ proposal-names ];
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define an IPSec policy.

Options *ipsec-policy-name*—Specify an IPSec policy name.

The remaining statements are explained separately.

Usage Guidelines See "Configuring the IPSec Policy" on page 737.

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

pre-shared-key

Syntax pre-shared-key (ascii-text *key* | hexadecimal *key*);

Hierarchy Level [edit security ike policy *ike-peer-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define a preshared key for an IKE policy.

Options preshared-key—Type of preshared key.

The key can be one of the following:

- ascii-text—ASCII text key.
- hexadecimal—Hexadecimal key.

Usage Guidelines See "Configuring the Preshared Key for an IKE Policy" on page 732.

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

proposal

See the following sections:

- proposal (IKE) on page 788
- proposal (IPSec) on page 788

proposal (IKE)

Syntax

```
proposal ike-proposal-name {
    authentication-algorithm (md5 | sha1);
    authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group2);
    encryption-algorithm (3des-cbc | des-cbc);
    lifetime-seconds seconds;
}
```

Hierarchy Level [edit security ike]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define an IKE proposal for a dynamic SA.

Options *ike-proposal-name*—Specifies a IKE proposal name.

The remaining statements are explained separately.

Usage Guidelines See "Configuring an IKE Proposal (Dynamic SAs Only)" on page 728.

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

proposal (IPSec)

Syntax

```
proposal ipsec-proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha1-96);
    encryption-algorithm (3des-cbc | des-cbc);
    lifetime-seconds seconds;
    protocol (ah | esp | bundle);
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define an IPSec proposal for a dynamic SA.

Options *ipsec-proposal-name*—Specifies an IPSec proposal name.

The statements are explained separately.

Usage Guidelines See “Configuring an IPSec Proposal” on page 734.

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

proposals

Syntax proposals [*proposal-names*];

Hierarchy Level [edit security ike policy *ike-peer-address*],
[edit security ipsec policy *ipsec-policy-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Associate one or more proposals with an IKE or IPSec policy.

Options *proposal-names*—Name of one or more proposals.

Usage Guidelines See “Associating Proposals with an IKE Policy” on page 732 and “Configuring the IPSec Policy” on page 737.

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

protocol

See the following sections:

- protocol (JUNOS) on page 789
- protocol (JUNOS-FIPS) on page 790

protocol (JUNOS)

Syntax protocol (ah | esp | bundle);

Hierarchy Level [edit security ipsec proposal *ipsec-proposal-name*],
[edit security ipsec security-association *sa-name* manual direction (JUNOS) (inbound | outbound | bidirectional)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the IPSec protocol for a manual or dynamic SA.

Options ah—Authentication Header protocol

bundle—AH and ESP protocols

esp—ESP protocol (the tunnel statement must be included at the [edit security ipsec security-association *sa-name* mode] hierarchy level)

Usage Guidelines See “Configuring the Protocol for a Manual SA” on page 723 and “Configuring the Protocol for a Dynamic IPSec SA” on page 736.

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

protocol (JUNOS-FIPS)

Syntax protocol esp;

Hierarchy Level [edit security ipsec internal security-association manual direction]

Description The protocol used for the internal Routing-Engine-to-Routing-Engine IPSec security association (SA) configuration.

Options Only esp is supported.

Usage Guidelines See “Configuring Internal IPSec for JUNOS-FIPS” on page 758.

Required Privilege Level Crypto Officer—To add and view this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

security-association

See the following sections:

- security-association (JUNOS) on page 791
- security-association (JUNOS-FIPS) on page 792

security-association (JUNOS)

Syntax

```
security-association sa-name {
    dynamic {
      ipsec-policy policy-name;
      replay-window-size (32 | 64);
    }
    manual {
      direction (JUNOS) (inbound | outbound | bi-directional) {
        authentication {
          algorithm (hmac-md5-96 | hmac-sha1-96);
          key (ascii-text key | hexadecimal key);
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
          algorithm (des-cbc | 3des-cbc);
          key (ascii-text key | hexadecimal key);
        }
        protocol ( ah | esp | bundle);
        spi spi-value;
    }
    mode (tunnel | transport);
}
```

Hierarchy Level [edit security ipsec]

Options *name*—Name of the security association.

The remaining statements are explained separately.

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure an IPSec security association.

Usage Guidelines See “Configuring Security Associations” on page 719.

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

security-association (JUNOS-FIPS)

Syntax

```
security-association {
    manual {
      direction (bidirectional | inbound | outbound) {
        protocol esp;
        spi spi-value;
        encryption {
          algorithm 3des-cbc;
          key ascii-text ascii-text-string;
        }
      }
    }
}
```

Hierarchy Level [edit security ipsec internal]

Description Define a security association (SA) for internal Routing-Engine-to-Routing-Engine communication.

Options The remaining statements are explained separately.

Usage Guidelines See “Configuring Internal IPSec for JUNOS-FIPS” on page 758.

Required Privilege Level Crypto Officer—To view and add this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

spi

See the following sections:

- spi (JUNOS) on page 793
- spi (JUNOS-FIPS) on page 793

spi (JUNOS)

Syntax spi *spi-value*;

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction (inbound | outbound | bi-directional)]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure SPI for an SA.

Options *spi-value*—An arbitrary value that uniquely identifies which SA to use at the receiving host (the destination address in the packet).
Range: 256 through 16639

NOTE: Use the auxiliary SPI when you configure the protocol statement to use the bundle option.

Usage Guidelines See “Configuring the Security Parameter Index” on page 724.

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

spi (JUNOS-FIPS)

Syntax spi *spi-value*;

Hierarchy Level [edit security ipsec internal security-association manual direction]

Description The Security Parameter Index (SPI) value used for the internal Routing-Engine-to-Routing-Engine IPSec security association (SA) configuration.

Options *spi-value*—Integer to use for this SPI.
Range: 256 through 16639

Usage Guidelines See “Configuring Internal IPSec for JUNOS-FIPS” on page 758.

Required Privilege Level Crypto Officer—To add and view this statement in the configuration.

See Also *JUNOS-FIPS Configuration Guide*

traceoptions

Syntax

```
traceoptions {
    file filename <files number> <size size>;
    flag all;
    flag database;
    flag general;
    flag ike;
    flag parse;
    flag policy-manager;
    flag routing-socket;
    flag timer;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure security tracing options.

To specify more than one tracing option, include multiple flag statements. The output of the security tracing options is placed in one file: /var/log/kmd.

Options files *number*—(Optional) Maximum number of trace files. When a trace file (for example, kmd) reaches its maximum size, it is renamed kmd.0, then kmd.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.
Range: 2 through 1000 files
Default: 10 files

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, kmd) reaches this size, it is renamed, kmd.0, then kmd.1 and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
Default: 1024 KB

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

- all—Trace all security events.
- database—Trace database events.
- general—Trace general events.
- ike—Trace IKE module processing.

- parse—Trace configuration processing.
- policy-manager—Trace policy manager processing.
- routing-socket—Trace routing socket messages.
- timer—Trace internal timer events.

Usage Guidelines See “Configuring Trace Options” on page 750.

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

Part 8

Router Chassis

Chapter 34

Router Chassis Configuration Guidelines

You can configure properties of the router chassis, including conditions that activate the red and yellow alarm LEDs on the router's craft interface and SONET/SDH framing and concatenation properties for individual Physical Interface Cards (PICs).

To configure router chassis properties, include the following statements at the [edit chassis] hierarchy level:

```
chassis {
    aggregated-devices {
      ethernet {
        device-count number;
      }
      sonet {
        device-count number;
      }
    }
    alarm {
      interface-type {
        alarm-name (red | yellow | ignore);
      }
    }
    config-button {
      no-clear;
      no-rescue;
    }
    fpc slot-number {
      pic pic-number {
        atm-cell-relay-accumulation;
        atm-l2circuit-mode (cell | aal5 | trunk trunk);
        vtmapping number;
        ce1 {
          e1 port-number {
            channel-group group-number timeslots slot-number;
          }
        }
        ct3 {
          port port-number {
            t1 link-number {
              channel-group group-number timeslots slot-number;
            }
          }
        }
        framing (sdh | sonet);
        idle-cell-format {
            itu-t;
            payload-pattern;
        }
        max-queues-per-interface (8 | 4);
        mlfr-uni-nni-bundles number;
        no-concatenate;
        q-pic-large-buffer:
        t1;
        vtmapping (itu-t | klm);
    }
}
lcc number {
    fpc number {
        pic number {
            atm-cell-relay-accumulation;
            atm-l2-circuit-mode (cell | aal5 | trunk trunk);
            framing (sdh | sonet);
            idle-cell-format {
                itu-t;
                payload-pattern payload-pattern-byte;
            }
            max-queues-per-interface (8 | 4);
            no-concatenate;
        }
    }
    offline;
    online-expected;
}
(packet-scheduling | no-packet-scheduling);
pem {
        minimum number;
}
no-concatenate;
redundancy {
    failover {
        on-loss-of-keepalives;
        on-disk-failure;
    }
    graceful-switchover (disable | enable);
    keepalive-time seconds;
    routing-engine slot-number (master | backup | disabled);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
routing-engine {
    on-disk-failure reboot;
}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
```

```
        (source-route | no-source-route);
        vrf-mtu-check;
    }
```

NOTE: The configuration statements at the [edit chassis lcc] hierarchy level apply only to a routing matrix. For information about a routing matrix, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852 and the *TX Matrix Platform Hardware Guide*.

This chapter describes the following tasks for configuring the router chassis:

- Minimum Chassis Configuration on page 802
- Configuring Aggregated Devices on page 802
- Configuring ATM Cell-Relay Accumulation Mode on an ATM1 PIC on page 803
- Configuring Conditions That Trigger Alarms on page 803
- Configuring SONET/SDH Framing on page 825
- Configuring Sparse DLCI Mode on page 826
- Configuring Channelized PIC Operation on page 827
- Configuring Channelized DS3-to-DS0 Naming on page 828
- Configuring Eight Queues on IQ Interfaces on page 830
- Configuring Channelized E1 Naming on page 830
- Configuring Channelized STM1 Interface Virtual Tributary Mapping on page 832
- Configuring ATM2 Intelligent Queuing Layer 2 Circuit Transport Mode on page 833
- Enabling ILMI for Cell Relay on page 834
- Configuring the Drop Policy for Traffic with Source-Route Constraints on page 834
- Configuring Packet Scheduling on page 835
- Configuring the Link Services PICs on page 835
- Configuring the Idle Cell Format on page 836
- Configuring an MTU Path Check for a Routing Instance on page 837
- Configuring Redundancy on page 838
- Configuring a Routing Engine to Reboot (or Failover) on Hard Disk Errors on page 849

- Configuring the CONFIG (Reset) Button on page 850
- Configuring Larger Delay Buffers on page 851
- Configuring an Entry-Level M320 Router on page 852
- TX Matrix Platform and T640 Routing Node Configuration Guidelines on page 852

Minimum Chassis Configuration

All of the statements at the [edit chassis] hierarchy level of the configuration are optional.

Configuring Aggregated Devices

JUNOS software supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard. To define the virtual links, you need to specify the associations between physical and logical devices within the [edit interfaces] hierarchy, and assign the correct number of logical devices by including the device-count statement at the [edit chassis aggregated-devices ethernet] and [edit chassis aggregated-devices sonet] hierarchy levels:

```
[edit chassis]
aggregated-devices {
      ethernet {
         device-count number;
      }
      sonet {
         device-count number;
      }
}
```

The maximum number of Ethernet logical interface you can configure is 128. The aggregated Ethernet interfaces are numbered from ae0 through ae127. The maximum number of SONET/SDH logical interfaces is 16. The aggregated SONET/SDH interfaces are numbered from as0 through as15.

For more information on physical and logical interfaces using aggregated links, including sample configurations, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring ATM Cell-Relay Accumulation Mode on an ATM1 PIC

You can configure Asynchronous Transfer Mode (ATM) 1 PIC to use cell-relay accumulation mode. In this mode, the incoming cells (1 to 8 cells) are packaged in to a single packet and forwarded to the label-switched path (LSP). At the edge router, this packet is divided in to individual cells and transmitted over the ATM interface.

NOTE: When you configure an ATM PIC to use cell-relay accumulation, all ports on the ATM PIC use cell-relay accumulation mode.

To configure an ATM PIC to use cell-relay accumulation mode, include the atm-cell-relay-accumulation statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ]
atm-cell-relay-accumulation;
```

On a TX Matrix platform, include the atm-cell-relay-accumulation statement at the [edit chassis lcc *numbe*r fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis lcc number fpc slot-number pic pic-number]
atm-cell-relay-accumulation;
```

For more information about configuring a TX Matrix platform, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

Configuring Conditions That Trigger Alarms

For the different types of PICs, you can configure which conditions trigger alarms and whether they trigger a red or yellow alarm. Red alarm conditions light the RED ALARM LED on the router’s craft interface and trigger an audible alarm if one is connected to the contacts on the craft interface. Yellow alarm conditions light the YELLOW ALARM LED on the router’s craft interface and trigger an audible alarm if one is connected to the craft interface.

NOTE: By default, any failure condition on the integrated-services interface (Adaptive Services PIC) triggers a red alarm.

To configure conditions that trigger alarms and that can occur on any interface of the specified type, include the alarm statement at the [edit chassis] hierarchy level:

```
[edit chassis]
alarm {
      interface-type {
         alarm-name (red | yellow | ignore);
      }
}
```

alarm-name is the name of an alarm. Table 35 on page 804 lists the systemwide alarms and the alarms for each interface type.

Table 35: Configurable PIC Alarm Conditions

Interface/System	Alarm Condition	Configuration Option
SONET/SDH and ATM	Link alarm indication signal	ais-l
	Path alarm indication signal	ais-p
	Signal degrade (SD)	ber-sd
	Signal fail (SF)	ber-sf
	Loss of cell delineation (ATM only)	locd
	Loss of framing	lof
	Loss of light	lol
	Loss of pointer	lop-p
	Loss of signal	los
	Phase locked loop out of lock	pll
	Synchronous transport signal (STS) payload label (C2) mismatch	plm-p
	Line remote failure indication	rfi-l
	Path remote failure indication	rfi-p
	STS path (C2) unequipped	uneq-p
E3/T3	Alarm indicator signal	ais
	Excessive numbers of zeros	exz
	Failure of the far end	ferf
	Idle alarm	idle
	Line code violation	lcv
	Loss of frame	lof
	Loss of signal	los
	Phase locked loop out of lock	pll
	Yellow alarm	ylw
Ethernet	Link has gone down	link-down
DS1	Alarm indicator signal	ais
	Yellow alarm	ylw
Integrated-services	Hardware or software failure	failure
Management-Ethernet	Link has gone down	link-down

Chassis Conditions That Trigger Alarms

Various conditions related to the chassis components trigger yellow and red alarms. You cannot configure these conditions. Table 36 through Table 42 list the alarms that the chassis components can generate. For information about chassis alarms for J-series Services Routers, see the *J-series Services Router Administration Guide*. For information about chassis alarms for the TX Matrix platform, see the *TX Matrix Platform Hardware Guide*.

Table 36 lists the alarms that the chassis components can generate on an M5 or M10 Internet router.

Table 36: Chassis Components Alarm Conditions on an M5 or M10 Router

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router boots from an alternate boot device, the hard disk. Typically, the router boots from the flash drive. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see "Boot Devices" on page 328.	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).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace failed fan tray.	Red
Forwarding Engine Board (FEB)	The control board has failed. If this occurs, the board attempts to reboot.	Replace failed FEB.	Red
Flexible PIC Concentrator (FPC)	An FPC has failed. If this occurs, the FPC attempts to reboot. If the FEB sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router's backplane from the front (generally, an FPC) is broken.	--	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
Power supplies	A power supply has been removed from the chassis.	Install missing power supply.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Table 37 lists the alarms that the chassis components can generate on an M7i or M10i Internet router.

Table 37: Chassis Components Alarm Conditions on an M7i or M10i Router

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router has a optional flash disk and boots from an alternate boot device. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see "Boot Devices" on page 328.	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).	Yellow
Compact FEB (CFEB)	For an M7i router, CFEB has failed. If this occurs, the board attempts to reboot.	Replace failed CFEB.	Red
	For an M10i router, both control boards have been removed or have failed.	Replace failed or missing CFEB.	Red
	Too many hard errors in CFEB memory.	Replace failed CFEB.	Red
	Too many soft errors in CFEB memory.	Replace failed CFEB.	Red
	A CFEB microcode download has failed.	Replace failed CFEB.	Red
Fan trays	A fan has failed.	Replace failed fan tray.	Red
	For an M7i router, a fan tray has been removed from the chassis.	Install missing fan tray.	Red
	For an M10i router, both fan trays are absent from the chassis.	Install missing fan tray.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router's midplane from the front is broken.	---	Red
Power supplies	A power supply has been removed.	Insert missing power supply.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
	For an M10i router, only one power supply is operating.	Insert or replace secondary power supply.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk. This alarm only applies, if you have an optional flash drive.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Table 38 lists the alarms that the chassis components can generate on an M20 Internet router.

Table 38: Chassis Components Alarm Conditions for an M20 Router

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The router boots from an alternate boot device, the hard disk. Typically, the router boots from the flash drive. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see “Boot Devices” on page 328.	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).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below requires speed.	Replace fan tray.	Red
FPC	An FPC has failed. If this occurs, the FPC attempts to reboot. If the System and Switch Board (SSB) sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs in to the router’s backplane from the front (generally, an FPC) is broken.	--	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default master Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
SSB	The control board has failed. If this occurs, the board attempts to reboot.	Replace failed control board.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Table 39 lists the alarms that the chassis components can generate on an M40 Internet router.

Table 39: Chassis Component Alarm Conditions for an M40 Router

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filter	Change air filter.	Change air filter.	------------------------
Alternative media	The router boots from an alternate boot device, the hard disk. Typically, the router boots from the flash drive. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see "Boot Devices" on page 328.	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).	Yellow
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	An FPC microcode download has failed.	Replace failed FPC.	Red
	An FPC has failed. If this occurs, the FPC attempts to reboot. If the SCB sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
	Too many hard errors in FPC memory.	Replace failed FPC.	Red
	Too many soft errors in FPC memory.	Replace failed FPC.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router's backplane from the front (generally, an FPC) is broken.	--	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply temperature sensor has failed.	Replace failed power supply or power entry module.	Yellow
	A power supply fan has failed.	Replace failed power supply fan.	Yellow
	A power supply has high temperature.	Replace failed power supply or power entry module.	Red
	A 5V power supply has failed.	Replace failed power supply or power entry module.	Red
	A 3.3V power supply has failed.	Replace failed power supply or power entry module.	Red
	A 2.5V power supply has failed.	Replace failed power supply or power entry module.	Red
	A power supply input has failed.	Check power supply input connection.	Red
	A power supply has failed.	Replace failed power supply or power entry module.	Red
Routing Engine	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default master Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
SCB	The System Control Board (SCB) has failed. If this occurs, the board attempts to reboot.	Replace failed SCB.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Table 40 lists the alarms that the chassis components can generate on an M40e or M160 Internet router.

Table 40: Chassis Component Alarm Conditions for an M40e or M160 Router

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filter	Change air filter.	Change air filter	--------------
Alternative media	The router boots from an alternate boot device, the hard disk. Typically, the router boots from the flash drive. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see “Boot Devices” on page 328.	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).	Yellow

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Connector Interface Panel (CIP)	A CIP is missing.	Insert CIP into empty slot.	Red
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	An FPC microcode download has failed.	Replace failed FPC.	Red
	An FPC has failed. If this occurs, the FPC attempts to reboot. If the MCS sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
	Too many hard errors in FPC memory.	Replace failed FPC.	Red
	Too many soft errors in FPC memory.	Replace failed FPC.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router's backplane from the front (generally, an FPC) is broken.	----------------------------------	Red
Miscellaneous Control Subsystem (MCS)	An MCS has an out of range or invalid temperature reading.	Replace failed MCS.	Yellow
	MCS0 has been removed.	Reinstall MCS0.	Yellow
	An MCS has failed.	Replace failed MCS.	Red
Packet Forwarding Engine Clock Generator (PCG)	A backup PCG is offline.	Set backup PCG online.	Yellow
	A PCG has an out of range or invalid temperature reading.	Replace failed PCG.	Yellow
	A PCG has been removed.	Insert PCG into empty slot.	Yellow
	A PCG has failed to come online.	Replace failed PCG.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default master Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
Switching and Forwarding Module (SFM)	A SFM has an out of range or invalid temperature reading on SPP.	Replace failed SFM.	Yellow
	A SFM has an out of range or invalid temperature reading on SPR.	Replace failed SFM.	Yellow
	A SFM is offline.	Set SFM online.	Yellow
	A SFM has failed.	Replace failed SFM.	Red
	A SFM has been removed from the chassis.	Insert SFM into empty slot.	Red
	All SFMs are offline or missing from the chassis.	Insert SFMs into empty slots or set all SFMs online.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Table 41 lists the alarms that the chassis components can generate on an M320 Internet router.

Table 41: Chassis Component Alarm Conditions for an M320 Router

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filters	Change air filter.	Change air filter.	----------------------
Alternative media	The router boots from an alternate boot device, the hard disk. Typically, the router boots from the flash drive. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see “Boot Devices” on page 328.	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).	Yellow
Control Board (CB)	A CB has been removed.	Insert CB into empty slot.	Yellow
	A CB temperature sensor alarm has failed.	Replace failed CB.	Yellow
	A CB has failed.	Replace failed CB.	Red
CIP	A CIP is missing.	Insert CIP into empty slot.	Red
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	A FPC microcode download has failed.	Replace failed FPC.	Red
	An FPC has failed. If this occurs, the FPC attempts to reboot. If the CB sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
	Too many hard errors in FPC memory.	Replace failed FPC.	Red
	Too many soft errors in FPC memory.	Replace failed FPC.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router's backplane from the front (generally, an FPC) is broken.	--------------------------------------	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
Routing Engine	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default master Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
Switch Interface Board (SIB)	A spare SIB is missing.	Insert spare SIB in to empty slot.	Yellow
	An SIB has failed.	Replace failed SIB.	Yellow
	A spare SIB has failed.	Replace failed SIB.	Yellow
	An SIB has an out of range or invalid temperature reading.	Replace failed SIB.	Yellow
	An SIB is missing.	Insert SIB into empty slot.	Red
	An SIB has failed.	Replace failed SIB.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	Chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Table 42 lists the alarms that the chassis components can generate on a T320 or T640 Internet routing platform.

Table 42: Chassis Component Alarm Conditions for the T320 or T640 Routing Platform

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Air filter	Change air filter.	Change air filter.	----------------
Alternative media	The router boots from an alternate boot device, the hard disk. Typically, the router boots from the flash drive. If you configure your router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see “Boot Devices” on page 328.	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).	Yellow
CB	A CB has been removed.	Insert CB into empty slot.	Yellow
	A CB temperature sensor alarm has failed.	Replace failed CB.	Yellow
	A CB has failed.	Replace failed CB.	Red
CIP	A CIP is missing.	Insert CIP into empty slot.	Red
Craft interface	The craft interface has failed.	Replace failed craft interface.	Red
Fan trays	One fan tray has been removed from the chassis.	Install missing fan tray.	Yellow
	Two or more fan trays have been removed from the chassis.	Install missing fan trays.	Red
	One fan in the chassis is not spinning or is spinning below required speed.	Replace fan tray.	Red
FPC	An FPC has an out of range or invalid temperature reading.	Replace failed FPC.	Yellow
	An FPC microcode download has failed.	Replace failed FPC.	Red
	An FPC has failed. If this occurs, the FPC attempts to reboot. If the CB sees that an FPC is rebooting too often, it shuts down the FPC.	Replace failed FPC.	Red
	Too many hard errors in FPC memory.	Replace failed FPC.	Red
	Too many soft errors in FPC memory.	Replace failed FPC.	Red
Hot swapping	Too many hot-swap interrupts are occurring. This message generally indicates that a hardware component that plugs into the router’s backplane from the front (generally, an FPC) is broken.	--	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	Error in reading or writing hard disk.	Reformat hard disk and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	Error in reading or writing compact flash.	Reformat compact flash and install bootable image. If this fails, replace failed Routing Engine.	Yellow
	System booted from default backup Routing Engine. If you manually switched mastership, ignore this alarm condition.	Install bootable image on default master Routing Engine. If this fails, replace failed Routing Engine.	Yellow
	System booted from hard disk.	Install bootable image on compact flash. If this fails, replace failed Routing Engine.	Yellow
	Compact flash missing in boot list.	Replace failed Routing Engine.	Red
	Hard disk missing in boot list.	Replace failed Routing Engine.	Red
	Routing Engine failed to boot.	Replace failed Routing Engine.	Red
Power supplies	A power supply has been removed from the chassis.	Insert power supply into empty slot.	Yellow
	A power supply has failed.	Replace failed power supply.	Red
SONET Clock Generator (SCG)	A backup SCG is offline.	Set backup SCG online.	Yellow
	An SCG has an out of range or invalid temperature reading.	Replace failed SCG.	Yellow
	An SCG has been removed.	Insert SCG into empty slot.	Yellow
	All SCGs are offline or missing.	Insert SCGs into empty slots or set all SCGs online.	Red
	An SCG has failed.	Replace failed SCG.	Red
SIB	A spare SIB is missing.	Insert spare SIB into empty slot.	Yellow
	An SIB has failed.	Replace failed SIB.	Yellow
	A spare SIB has failed.	Replace failed SIB.	Yellow
	A SIB has an out of range or invalid temperature reading.	Replace failed SIB.	Yellow
	An SIB is missing.	Insert SIB into empty slot.	Red
	An SIB has failed.	Replace failed SIB.	Red
Switch Processor Mezzanine Board (SPMB)	A local SPMB is offline.	Reset control board. If this fails, replace control board.	Red

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Temperature	The chassis temperature has exceeded 55 degrees C, the fans have been turned on to full speed, and one or more fans have failed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and the fans have been turned on to full speed.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Yellow
	The chassis temperature has exceeded 65 degrees C and a fan has failed. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	Chassis temperature has exceeded 75 degrees C. If this condition persists for more than 4 minutes, the router shuts down.	■ Check room temperature. ■ Check air filter and replace it. ■ Check air flow. ■ Check fan.	Red
	The temperature sensor has failed.	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).	Red

Backup Routing Engine Alarms

For routers with master and backup Routing Engines, a master Routing Engine can generate alarms for events that occur on a backup Routing Engine. Table 43 lists chassis alarms generated for a backup Routing Engine.

NOTE: Because the failure occurs on the backup Routing Engine, alarm severity for some events (such as Ethernet interface failures) is yellow instead of red.

Table 43: Backup Routing Engine Alarms

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Alternative media	The backup Routing Engine boots from an alternate boot device, the hard disk. Typically, routers boot from the flash drive. If you configure your backup Routing Engine router to boot from the hard disk, ignore this alarm condition. For more information about alternate boot devices, see "Boot Devices" on page 328.	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).	Yellow
Boot Device	The boot device (compact flash or hard disk) is missing in boot list on the backup Routing Engine.	Replace failed backup Routing Engine.	Red
Ethernet	The Ethernet management interface (fxp0) on the backup Routing Engine is down.	■ Check the interface cable connection. ■ Reboot the system. ■ If the alarm reoccurs, 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)	Yellow
FRU Offline	When the backup routing engine stops communicating with the master routing engine.	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).	Yellow
Hard Disk	Error in reading or writing hard disk on the backup Routing Engine.	Reformat hard disk and install bootable image. If this fails, replace failed backup Routing Engine.	Yellow
Multibit Memory ECC	The backup Routing Engine reports an multi-bit ECC error.	■ Reboot the system with the board reset button on the backup Routing Engine. ■ If the alarm reoccurs, 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)	Yellow

Silencing External Devices

You can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button located on the craft interface front panel. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after an external device is silenced reactivate the external device.

Configuring SONET/SDH Framing

By default, SONET/SDH PICs use SONET framing. For a discussion of the differences between the two standards, see the *JUNOS Network Interfaces Configuration Guide*. To configure a PIC to use SDH framing, include the framing statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level, specifying the sdh option:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number framing sdh
[edit chassis]
user@host# show
fpc slot-number {
    pic pic-number {
      framing sdh;
    }
}
```

On a TX Matrix platform, include the framing statement at the [edit chassis lcc *number* fpc *slot-number* pic *pic-number*] hierarchy level, specifying the sdh option:

```
[edit chassis lcc number]
user@host# set fpc slot-number pic pic-number framing sdh
[edit chassis lcc number]
user@host# show
fpc slot-number {
    pic pic-number {
      framing sdh;
    }
}
```

To explicitly configure a PIC to use SONET framing, include the framing statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level, specifying the sonet option:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number framing sonet
[edit chassis]
user@host# show
fpc slot-number {
    pic pic-number {
      framing sonet;
    }
}
```

On a TX Matrix platform, include the framing statement at the [edit chassis lcc *number* fpc *slot-number* pic *pic-number*] hierarchy level, specifying the sonet option:

```
user@host# set fpc slot-number pic pic-number framing sonet
[edit chassis lcc number]
user@host# show
fpc slot-number {
     pic pic-number {
        framing sonet;
     }
}
```

For information about configuring a TX Matrix platform, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

Configuring Sparse DLCI Mode

By default, original channelized DS3 and original channelized STM1-to-E1 (or T1) interfaces can support a maximum of 64 data-link connection identifiers (DLCIs) per channel—as many as 1792 DLCIs per DS3 interface or 4032 DLCIs per STM1 interface (0 through 63).

In sparse DLCI mode, the full DLCI range (1 through 1022) is supported. This allows you to use circuit cross-connect (CCC) and translation cross-connect (TCC) features by means of Frame Relay on T1 and E1 interfaces. For more information about CCC and DLCIs, see the *JUNOS Network Interfaces Configuration Guide*.

NOTE: Sparse DLCI mode requires a Channelized STM1 or Channelized DS3 PIC.

DLCI 0 is reserved for Local Management Interface (LMI) signaling.

Channelized T3 intelligent queuing (IQ) and STM1 IQ interfaces support a maximum of 64 DLCIs, numbered 0 through 1022, and therefore do not require sparse mode.

To configure the router to use sparse DLCI mode, include the sparse-dlcis statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ]
sparse-dlcis;
```

Configuring Channelized PIC Operation

By default, SONET PICs (interfaces with names so-*fpc/pic/port*) operate in concatenated mode, a mode in which the bandwidth of the interface is in a single channel.

To configure a PIC to operate in channelized (multiplexed) mode, include the no-concatenate statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number no-concatenate
[edit chassis]
user@host# show
fpc slot-number {
      pic pic-number {
         no-concatenate;
      }
}
```

On a TX Matrix platform, include the no-concatenate statement at the [edit chassis lcc *number* fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis lcc number]
user@host# set fpc slot-number pic pic-number no-concatenate
[edit chassis lcc number]
user@host# show
fpc slot-number {
      pic pic-number {
         no-concatenate;
      }
}
```

When configuring and displaying information about interfaces that are operating in channelized mode, you must specify the channel number in the interface name (*physical:channel*); for example, so-2/2/0:0 and so-2/2/0:1. For more information about interface names, see the *JUNOS Network Interfaces Configuration Guide*. For information about the TX Matrix platform, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

Concatenated and Nonconcatenated Mode

On SONET OC48 interfaces that are configured for channelized (multiplexed) mode, the bytes e1-quiet and bytes f1 options in the sonet-options statement have no effect. The bytes f2, bytes z3, bytes z4, and path-trace options work correctly on channel 0. These bytes work in the transmit direction only on channels 1, 2, and 3.

The M160 four-port SONET/SDH OC12 PIC can run each of the OC12 links in concatenated mode only and requires a Type 2 M160 FPC. Similarly, the four-port SONET/SDH OC3 PIC cannot run in nonconcatenated mode on any platform.

Configuring Channelized DS3-to-DS0 Naming

You can configure 28 T1 channels per T3 interface. Each T1 link can have up to eight channel groups, and each channel group can hold any combination of DS0 timeslots. To specify the T1 link and DS0 channel group number in the name, use colons (:) as separators. For example, a Channelized DS3-to-DS0 PIC might have the following physical and virtual interfaces:

```
ds-0/0/0:x:y
```

where *x* is a T1 link ranging from 0 through 27 and *y* is a DS0 channel group ranging from 0 through 7 (see Table 44 on page 829 for more information about ranges).

You can use any of the values within the range available for *x* and *y*; you do not have to configure the links sequentially. The software applies the interface options you configure according to the following rules:

- You can configure t3-options for t1 link 0 and channel group 0 only; for example, ds-/0/0/0:0:0.
- You can configure t1-options for any t1 link value, but only for channel group 0; for example, ds-0/0/0:x:0.
- There are no restrictions on changing the default ds0-options.
- If you delete a configuration you previously committed for channel group 0, the options return to the default values.

To configure the channel groups and timeslots for a channelized DS3 interface, include the channel-group and timeslots statements at the [edit chassis fpc slot-number pic *pic-number* ct3 port *port-number* t1 *link-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
fpc slot-number {
      pic pic-number {
         ct3 {
            port port-number {
               t1 link-number {
                  channel-group group-number timeslots slot-number;
               }
            }
         }
      }
}
```

NOTE: If you commit the interface name but do not include the [edit chassis] configuration, the Channelized DS3-to-DS0 PIC behaves like a Channelized DS3-to-DS1 PIC: none of the DS0 functionality is accessible.

Table 44 shows the ranges for each of the quantities in the preceding configuration.

Table 44: Ranges for Channelized DS3-to-DS0 Configuration

Item	Variable	Range
FPC slot	*slot-number*	0 through 7(see note below)
PIC slot	*pic-number*	0 through 3
Port	*port-number*	0 through 1
T1 link	*link-number*	0 through 27
DS0 channel group	*group-number*	0 through 7
timeslot	*slot-number*	1 through 24

NOTE: The FPC slot range depends on the platform. The maximum range of 0 through 7 applies to M40 routers; for M20 routers, the range is 0 through 3; for M10 routers the range is 0 through 1; for M5 routers, the only applicable value is 0. The Multichannel DS3 (Channelized DS3-to-DS0) PIC is not supported on M160 routers.

Bandwidth limitations restrict the interface to a maximum of 128 channel groups per T3 port, rather than the theoretical maximum of 8 x 28 = 224.

There are 24 timeslots on a T1 interface. You can designate any combination of timeslots for usage, but you can use each timeslot number on only one channel group within the same T1 link.

To use timeslots 1 through 10, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
channel-group group-number timeslots 1-10;
```

To use timeslots 1 through 5, timeslot 10, and timeslot 24, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
channel-group group-number timeslots 1-5,10,24;
```

Note that spaces are not allowed when you specify timeslot numbers. For more information about these interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring Eight Queues on IQ Interfaces

By default, IQ PICs on T-series and M320 routing platforms are restricted to a maximum of four egress queues per interface. To configure a maximum of eight egress queues on IQ interfaces, include the max-queues-per-interface statement at the [edit chassis fpc *slot-numbe*r pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
max-queues-per-interface (8 | 4);
```

On a TX Matrix platform, include the max-queues-per-interface statement at the [edit chassis lcc *number* fpc *slot-numbe*r pic *pic-number*] hierarchy level:

```
[edit chassis lcc number fpc slot-number pic pic-number]
max-queues-per-interface (8 | 4);
```

NOTE: The configuration at the [edit class-of-service] hierarchy level must also support eight queues per interface.

The maximum number of queues per IQ PIC can be 4 or 8.

If you include the max-queues-per-interface statement, all ports on the IQ PIC use configured mode and all interfaces on the IQ PIC have the same maximum number of queues.

When you include the max-queues-per-interface statement and commit the configuration, all physical interfaces on the IQ PIC are deleted and readded. Also, the PIC is taken offline and then brought back online immediately. You do not need to take the PIC offline and online manually. You should change modes between four queues and eight queues only when there is no active traffic going to the IQ PIC.

For more information about how to configure eight queues on each interface, see the *JUNOS Class of Service Configuration Guide*. For information about the TX Matrix platform, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

Configuring Channelized E1 Naming

Each Channelized E1 PIC has 10 E1 ports that you can channelize to the *N*xDS0 level. Each E1 interface has 32 timeslots (DS0), in which timeslot 0 is reserved. You can combine one or more of these DS0 (channels) to create a channel group (*N*xDS0). There can be a maximum of 24 channel groups per E1 interface. Thus, you can configure as many as 240 channel groups per PIC (10 ports x 24 channel groups per port).

To specify the DS0 channel group number in the interface name, include a colon (:) as a separator. For example, a Channelized E1 PIC might have the following physical and virtual interfaces:

```
ds-0/0/0:x
```

where *x* is a DS0 channel group ranging from 0 through 23 (see Table 45 on page 831 for more information about ranges).

You can use any of the values within the range available for *x*; you do not have to configure the links sequentially. The software applies the interface options you configure according to the following rules:

- You can configure the e1-options statement for channel group 0 only; for example, ds-0/0/0:0.
- There are no restrictions on changing the default ds0-options.
- If you delete a configuration you previously committed for channel group 0, the options return to the default values.

To configure the channel groups and timeslots for a Channelized E1 interface, include the channel-group and timeslots statements at the [edit chassis fpc *slot-number* pic *pic-number* ce1 e1 *port-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]
    fpc slot-number {
      pic pic-number {
        ce1 {
          e1 port-number {
            channel-group group-number timeslots slot-number;
          }
        }
      }
}
```

NOTE: If you commit the interface name but do not include the [edit chassis] configuration, the Channelized E1 PIC behaves like a standard E1 PIC: none of the DS0 functionality is accessible.

Table 45 shows the ranges for each of the quantities in the preceding configuration.

Table 45: Ranges for Channelized E1 Configuration

Item	Variable	Range
FPC slot	*slot-number*	0 through 7 (see note below)
PIC slot	*pic-number*	0 through 3
E1 port	*port-number*	0 through 9
DS0 channel group	*group-number*	0 through 23
Timeslot	*slot-number*	1 through 32

NOTE: The FPC slot range depends on the platform. The maximum range of 0 through 7 applies to M40 routers; for M20 routers, the range is 0 through 3; for M10 routers the range is 0 through 1; for M5 routers, the only applicable value is 0. The Channelized E1 PIC is not supported on M160 routers.

The theoretical maximum number of channel groups possible per PIC is 10 x 24 = 240. This is within the maximum bandwidth available.

There are 32 timeslots on an E1 interface. You can designate any combination of timeslots for usage.

To use timeslots 1 through 10, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]
channel-group group-number timeslots 1-10;
```

To use timeslots 1 through 5, timeslot 10, and timeslot 24, designate *slot-number* as follows:

```
[edit chassis fpc slot-number pic pic-number ce1 e1 port-number]
channel-group group-number timeslots 1-5,10,24;
```

Note that spaces are not allowed when you specify timeslot numbers.

For further information about these interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring Channelized STM1 Interface Virtual Tributary Mapping

By default, virtual tributary mapping uses KLM mode. You can configure virtual tributary mapping to use KLM or ITU-T mode. On the original Channelized STM1 PIC, to configure virtual tributary mapping, include the vtmapping statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
vtmapping (klm | itu-t);
```

For the Channelized STM1 PIC with IQ, you can configure virtual tributary mapping by including the vtmapping statement at the [edit interfaces cau4 fpc *slot-number* pic *pic-number* sonet-options] hierarchy level. For more information, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring ATM2 Intelligent Queuing Layer 2 Circuit Transport Mode

On ATM2 IQ PICs only, you can configure Layer 2 circuit cell relay, Layer 2 circuit ATM Adaptation Layer 5 (AAL5), or Layer 2 circuit trunk mode.

Layer 2 circuit cell relay and Layer 2 circuit AAL5 are defined in the Internet draft draft-martini-l2circuit-encap-mpls-04.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks.*

Layer 2 circuit trunk mode allows you to send ATM cells over Multiprotocol Label Switching (MPLS) trunking.

The four transport modes are defined as follows:

- To tunnel IP packets over an ATM backbone, use the default standard AAL5 transport mode.
- To tunnel a stream of AAL5-encoded ATM segmentation-and-reassembly protocol data units (SAR-PDUs) over an MPLS or IP backbone, use Layer 2 circuit AAL5 transport mode.
- To tunnel a stream of ATM cells over an MPLS or IP backbone, use Layer 2 circuit cell-relay transport mode.
- To transport ATM cells over an MPLS core network that is implemented on some other vendor switches, use Layer 2 circuit trunk mode.

NOTE: You can transport AAL5-encoded traffic with Layer 2 circuit cell-relay transport mode, because Layer 2 circuit cell-relay transport mode ignores the encoding of the cell data presented to the ingress interface.

When you configure AAL5 mode Layer 2 circuits, the control word carries cell loss priority (CLP) information by default.

By default, ATM2 IQ PICs are in standard AAL5 transport mode. Standard AAL5 allows multiple applications to tunnel the protocol data units of their Layer 2 protocols over an ATM virtual circuit. To configure the Layer 2 circuit transport modes, include the atm-l2circuit-mode statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
atm-l2circuit-mode (cell | aal5 | trunk trunk);
```

On a TX Matrix platform, include the atm-l2circuit-mode statement at the [edit chassis lcc *numbe*r fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis lcc number fpc slot-number pic pic-number]
atm-l2circuit-mode (cell | aal5 | trunk trunk);
```

aal5 tunnels a stream of AAL5-encoded ATM cells over an IP backbone.

cell tunnels a stream of ATM cells over an IP backbone.

trunk transports ATM cells over an MPLS core network that is implemented on some other vendor switches. Trunk mode can be user-to-network interface (UNI) or network-to-network interface (NNI).

NOTE: To determine which vendors support Layer 2 circuit trunk mode, contact Juniper Networks customer support.

For more information about ATM Layer 2 circuit transport mode, see the *JUNOS Network Interfaces Configuration Guide* and the *JUNOS Feature Guide*. For information about the TX Matrix platform, see "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

Enabling ILMI for Cell Relay

Integrated Local Management Interface (ILMI) is supported on AAL5 interfaces, regardless of transport mode. To enable ILMI on interfaces with cell-relay encapsulation, you must configure an ATM2 IQ PIC to use Layer 2 circuit trunk transport mode.

To configure ILMI on an interface with cell-relay encapsulation, include the following statements:

```
[edit chassis fpc slot-number pic pic-number]
atm-l2circuit-mode trunk trunk;

[edit interfaces at-fpc/pic/port]
encapsulation atm-ccc-cell-relay;
atm-options {
    ilmi;
    pic-type atm2;
}
unit logical-unit-number {
    trunk-id number;
}
```

For an example on how to enable ILMI for cell relay, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring the Drop Policy for Traffic with Source-Route Constraints

By default, the router forwards IP traffic that has either loose or strict source-route constraints. However, you might want the router to use only the IP destination address on transit traffic for forwarding decisions. You can configure the router to discard IP traffic with source-route constraints by including the no-source-route statement at the [edit chassis] hierarchy level:

```
[edit chassis]
no-source-route;
```

Configuring Packet Scheduling

By default, packet scheduling is disabled. To configure a router to operate in packet-scheduling mode, include the packet-scheduling statement at the [edit chassis] hierarchy level:

```
[edit chassis]
packet-scheduling;
```

To explicitly disable the packet-scheduling statement, include the no-packet-scheduling statement at the [edit chassis] hierarchy level:

```
[edit chassis]
no-packet-scheduling;
```

When you enable packet-scheduling mode, the Packet Director application-specific integrated circuit (ASIC) schedules packet dispatches to compensate for transport delay differences. This preserves the interpacket gaps as the packets are distributed from the Packet Director ASIC to the Packet Forwarding Engine.

Whenever you change the configuration for packet-scheduling, the system stops all SFMs and FPCs and restarts them in the new mode.

NOTE: Packet scheduling is for M160 routers only.

Configuring the Link Services PICs

The Multilink Protocol enables you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members.

The Link Services PIC supports the following Multilink Protocol encapsulation types at the logical unit level:

- Multilink Point-to-Point Protocol (MLPPP)
- Multilink Frame Relay (MLFR FRF.15)

The Link Services PIC also supports the Multilink Frame Relay UNI and NNI (MLFR FRF.16) encapsulation type at the physical interface level.

MLFR (FRF.16) is supported on a channelized interface, *ls-fpc/pic/port*:*channel*, which denotes a single MLFR (FRF.16) bundle. For MLFR (FRF.16), multiple links are combined to form one logical link. Packet fragmentation and reassembly occur on a per-virtual circuit (VC) basis. Each bundle can support multiple VCs. The physical connections must be E1, T1, channelized DS3 to DS1, channelized DS3 to DS0, channelized E1, channelized STM 1, or channelized IQ interfaces.

The default number of bundles per Link Services PIC is 16, ranging from ls-fpc/pic/port:0 to ls-fpc/pic/port:15.

To configure the number of bundles on a Link Services PIC, include the mlfr-uni-nni-bundles statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
mlfr-uni-nni-bundles number;
```

The maximum number of MLFR UNI NNI bundles each Link Services PIC can accommodate is 128. A link can associate with one link services bundle only. For more information, see the *JUNOS Services Interfaces Configuration Guide*.

NOTE: The Link Services PIC is not compatible with the M160 or T-series routing platforms.

Multiclass Extension to MLPPP (RFC 2686)

This extension enables multiple classes of service using MLPPP. For more information, see RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. The JUNOS software PPP implementation does not support the negotiation of address field compression and protocol field compression PPP NCP options. The software always send a full 4-byte PPP header.

Configuring the Idle Cell Format

ATM devices send idle cells to enable the receiving ATM interface to recognize the start of each new cell. The receiving ATM device does not act on the contents of idle cells and does not pass them up to the ATM layer in the ATM protocol stack.

By default, the idle cell format for ATM cells is (4 bytes): 0x00000000. For ATM 2 PICs only, you can configure the format of the idle cell header and payload bytes.

To configure the idle cell header to use the International Telecommunications Union (ITU-T) standard of 0x00000001, include the itu-t statement at the [edit chassis fpc *slot-number* pic *number* idle-cell-format] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number idle-cell-format]
itu-t;
```

On a TX Matrix platform, include the itu-t statement at the [edit chassis lcc *numbe*r fpc *slot-number* pic *pic-number* idle-cell-format] hierarchy level:

```
[edit chassis lcc number fpc slot-number pic pic-number idle-cell-format]
itu-t;
```

By default, the payload pattern is cell payload (48 bytes). To configure the idle cell payload pattern, include the payload-pattern statement at the [edit chassis fpc *slot-number* pic *number* idle-cell-format] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number idle-cell-format]
payload-pattern payload-pattern-byte;
```

On a TX Matrix platform, include the payload-pattern statement at the [edit chassis lcc *numbe*r fpc *slot-number* pic *pic-number* idle-cell-format] hierarchy level:

```
[edit chassis lcc number fpc slot-number pic pic-number]
payload-pattern payload-pattern-byte;
```

The payload pattern byte can range from 0x00 through 0xff.

For information about the TX Matrix platform, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

Configuring an MTU Path Check for a Routing Instance

By default, the maximum transmission unit (MTU) check for a routing instance is not enabled.

NOTE: The MTU check is automatically present for interfaces belonging to the main router.

On M-series routers (except the M320 router) you can configure MTU path checks on the outgoing interface for unicast traffic routed on a virtual private network (VPN) routing and forwarding (VRF) routing instance. When you enable MTU check, the routing platform sends an Internet Control Message Protocol (ICMP) message when the size of a unicast packet traversing a VRF routing instance or virtual-router routing instance has exceeded the MTU size and when an IP packet is set to "do not fragment". The ICMP message uses the routing instance local address as its source address.

For an MTU check to work in a routing instance, you must include the vrf-mtu-check statement at the [edit chassis] hierarchy level and assign at least one interface containing an IP address to the routing instance.

To configure path MTU checks, do the following:

- Enabling MTU Check for a Routing Instance on page 837
- Assigning an IP Address to an Interface in the Routing Instance on page 838

Enabling MTU Check for a Routing Instance

To enable MTU check for a routing instance, include the vrf-mtu-check statement at the [edit chassis] hierarchy level:

```
[edit chassis]
vrf-mtu-check;
```

Assigning an IP Address to an Interface in the Routing Instance

To assign an IP address to an interface in the VRF or virtual-router routing instance, configure the local address for that routing instance. A local address is any IP address derived from an interface that is assigned to the routing instance.

To assign an interface to a routing instance, include the interface statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
interface interface-name;
```

To configure an IP address for a loopback interface, include the address statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet]
address address;
```

NOTE: If you are assigning Internet Protocol Security (IPSec) or generic routing encapsulation (GRE) tunnel interfaces without IP addresses in the routing instance, include a loopback interface to the routing instance. To do this, include the lo0.*n* option at the [edit routing-instances *routing-instance-name* interface] hierarchy level. *n* cannot be 0, because lo0.0 is reserved for the main router (and not appropriate for use with routing instances). Also, an IP address must be assigned to this loopback interface in order to work. To set an IP address for a loopback interface, include the address statement at the [edit interfaces lo0 unit *logical-unit-number* family inet] hierarchy level.

For more information about assigning an IP address to an interface in the VRF, see the *JUNOS VPNs Configuration Guide*.

Configuring Redundancy

For routers that have multiple Routing Engines or multiple SFMs or SSBs, you can configure redundancy properties. A separate log file is provided for redundancy logging, located at /var/log/mastership.

This section describes the following tasks for configuring redundancy:

- Configuring Routing Engine Redundancy on page 839
- Default Routing Engine Redundancy Behavior on page 846
- Configuring SFM Redundancy on page 847
- Configuring an SFM to Stay Offline on page 847

- Configuring SSB Redundancy on page 848
- Running Different JUNOS Software Releases on the Routing Engines on a TX Matrix Platform on page 848

For information about how to synchronize Routing Engines, see “Synchronizing Routing Engines” on page 261.

Configuring Routing Engine Redundancy

For routers with two Routing Engines, you can configure which Routing Engine is the master and which is the backup. By default, the Routing Engine in slot 0 is the master (RE0) and the one in slot 1 is the backup (RE1).

To modify the default configuration, include the routing-engine statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
routing-engine slot-number (master | backup | disabled);
```

slot-number can be 0 or 1. To configure the Routing Engine to be the master, specify the master option. To configure it to be the backup, specify the backup option. To switch between the master and the backup Routing Engines, you must modify the configuration and then activate it by issuing the commit command.

NOTE: All master Routing Engines on a routing matrix must use the same version of JUNOS software. The software version must be release 7.0 or later. For information about the routing matrix, see “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

You can use either the console port or the management Ethernet (fxp0) port to establish connectivity between the two Routing Engines. You can then copy or ftp the configuration from the master to the backup, and load the file and commit it in the normal way.

To make a vty connection to the other Routing Engine using the router’s internal Ethernet network, issue the following command:

```
user@host > request routing-engine login (other-routing-engine | re0 | re1)
```

On a TX Matrix platform, to make connections to the other Routing Engines using the router’s internal Ethernet network, issue the following command:

```
user@host > request routing-engine login ( backup | lcc number | master |
other-routing-engine | re0 | re1)
```

For more information about the request routing-engine login command, see the *JUNOS System Basics and Services Command Reference*.

NOTE: If your routing platform contains two Routing Engines, you can halt the primary and backup Routing Engine at the same time. To halt both Routing Engines simultaneously, issue the request system halt both-routing-engines command. If you want to reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) and then the master Routing Engine.

CAUTION: Halt the primary and backup Routing Engine before you remove or shut off the power to the router; otherwise you might need to reinstall the JUNOS software.

You can configure Routing Engine redundancy in the following ways:

- Copying a Configuration File from One Routing Engine to the Other on page 840
- Loading a Package from the Other Routing Engine on page 842
- Changing to the Backup Routing Engine if It Detects Loss of KeepAlive Signal on page 844
- Changing to the Backup Routing Engine if It Detects a Hard Disk Error on the Master Routing Engine on page 842
- Changing to the Backup Routing Engine Without Interruption to Packet Forwarding (Graceful Routing Engine Switchover) on page 842
- Changing to the Backup Routing Engine if It Detects Loss of KeepAlive Signal on page 844

Copying a Configuration File from One Routing Engine to the Other

To copy a configuration file from one Routing Engine to the other, you use the existing file copy command:

```
user@host > file copy source destination
```

In this case, *source* is the name of the configuration file. These files are stored in the directory /config. The active configuration is /config/juniper.conf, and older configurations are in /config/juniper.conf {1...9}. The *destination* is a file on the other Routing Engine.

The following is an example of copying a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

The following is an example of copying a configuration file from Routing Engine 0 to Routing Engine 1 on a TX Matrix platform:

```
user@host>file copy /config/juniper.conf scc-re1:/var/tmp/copied-juniper.conf
```

To load the file into configuration mode, use the load replace configuration mode command:

```
user@host% load replace /var/tmp/copied-juniper.conf
```

CAUTION: Make sure you change any IP addresses specified in fxp0 on Routing Engine 0 to addresses appropriate for Routing Engine 1.

You can use configuration groups to ensure that the correct IP addresses are used for each Routing Engine and to maintain a single configuration file for both Routing Engines.

The following example defines configuration groups re0 and re1 with separate IP addresses. These well-known configuration group names take effect only on the appropriate Routing Engine.

```
groups {
    re0 {
        system {
            host-name my-re0;
        }
        interfaces {
            fxp0 {
                description "10/100 Management interface";
                unit 0 {
                    family inet {
                        address 10.255.2.40/24;
                    }
                }
            }
        }
    }
    re1 {
        system {
            host-name my-re1;
        }
        interfaces {
            fxp0 {
                description "10/100 Management interface";
                unit 0 {
                    family inet {
                        address 10.255.2.41/24;
                    }
                }
            }
        }
    }
}
```

For more information about the configuration groups feature, see “Configuration Groups” on page 615.

Loading a Package from the Other Routing Engine

You can load a package from the other Routing Engine onto the local Routing Engine using the existing request system software add *package-name* command:

```
user@host > request system software add re(0|1):/filename
```

In the re portion of the URL, specify the number of the other Routing Engine. In the *filename* portion of the URL, specify the path to the package. Packages are typically in the directory /var/sw/pkg.

Changing to the Backup Routing Engine if It Detects a Hard Disk Error on the Master Routing Engine

Once you have configured a backup Routing Engine, you can direct it to assume mastership automatically if it detects a hard disk error from the master Routing Engine. To enable this feature, include the on-disk-failure statement at the [edit chassis redundancy failover] hierarchy level:

```
[edit chassis redundancy failover]
on-disk-failure;
```

Changing to the Backup Routing Engine Without Interruption to Packet Forwarding (Graceful Routing Engine Switchover)

For routers with two Routing Engines, you can configure graceful Routing Engine switchover. During graceful switchover, there is no interruption to packet forwarding.

When you enable this feature, the backup Routing Engine automatically synchronizes its configuration and state with the master Routing Engine. Any update to the master Routing Engine state is replicated on the backup Routing Engine. When the backup Routing Engine assumes mastership, the Packet Forwarding Engine deletes its socket connection with the old master Routing Engine and reconnects with the new master Routing Engine. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.

NOTE: When graceful Routing Engine switchover is configured, socket reconnection occurs seamlessly without interruption to packet forwarding. Without graceful Routing Engine switchover, socket reconnection occurs only after the Packet Forwarding Engine reboots.

The master Routing Engine sends periodic keepalives to the backup Routing Engine. If the backup Routing Engine does not receive a keepalive after 2 seconds (the default value) from the master Routing Engine, it assumes that the master Routing Engine has failed and assumes mastership. The backup Routing Engine does not receive a keepalive signal when the master Routing Engine has failed or is removed. If this happens, the backup Routing Engine assumes mastership. When you reboot the master Routing Engine, mastership switches over to the backup Routing Engine.

NOTE: When graceful Routing Engine switchover is configured, the keepalive interval is 2 seconds; you cannot manually reset it. For more information about setting the keepalive interval, see "Changing to the Backup Routing Engine if It Detects Loss of KeepAlive Signal" on page 844.

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, for example, accounting and traceoptions, are not replicated to the backup Routing Engine. Local statistics, for example, rpd and lsp statistics, are not maintained.

When graceful Routing Engine switchover occurs, some offline field-replaceable units (FRUs) might come online. On T-series routing platforms, SIBs restart one at a time.

NOTE: You can configure graceful Routing Engine switchover for routers in which Adaptive Services (AS) Physical Interface Cards (PICs) are installed. When a Routing Engine switchover occurs, features using adaptive services are interrupted momentarily, much as when a PIC or daemon is restarted on a single Routing Engine. Features that do not use adaptive services continue uninterrupted. After switchover, all features are restored and packet forwarding continues.

If you modify the configuration after you have enabled graceful Routing Engine switchover, you must issue the commit synchronize command to synchronize both Routing Engines. We recommend issuing the commit synchronize command on the master Routing Engine. If you issue this command on the backup Routing Engine, the JUNOS software displays a warning and commits the candidate configuration. You cannot issue the commit command without the commit synchronize option after you enabled graceful Routing Engine switchover. If you issue this command, the JUNOS software displays a warning. For information about the commit synchronize command, see "Synchronizing Routing Engines" on page 261.

A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine. When you enable graceful Routing Engine switchover, the command-line interface (CLI) indicates which Routing Engine you are using. For example:

```
{master} [edit]
user@host#
```

NOTE: You must use the same version of JUNOS software on both Routing Engines. If you are performing a software upgrade, disable graceful Routing Engine switchover.

To enable switchover when a software process fails, include the other-routing-engine option at the [edit system processes *process-name* failure] hierarchy level. For example, if you want graceful Routing Engine switchover to take place when the routing process fails, include the other-routing-engine option at the [edit system processes routing failure] hierarchy level.

You must enable graceful restart on all the protocols you have configured at the [edit protocols] hierarchy level. If you have configured a protocol that does not support graceful restart, graceful Routing Engine switchover might not work. For information about graceful restart, see the *JUNOS Routing Protocols Configuration Guide*.

By default, graceful Routing Engine switchover is disabled. To enable it, include the graceful-switchover statement and specify enable at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
graceful-switchover (disable | enable)
```

Changing to the Backup Routing Engine if It Detects Loss of KeepAlive Signal

Once you have configured a backup Routing Engine, you can direct it to assume mastership automatically if it detects loss of keepalive signal from the master.

NOTE: This failover method is not the same as graceful Routing Engine switchover. It can be useful for networks that contain a mixture of routers that do not all support graceful Routing Engine switchover.

Routing Engine failover results in an interruption to packet forwarding; graceful Routing Engine switchover does not.

To enable this failover method, include the on-loss-of-keepalives statement at the [edit chassis redundancy failover] hierarchy level:

```
[edit chassis redundancy failover]
on-loss-of-keepalives;
```

By default, failover occurs after 300 seconds (5 minutes). You can change this value. To change the keepalive time period, include the keepalive-time statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
keepalive-time seconds;
```

The range for keepalive-time is from 2 through 10,000 seconds.

The sequence of events is shown in the following example:

1. Manually configure a keepalive-time in seconds.

2. After the Packet Forwarding Engine connection to the primary Routing Engine is lost, the keepalive timer expires (based on the value you set for the keepalive-time statement). At this point, packet forwarding is interrupted.

3. After 2 seconds of keepalive loss, a message is logged.

4. After 2 seconds of keepalive loss, the backup Routing Engine attempts to assume mastership. An alarm is generated whenever the backup is active and the display is updated with current status.

5. The backup Routing Engine assumes mastership and continues to function as master. Packet forwarding is restored.

   If the former master Routing Engine returns to service after a failover to the backup Routing Engine, it becomes a backup Routing Engine. To return the former master Routing Engine back to a master Routing Engine, you can use the request chassis routing-engine master CLI command.

   If at any time one of the Routing Engines is not present, the other one becomes master automatically, regardless of how redundancy is configured.

NOTE: Routing Engine failover is not the same as graceful Routing Engine switchover. Failover results in an interruption to packet forwarding; graceful Routing Engine switchover does not.

To configure a backup Routing Engine to assume mastership automatically without any interruption to packet forwarding, see “Changing to the Backup Routing Engine Without Interruption to Packet Forwarding (Graceful Routing Engine Switchover)” on page 842.

If you are configuring a TX Matrix platform, see “Running Different JUNOS Software Releases on the Routing Engines on a TX Matrix Platform” on page 848.

Creating Core Dumps

A core dump is a useful tool for isolating the cause of a problem. On Juniper Networks routing platforms, core dumping is enabled by default. The directory /var/tmp contains core files. The JUNOS software saves the current core file (0) and the four previous core files, which are numbered 1 through 4 (from newest to oldest).

When replication errors or problems occur on a router that supports graceful switchover, a live core dump is produced for both the master Routing Engine and the backup Routing Engine. The dump is written to the /var/crash/dump.live file. Comparing the core dumps of both kernels can be useful for debugging.

Default Routing Engine Redundancy Behavior

By default, the JUNOS software uses re0 as the master Routing Engine and re1 as the backup Routing Engine. Unless otherwise specified in the configuration, re0 will always assume mastership if the acting-master Routing Engine is rebooted.

Take the following steps to see how the default Routing Engine redundancy setting works:

1. Make sure the router is running on re0 as the master Routing Engine.

2. Manually switch the state of Routing Engine mastership. re0 is now the backup Routing Engine and re1 is the master Routing Engine. For information about switching Routing Engine mastership, see the *JUNOS System Basics and Services Command Reference*.

NOTE: On the next reboot of the master Routing Engine, the JUNOS software returns the router to the default state because you have not configured the Routing Engines to maintain this state after a reboot.

3. Reboot the master Routing Engine re1. When you do this, the Routing Engine boots up and reads the configuration. Because you have not specified in the configuration which Routing Engine is the master, re1 uses the default configuration as the backup. Now both re0 and re1 are in a backup state. The JUNOS software detects this conflict and, to prevent a no-master state, reverts to the default configuration to direct re0 to assume mastership.

CAUTION: Before you remove the Routing Engine or shut the power off to a routing platform that has two Routing Engines, you must first halt the backup Routing Engine and then the master Routing Engine. To halt the Routing Engine, issue the request system halt command.

CAUTION: Halt the primary and backup Routing Engine before you remove it or shut off the power; otherwise you might need to reinstall the JUNOS software.

Configuring SFM Redundancy

For M40e Internet routers with two SFMs, you can configure which SFM is the master and which is the backup. By default, the SFM in slot 0 is the master and the one in slot 1 is the backup. To modify the default configuration, include the sfm statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
sfm slot-number (always | preferred);
```

slot-number can be 0 or 1.

always defines the SFM as the sole device.

preferred defines a preferred SFM.

NOTE: SFM redundancy is for M40e routers only.

Configuring an SFM to Stay Offline

By default, if you use the request chassis sfm CLI command to take an SFM offline, the SFM will attempt to restart when you enter a commit CLI command. To prevent a restart, you can configure an SFM to stay offline. This feature is useful for repair situations.

To configure an SFM to stay offline, include the sfm statement at the [edit chassis] hierarchy level:

```
[edit chassis]
sfm slot-number {
    power off;
}
```

- *slot number*—Slot number in which the SFM is installed.
- power off—Take the SFM offline and configure it to remain offline.

For example, the following statement takes an SFM in slot 3 offline:

```
[edit chassis]
sfm 3 power off;
```

Use the show chassis sfm CLI command to confirm the offline status:

```
user@host# show chassis sfm

                Temp    CPU Utilization (%) Memory Utilization (%)
Slot    State   (C)     Total   Interrupt   DRAM (MB)   Heap    Buffer
0       Online  34      2       0           64          16      47
1       Online  38      2       0           64          16      47
2       Online  42      2       0           64          16      47
3       Offline --- Configured power off ---
```

To bring the SFM back online, delete the edit chassis sfm statement, then commit the configuration.

Configuring SSB Redundancy

For M20 routers with two SSBs, you can configure which SSB is the master and which is the backup. By default, the SSB in slot 0 is the master and the one in slot 1 is the backup. To modify the default configuration, include the ssb statement at the [edit chassis redundancy] hierarchy level:

```
[edit chassis redundancy]
ssb slot-number (always | preferred);
```

slot-number can be 0 or 1.

always defines the ssb as the sole device.

preferred defines a preferred ssb.

NOTE: SSB redundancy is for M20 routers only.

Running Different JUNOS Software Releases on the Routing Engines on a TX Matrix Platform

On a routing matrix, all master Routing Engines in the TX Matrix platform and connected T640 routing nodes must run the same JUNOS software release. Likewise, all backup Routing Engines in a routing matrix must run the same JUNOS software release. If they do not, there are consequences described below:

- If the on-loss-of-keepalives statement is included at the [edit chassis redundancy failure] hierarchy level, consider the following:

    - If you or a host subsystem initiates a change in mastership to the backup Routing Engine in the TX Matrix platform, the master Routing Engines in the T640 routing nodes detect a software release mismatch with the new master Routing Engine in the TX Matrix platform and switch mastership to their backup Routing Engines.

      In contrast, if you run the same JUNOS software release on all master and backup Routing Engines in the routing matrix, a change in mastership to any backup Routing Engine in the routing matrix does not cause a change in mastership in any other chassis in the routing matrix.

    - If you attempt to initiate a change in mastership to a backup Routing Engine in a T640 routing node, the new master Routing Engine in the T640 routing node detects a software release mismatch with the master Routing Engine in the TX Matrix platform and relinquishes mastership to the original master Routing Engine. (Routing Engine mastership in the TX Matrix platform does not switch in this case.)

  - If a host subsystem initiates a change in mastership to a backup Routing Engine in a T640 routing node because the master Routing Engine has failed, the T640 routing node is logically disconnected from the TX Matrix platform. To reconnect the T640 routing node, initiate a change in mastership to the backup Routing Engine in the TX Matrix platform, or replace the failed Routing Engine in the T640 routing node and switch mastership to it (the replacement Routing Engine must be running the same software release as the master Routing Engine in the TX Matrix platform).

- If the on-loss-of-keepalives statement is not included at the [edit chassis redundancy failure] hierarchy level, consider the following:

  - If you initiate a change in mastership to the backup Routing Engine in the TX Matrix platform, all T640 routing nodes are logically disconnected from the TX Matrix platform. To reconnect the T640 routing nodes, switch mastership of all master Routing Engines in the T640 routing nodes to their backup Routing Engines.

  - If you initiate a change in mastership to a backup Routing Engine in a T640 routing node, the T640 routing node is logically disconnected from the TX Matrix platform. To reconnect the T640 routing node, switch mastership of the new master Routing Engine in the T640 routing node back to the original master Routing Engine.

For more information about the on-loss-of-keepalives statement, see “Changing to the Backup Routing Engine if It Detects Loss of KeepAlive Signal” on page 844 and the *TX Matrix Platform Hardware Guide*. For information about the request chassis routing-engine master command, see the *JUNOS System Basics and Services Command Reference*.

Configuring a Routing Engine to Reboot (or Failover) on Hard Disk Errors

When a hard disk error occurs, a Routing Engine may enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding.

To recover from this situation, you can configure a single Routing Engine to reboot automatically when a hard disk error occurs. To enable this feature, include the on-disk-failure reboot statement at the [edit chassis routing-engine] hierarchy level.

```
[edit chassis routing-engine]
on-disk-failure reboot;
```

For dual Routing Engine environments, you can configure a backup Routing Engine to assume mastership automatically, if it detects a hard disk error on the master Routing Engine. To enable this feature, include the on-disk-failure statement at the [edit chassis redundancy failover] hierarchy level:

```
[edit chassis redundancy failover]
on-disk-failure;
```

Configuring the CONFIG (Reset) Button

On J-series Services Routers, if the current configuration fails, you can load a rescue configuration or the factory default configuration by pressing the CONFIG button:

- Rescue configuration—When you press and quickly release the CONFIG button, the configuration LED blinks green and the rescue configuration is loaded and committed. The rescue configuration is user defined and must be set previously for this operation to be successful.

- Factory defaults—When you hold the CONFIG button for more than 15 seconds, the configuration LED blinks red and the router is set back to the factory default configuration.

CAUTION: When you set the router back to the factory default configuration, *the current committed configuration and all previous revisions of the router's configuration are deleted*.

To limit how the CONFIG button resets a router configuration, include one or both of the following statements at the [edit chassis] hierarchy level:

```
[edit chassis]
config-button {
    no-clear;
    no-rescue;
}
```

no-clear—Prevents resetting the router to the factory default configuration. You can still press and quickly release the button to reset to the the rescue configuration (if one was set previously).

no-rescue—Prevents resetting the router to the rescue configuration. You can still press and hold the button for more than 15 seconds to reset to the factory default configuration.

When both the no-clear and no-rescue statements are present, the CONFIG button does not reset to either configuration.

Configuring Larger Delay Buffers

By default, T1, E1, and *N*xDS0 interfaces configured on Channelized IQ PICs and Gigabit Ethernet VLANs configured on Gigabit Ethernet IQ PICs are limited to 100,000 microseconds of delay buffer. For these interfaces, it might be necessary to configure a larger buffer size to prevent congestion and packet dropping.

To ensure traffic is queued and transmitted properly, you can configure a buffer size larger than the default maximum. Include the q-pic-large-buffer statement at the [edit chassis fpc s*lot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
q-pic-large-buffer;
```

This statement sets the maximum buffer size. (See Table 46.)

Table 46: Maximum Delay Buffer with ‘q-pic-large-buffer’ Enabled by Interface Type

Interface Types	Maximum Delay Buffer with ‘q-pic-large-buffer’
E1 and T1	500,000 microseconds
*N*xDSO:	
1xDSO through 3xDS0	4,000,000 microseconds
4xDSO through 7xDS0	2,000,000 microseconds
8xDSO through 15xDS0	1,000,000 microseconds
16xDSO through 32xDS0	500,000 microseconds
Gigabit Ethernet IQ VLANs:	
With shaping rate up to 10 mbps	400,000 microseconds
With shaping rate up to 20 mbps	300,000 microseconds
With shaping rate up to 30 mbps	200,000 microseconds
With shaping rate up to 40 mbps	150,000 microseconds

For information on configuring the buffer size, see the *JUNOS Class of Service Configuration Guide*.

Configuring an Entry-Level M320 Router

An M320 router can include an entry-level configuration with a minimum number of SIBs and PEMs. With this configuration, the router may have fewer than four SIBs or four PEMs.

To prevent unwanted alarms from occurring with this entry-level configuration, include the pem minimum and sib minimum statements at the [edit chassis] hierarchy level:

```
[edit chassis]
pem {
      minimum number;
}
sib {
      minimum number;
}
```

minimum *number* can be 0 through 3. With this configuration, SIB absent or PEM absent alarms are generated only if the SIB or PEM count falls below the minimum specified. For example, set this number to 2 for an entry-level configuration with 2 SIBs and 2 PEMs.

TX Matrix Platform and T640 Routing Node Configuration Guidelines

To configure a T640 routing node that is connected to a TX Matrix platform within a routing matrix, include the following statements at the [edit chassis lcc *number*] hierarchy level:

```
[edit chassis lcc number ]
fpc slot-number {
      pic pic-number {
         atm-cell-relay-accumulation;
         atm-l2-circuit-mode (cell | aal5 | trunk trunk);
         framing (sdh | sonet);
         idle-cell-format {
            itu-t;
            payload-pattern payload-pattern-byte;
         }
         max-queues-per-interface (8 | 4);
         no-concatenate;
      }
}
offline;
online-expected;
}
```

This section includes only configuration guidelines that are unique to the TX Matrix platform and its connected T640 routing nodes. The remaining statements are explained separately in this chapter.

This section contains the following topics:

Routing Matrix Overview

A routing matrix is a multichassis architecture that consists of a TX Matrix platform and from one to four T640 routing nodes. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix platform controls all the T640 routing nodes in the routing matrix, as shown in Figure 12.

Figure 12: Routing Matrix

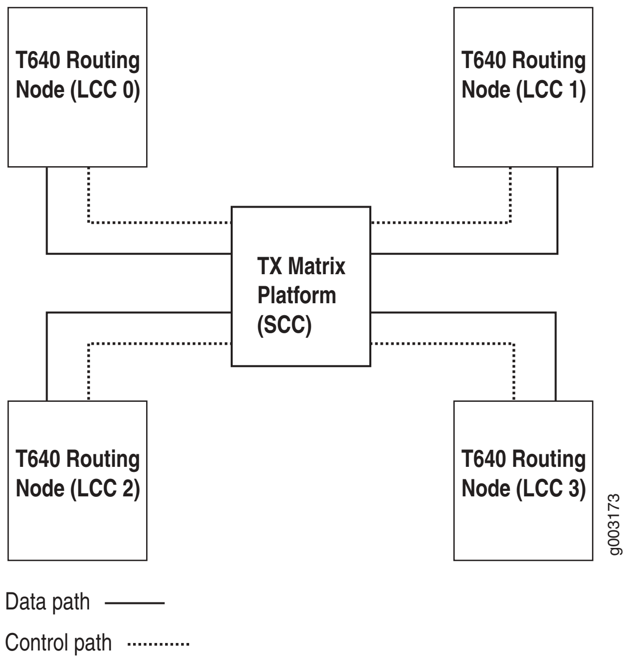
T640 Routing Node (LCC 0)
T640 Routing Node (LCC 1)
TX Matrix Platform (SCC)
T640 Routing Node (LCC 2)
T640 Routing Node (LCC 3)
g003173
Data path
Control path

You configure and manage the TX Matrix platform and its T640 routing nodes in the routing matrix through the CLI on the TX Matrix platform. This means that the configuration file on the TX Matrix platform is used for the entire routing matrix.

Because all configuration, troubleshooting, and monitoring is performed through the TX Matrix platform, we do not recommend accessing its T640 routing nodes directly (through the console port or management Ethernet [fxp0]). If you do, the following messages appear when you first start the CLI through a T640 routing node:

```
% cli
warning: This chassis is a Line Card Chassis (LCC) in a multichassis system.
warning: Use of interactive commands should be limited to debugging.
warning: Normal CLI access is provided by the Switch Card Chassis (SCC).
warning: Use 'request routing-engine login scc' to log into the SCC.
{master}
```

These messages appear because any configuration you commit on a T640 routing node is not propagated to the TX Matrix platform or other T640 routing nodes. For details, see "Committing Configurations" on page 855.

Running Different JUNOS Software Releases

On a routing matrix, if you elect to run different JUNOS software releases on the TX Matrix platform and T640 Routing Engines, a change in Routing Engine mastership can cause one or all T640 routing nodes to be logically disconnected from the TX Matrix platform. For more information, see "Running Different JUNOS Software Releases on the Routing Engines on a TX Matrix Platform" on page 848.

NOTE: The routing matrix supports Release 7.0 and later versions of the JUNOS software. All the master Routing Engines on the routing matrix must use the same software version. For information about hardware and software requirements, see the *TX Matrix Platform Hardware Guide.*

Software Upgrades and Reinstallation

By default, when you upgrade or reinstall software on the TX Matrix platform, the new software image is distributed to the connected T640 routing nodes. Software installed on a primary TX Matrix platform is distributed to all connected primary T640 nodes and the backup is distributed to all connected backup nodes.

Rebooting Process

When you reboot the TX Matrix platform master Routing Engine, all the master Routing Engines in the connected T640 routing nodes reboot. In addition, you can selectively reboot the master Routing Engine or any of the connected T640 routing nodes.

Committing Configurations

In a routing matrix, all configuration must be performed on the TX Matrix platform. Any configuration you commit on a T640 routing node is not propagated to the TX Matrix platform or other T640 routing nodes. Only configuration changes you commit on the TX Matrix platform are propagated to all T640 routing nodes. A commit on a TX Matrix platform overrides any changes you commit on a T640 routing node.

If you issue the commit command, you commit the configuration to all the master Routing Engines in the routing matrix.

```
user@host# commit
scc-re0:
configuration check succeeds
lcc0-re0:
commit complete
lcc1-re0:
commit complete
scc-re0:
commit complete
```

NOTE: If a commit operation fails on any node, then the commit operation is not completed for the entire TX Matrix platform.

If you issue the commit synchronize command on the TX Matrix platform, you commit the configuration to all the master and backup Routing Engines in the routing matrix.

```
user@host# commit synchronize
scc-re0:
configuration check succeeds
lcc0-re1:
commit complete
lcc0-re0:
commit complete
lcc1-re1:
commit complete
lcc1-re0:
commit complete
scc-re1:
commit complete
scc-re0:
commit complete
```

Configuring a T640 Routing Node Within a Routing Matrix

A routing matrix supports the same chassis configuration statements as a standalone routing platform (except ce1, ct3, mlfr-uni-nni-bundles, sparse-dlcis, and vtmapping). By including the lcc statement at the [edit chassis] hierarchy level, you configure PIC-specific features, such as framing, on specific T640 routing nodes. In addition, a routing matrix has two more chassis configuration statements, online-expected and offline.

To configure a T640 routing node that is connected to a TX Matrix platform, include the lcc statement at the [edit chassis] hierarchy level:

```
[edit chassis]
lcc number;
```

number can be 0 through 3.

To configure a T640 routing node within a routing matrix, include the following statements:

```
[edit chassis lcc number]
fpc slot-number { # Use the hardware FPC slot number
    pic pic-number {
      atm-cell-relay-accumulation;
      atm-l2circuit-mode (cell | aal5 | trunk trunk);
      framing (sdh | sonet);
      idle-cell-format {
        itu-t;
        payload-pattern payload-pattern-byte;
      }
      max-queues-per-interface (8 | 4);
      no-concatenate;
    }
}
offline;
online-expected;
```

NOTE: For the FPC slot number, specify the actual hardware slot number (numbered 0 through 7) as labeled on the T640 routing node chassis. Do not use the corresponding software FPC number shown in the Table 47 on page 857.

For information about how to configure the online-expected and offline configuration statements, see “Configuring the Online Expected Alarm” on page 858.

Chassis and Interface Names

The output from some CLI commands uses the terms SCC and scc (for *switch-card chassis*) to refer to the TX Matrix platform. Similarly the terms LCC, and lcc as a prefix (for *line-card chassis*) refer to a T640 routing node in a routing matrix.

T640 routing noticed are assigned LCC index numbers, 0 through 3, depending on the hardware setup to the TX Matrix platform. A routing matrix can have up to four T640 routing nodes, and each T640 routing node has up to eight FPCs. Therefore, the routing matrix can have up to 32 FPCs (0 through 31). The FPCs are configured at the [edit chassis lcc *number*] hierarchy level.

In the JUNOS CLI, an interface name has the following format:

type-fpc/pic/port

When you specify the FPC number, the JUNOS software determines which T640 routing node contains the specified FPC based on the following assignment:

- On LCC 0, FPC hardware slots 0 through 7 correspond to FPC software numbers 0 through 7.
- On LCC 1, FPC hardware slots 0 through 7 correspond to FPC software numbers 8 through 15.
- On LCC 2, FPC hardware slots 0 through 7 correspond to FPC software numbers 16 through 23.
- On LCC 3, FPC hardware slots 0 through 7 correspond to FPC software numbers 24 through 31.

To convert FPC numbers in the T640 routing nodes to the correct FPC in a routing matrix, use the conversion chart shown in Table 47. You can use the converted FPC number to configure the interfaces on the TX Matrix platform in your routing matrix.

Table 47: T640 to Routing Matrix FPC Conversion Chart

FPC Numbering	T640 Routing Nodes							
	LCC 0							
T640 FPC Slots	0	1	2	3	4	5	6	7
Routing Matrix FPC Slots Equivalent	0	1	2	3	4	5	6	7
	LCC 1							
T640 FPC Slots	0	1	2	3	4	5	6	7
Routing Matrix FPC Slots Equivalent	8	9	10	11	12	13	14	15
	LCC 2							
T640 FPC Slots	0	1	2	3	4	5	6	7
Routing Matrix FPC Slots Equivalent	16	17	18	19	20	21	22	23
	LCC 3							
T640 FPC Slots	0	1	2	3	4	5	6	7
Routing Matrix FPC Slots Equivalent	24	25	26	27	28	29	30	31

Some examples include:

- In a routing matrix that contains lcc 0 through lcc 2, so-20/0/1 refers to FPC slot 4 of lcc 2.

- If you have a Gigabit Ethernet interface installed in FPC slot 7, PIC slot 0, port 0 of T640 routing node LCC 3, you can configure this interface on the TX Matrix platform by including the ge-31/0/0 statement at the [edit interfaces] hierarchy level.

```
[edit]
interfaces {
    ge-31/0/0 {
        unit 0 {
            family inet {
                address ip-address;
            }
        }
    }
}
```

For more information about the interface-naming conventions for a routing matrix, see the *JUNOS Network Interfaces Configuration Guide*. For information about CLI enhancements for the TX Matrix platform and its connected T640 routing nodes, see “Routing Matrix CLI Enhancements” on page 204.

Configuring the Online Expected Alarm

By default, the JUNOS software allows all the T640 routing nodes in the routing matrix to come online. The JUNOS software also allows you to configure all the T640 routing nodes so that if they do not come online, an alarm is sent by the TX Matrix platform. To configure this, include the online-expected statement at the [edit chassis lcc *number*] hierarchy level:

```
[edit chassis lcc number]
online-expected;
```

If you do not want a T640 routing node to be part of the routing matrix, you can configure it to be offline**.** This is useful when you are performing maintenance on a T640 routing node. When the T640 routing is ready to come back online, delete the offline configuration statement.

To configure a T640 routing so that it is offline, include the offline statement at the [edit chassis lcc *number*] hierarchy level:

```
[edit chassis lcc number]
offline;
```

NOTE: If you do not configure the online-expected or offline statement, any T640 routing node that is part of the routing matrix is allowed to come online. However, if a T640 routing node does not come online, the TX Matrix platform does not generate an alarm.

Creating Configuration Groups

For routers that include two Routing Engines, you can specify two special group names—re0 and re1. These two special group names apply to the Routing Engines in slots 0 and 1 of the TX Matrix platform. In addition, the routing matrix supports group names for the Routing Engines for each T640 routing node: lcc*n*-re0 and lcc*n*-re1. *n* identifies a T640 routing node from 0 through 3. For more information about configuration groups, see "Creating a Configuration Group" on page 618, and "Example: Creating and Applying Configuration Groups on a TX Matrix Platform" on page 621.

Configuring System Log Messages

You configure the T640 routing nodes to forward their system log messages to the TX Matrix platform at the [edit system syslog host scc-master] hierarchy level. For information about how to configure system log messages in a routing matrix, see "Configuring System Log Messages" on page 427 and "Configuring System Logging for a Routing Matrix" on page 446.

Chapter 35

Summary of Router Chassis Configuration Statements

The following sections explain each of the chassis configuration statements. The statements are organized alphabetically.

aggregated-devices

Syntax

```
aggregated-devices {
      ethernet {
         device-count number;
      }
      sonet {
         device-count number;
      }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure properties for aggregated devices on the router.

Options The statements are explained separately in this chapter.

Usage Guidelines See “Configuring Aggregated Devices” on page 802.

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

alarm

Syntax

```
alarm {
    interface-type {
        alarm-name (red | yellow | ignore);
    }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the RED ALARM LED on the router's craft interface and trigger an audible alarm if one is connected to the contact on the craft interface. Yellow alarm conditions light the YELLOW ALARM LED on the router's craft interface and trigger an audible alarm if one is connected to the craft interface.

To configure more than one alarm, include multiple *alarm-name* lines.

Options *alarm-name*—Alarm condition. For a list of conditions, see Table 35 on page 804.

ignore—The specified alarm condition does not set off any alarm.

interface-type—Type of interface on which you are configuring the alarm. It can be one of the following: atm, ethernet, sonet, or t3.

red—The specified alarm condition sets off a red alarm.

yellow—The specified alarm condition sets off a yellow alarm.

Usage Guidelines See "Chassis Conditions That Trigger Alarms" on page 805.

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

atm-cell-relay-accumulation

Syntax atm-cell-relay-accumulation;

Hierarchy Level (routing matrix) [edit chassis fpc *slot-numbe*r pic *pic-number*],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure an Asynchronous Transfer Mode (ATM) Physical Interface Card (PIC) in cell-relay accumulation mode.

Usage Guidelines See "Configuring ATM Cell-Relay Accumulation Mode on an ATM1 PIC" on page 803.

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

See Also "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

atm-l2circuit-mode

Syntax atm-l2circuit-mode (cell | aal5 | trunk *trunk*);

Hierarchy Level (Routing Matrix) [edit chassis fpc *slot-number* pic *pic-number*],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the ATM2 intelligent queuing (IQ) Layer 2 circuit transport mode.

Options aal5—Tunnel a stream of ATM cells encoded with ATM Adaptation Layer (AAL5) over an IP Multiprotocol Label Switching (MPLS) backbone.

cell—Tunnel a stream of ATM cells over an IP MPLS backbone.

trunk *trunk*—Transport ATM cells over an MPLS core network that is implemented on some other vendor switches. Trunk mode can be UNI or NNI.

NOTE: To determine which vendors support Layer 2 circuit trunk mode, contact Juniper Networks Customer Support.

Default: aal5

Usage Guidelines See "Configuring ATM2 Intelligent Queuing Layer 2 Circuit Transport Mode" on page 833.

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

See Also "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

ce1

Syntax

```
ce1 {
    e1 port-number {
        channel-group group-number timeslots slot-number;
    }
}
```

Hierarchy Level [edit chassis fpc *slot-numbe*r pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure channelized E1 port and channel specifications.

Options *port-number*—Any valid E1 port number on the host system.

The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring Channelized E1 Naming" on page 830.

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

channel-group

Syntax channel-group *group-number;*

Hierarchy Level [edit chassis fpc *slot-numbe*r pic *pic-number* ce1 e1 *port-number*],
[edit chassis fpc *slot-numbe*r pic *pic-number* ct3 port *port-number* t1 *link-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the DS0 channel number.

Options *group-number*—DS0 channel group.

Range: 0 through 7 for DS0 naming, and 0 through 23 for E1 naming.

Usage Guidelines See "Configuring Channelized DS3-to-DS0 Naming" on page 828 and "Configuring Channelized E1 Naming" on page 830.

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

chassis

Syntax

```
chassis { ... }
```

Hierarchy Level [edit]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure router chassis properties.

Usage Guidelines See “Router Chassis Configuration Guidelines” on page 799.

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

config-button

Syntax

```
config-button {
    no-clear;
    no-rescue;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in JUNOS Release 7.4.

Description (J-series Services Routers only) Configure the CONFIG button on the router to prevent resetting the router to the factory default or rescue configuration.

Options no-clear—Prevents resetting the router to the factory default configuration. You can still press and quickly release the button to reset to the rescue configuration (if one was set previously).

no-rescue—Prevents resetting the router to the rescue configuration. You can still press and hold the button for more than 15 seconds to reset to the factory default configuration.

When both the no-clear and no-restatements are present, the CONFIG button is deactivated for all types of reset.

Usage Guidelines See “Configuring the CONFIG (Reset) Button” on page 850.

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

ct3

Syntax

```
ct3 {
      port port-number {
         t1 link-number {
            channel-group group-number timeslots slot-number;
         }
      }
}
```

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure channelized T3 port and channel specifications.

Options port *port-number*—Any valid T3 port number on the host system.

t1 *link-number*—T1 link.
Range: 0 through 27

The remaining statements are explained separately in this chapter.

Usage Guidelines See “Configuring Channelized DS3-to-DS0 Naming” on page 828.

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

device-count

Syntax device-count *number*;

Hierarchy Level [edit chassis aggregated-devices ethernet]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the number of aggregated logical devices available to the router.

Usage Guidelines See “Configuring Aggregated Devices” on page 802.

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

e1

Syntax

```
e1 port-number {
    channel-group group-number timeslots slot-number;
}
```

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* ce1]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the channelized E1 port number on the PIC.
Range: 0 through 9

Usage Guidelines See “Configuring Channelized E1 Naming” on page 830.

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

ethernet

Syntax

```
ethernet {
    device-count number;
}
```

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure properties for Ethernet aggregated devices on the router.

Usage Guidelines See “Configuring Aggregated Devices” on page 802.

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

fpc

See the following topics:

- fpc (M320, T320, T640 Routing Platforms) on page 868
- fpc (TX Matrix Platform) on page 869

fpc (M320, T320, T640 Routing Platforms)

Syntax

```
fpc slot-number {
    pic pic-number {
      ce1 {
        e1 port-number {
          channel-group group-number timeslots slot-number;
        }
      }
      ct3 {
        port port-number {
          t1 link-number {
            channel-group group-number timeslots slot-number;
          }
        }
      }
      framing (sdh | sonet);
      idle-cell-format {
        itu-t;
        payload-pattern payload-pattern-byte;
      }
      max-queues-per-interface (8 | 4);
      no-concatenate;
      q-pic-large-buffer;
    }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure properties for the PICs in individual Flexible PIC Concentrators (FPCs).

Options *slot-number*—Slot number in which the FPC is installed.
Range: 0 through 7

The remaining statements are explained separately in this chapter.

Usage Guidelines See “Configuring SONET/SDH Framing” on page 825 and “Configuring Channelized PIC Operation” on page 827.

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

fpc (TX Matrix Platform)

Syntax

```
fpc slot-number {
    pic pic-number {
        atm-cell-relay-accumulation;
        atm-l2circuit-mode (cell | aal5 | trunk trunk);
        framing (sdh | sonet);
        idle-cell-format {
            itu-t;
            payload-pattern payload-pattern-byte;
        }
        max-queues-per-interface (8 | 4);
        no-concatenate;
        q-pic-large-buffer;
      }
    }
online-expected;
offline;
}
```

Hierarchy Level [edit chassis lcc *number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description On a TX Matrix platform, configure properties for the PICs in individual FPCs.

Options *slot-number*—Slot number in which the FPC is installed.

Range: 0 through 7

The remaining statements are explained separately in this chapter.

Usage Guidelines See “Configuring SONET/SDH Framing” on page 825 and “Chassis and Interface Names” on page 856.

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

See Also “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

framing

Syntax framing (sdh | sonet);

Hierarchy Level (routing matrix) [edit chassis fpc *slot-number* pic *pic-number*],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description On SONET/SDH PICs only, configure the framing type.

Options sdh—SDH framing.

sonet—SONET framing.
Default: sonet

Usage Guidelines See “Configuring SONET/SDH Framing” on page 825.

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

graceful-switchover

Syntax graceful-switchover (disable | enable);

Hierarchy Level [edit chassis redundancy]

Release Information Statement introduced before JUNOS Release 7.4.

Description For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.

Options disable—Disables graceful Routing Engine switchover.

enable—Enables graceful Routing Engine switchover.
Default: disable

Usage Guidelines See “Changing to the Backup Routing Engine Without Interruption to Packet Forwarding (Graceful Routing Engine Switchover)” on page 842.

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

idle-cell-format

Syntax

```
idle-cell-format {
    itu-t;
    payload-pattern payload-pattern-byte;
}
```

Hierarchy Level (routing matrix)

[edit chassis fpc *slot-number* pic *pic-number* idle-cell-format],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number* idle-cell-format]

Release Information Statement introduced before JUNOS Release 7.4.

Description For ATM2 PICs only, configure the format of the idle cell header and payload bytes.

Options itu-t—Configure the idle cell header to use the International Telecommunications Union (ITU-T) standard of 0x00000001.
Default: (4 bytes): 0x00000000

payload-pattern-byte—Configure the idle cell payload pattern. The payload pattern byte can range from 0x00 through 0xff.
Default: cell payload (48 bytes)

Usage Guidelines See "Configuring the Idle Cell Format" on page 836.

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

See Also "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

keepalive-time

Syntax keepalive-time *seconds*;

Hierarchy Level [edit chassis redundancy]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the time period that must elapse before the backup router assumes mastership if it detects loss of the keepalive signal.

NOTE: When graceful Routing Engine switchover is configured, the keepalive signal is automatically enabled and the switchover time is set to 2 seconds. You cannot manually reset the keepalive time.

Usage Guidelines See "Configuring Routing Engine Redundancy" on page 839.

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

lcc

Syntax

```
lcc number {
    fpc slot-number {
        pic pic-number {
            atm-cell-relay-accumulation;
            atm-l2circuit-mode (cell | aal5 | trunk trunk);
            framing (sdh | sonet);
            idle-cell-format {
                itu-t;
                payload-pattern payload-pattern-byte;
            }
            max-queues-per-interface (8 | 4);
            no-concatenate;
        }
    }
    online-expected;
    offline;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a T640 routing node on a routing matrix.

Options *number*—Specifies a T640 routing node on a routing matrix.

Range: 0 through 3

The remaining statements are explained separately.

Usage Guidelines See “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852 and “Configuring a T640 Routing Node Within a Routing Matrix” on page 856.

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

See Also *TX Matrix Platform Hardware Guide.*

max-queues-per-interface

Syntax max-queues-per-interface (8 | 4);

Hierarchy Level (routing matrix) [edit chassis fpc *slot-numbe*r pic *pic-number*],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description On M320, T320, and T640 routing platforms, and TX Matrix platforms, configure eight egress queues on IQ interfaces.

Usage Guidelines See "Configuring Eight Queues on IQ Interfaces" on page 830.

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

See Also "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

mlfr-uni-nni-bundles

Syntax mlfr-uni-nni-bundles *number*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure link services management properties.

Options *number*—Number of Multilink Frame Relay user-to-network interface network-to-network interface (UNI-NNI) (FRF.16) bundles to allocate on a Link Services PIC.
Range: 1 through 128
Default: 16

Usage Guidelines See "Configuring the Link Services PICs" on page 835. See also the *JUNOS Services Interfaces Configuration Guide*.

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

no-concatenate

Syntax no-concatenate;

Hierarchy Level (routing matrix) [edit chassis fpc *slot-number* pic *pic-number*],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Do not concatenate (multiplex) the output of a SONET/SDH PIC (an interface with a name so-*fpc/pic/port*).

When configuring and displaying information about interfaces that are operating in channelized mode, you must specify the channel number in the interface name (*physical:channel*); for example, so-2/2/0:0 and so-2/2/0:1. For more information about interface names, see the *JUNOS Network Interfaces Configuration Guide*.

On SONET OC48 interfaces that are configured for channelized (multiplexed) mode, the bytes e1-quiet and bytes f1 options in the sonet-options statement have no effect. The bytes f2, bytes z3, bytes z4, and path-trace options work correctly on channel 0. They work in the transmit direction only on channels 1, 2, and 3.

Default Output is concatenated (multiplexed).

Usage Guidelines See “Configuring Channelized PIC Operation” on page 827.

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

See Also *JUNOS Network Interfaces Configuration Guide* and “TX Matrix Platform and T640 Routing Node Configuration Guidelines” *on page 852*.

offline

Syntax offline;

Hierarchy Level [edit chassis lcc *number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description (Routing matrix only) Configure a T640 routing node so that it is not part of the routing matrix.

Usage Guidelines See “Configuring the Online Expected Alarm” on page 858.

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

See Also online-expected on page 875 and “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

on-disk-failure

Syntax on-disk-failure [*reboot*];

Hierarchy Level [edit chassis redundancy failover],
[edit chassis routing-engine]

Release Information Statement introduced before JUNOS Release 7.4.

Description (Dual Routing Engines) Instruct the backup router to assume mastership if it detects hard disk errors on the master Routing Engine.

(Single Routing Engine) Reboot the Routing Engine when a hard disk error occurs.

Usage Guidelines See “Configuring a Routing Engine to Reboot (or Failover) on Hard Disk Errors” on page 849.

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

online-expected

Syntax online-expected;

Hierarchy Level [edit chassis lcc *number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description (Routing matrix only) Configure a T640 routing node so that if it does not come online, an alarm is sent to the TX Matrix platform.

Usage Guidelines See“Configuring the Online Expected Alarm” on page 858.

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

See Also offline on page 874 and “TX Matrix Platform and T640 Routing Node Configuration Guidelines” on page 852.

on-loss-of-keepalives

Syntax on-loss-of-keepalives;

Hierarchy Level [edit chassis redundancy failover]

Release Information Statement introduced before JUNOS Release 7.4.

Description Instruct the backup router to assume mastership if it detects a loss of keepalive signal from the master Routing Engine.

Usage Guidelines See “Configuring Routing Engine Redundancy” on page 839.

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

packet-scheduling

Syntax (packet-scheduling | no-packet-scheduling);

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable packet-scheduling mode, in which the Packet Director application-specific integrated circuit (ASIC) schedules packet dispatches to compensate for transport delay differences. This preserves the interpacket gaps as the packets are distributed from the Packet Director ASIC to the Packet Forwarding Engine.

NOTE: The packet-scheduling feature is available on M160 routers only.

Options no-packet-scheduling—Do not schedule packets.

packet-scheduling—Schedule packets to preserve interpacket gaps.

Default no-packet-scheduling

Usage Guidelines See “Configuring Packet Scheduling” on page 835.

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

pem

Syntax

```
pem {
    minimum number;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure the minimum number of PEMs on an M320 router. With this configuration, PEM absent alarms are generated only if the PEM count falls below the minimum specified.

Options *number*—Minimum number of PEMs on the router.
Range: 0 through 3

Usage Guidelines See “Configuring an Entry-Level M320 Router” on page 852.

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

See Also sib on page 883

pic

See the following topics:

- pic (M-series and T-series Routing Platforms) on page 878
- pic (TX Matrix Platform) on page 879

pic (M-series and T-series Routing Platforms)

Syntax

```
pic pic-number {
    ce1 {
        e1 port-number {
            channel-group group-number timeslots slot-number;
        }
    }
    ct3 {
        port port-number {
            t1 link-number {
                channel-group group-number timeslots slot-number;
            }
        }
    }
    framing (sdh | sonet);
    idle-cell format {
        itu-t;
        payload-pattern payload-pattern-byte;
    }
    max-queues-per-interface (8 | 4);
    no-concatenate;
}
```

Hierarchy Level [edit chassis fpc *slot-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure properties for an individual PIC.

Options *pic-number*—Slot number in which the PIC is installed.

Range: 0 through 3

The remaining statements are explained separately in this chapter.

Usage Guidelines See “Configuring SONET/SDH Framing” on page 825, “Configuring Channelized PIC Operation” on page 827, “Configuring Channelized DS3-to-DS0 Naming” on page 828, and “Configuring Channelized E1 Naming” on page 830.

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

pic (TX Matrix Platform)

Syntax

```
pic pic-number {
    atm-cell-relay-accumulation;
    atm-l2circuit-mode (cell | aal5 | trunk trunk);
    framing (sdh | sonet);
    idle-cell-format {
        itu-t;
        payload-pattern payload-pattern-byte;
    }
    max-queues-per-interface (8 | 4);
    no-concatenate;
    q-pic-large-buffer;
}
```

Hierarchy Level [edit chassis lcc *number* fpc *slot-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description On a TX Matrix platform, configure properties for an individual PIC.

Options *pic-number*—Slot number in which the PIC is installed.
Range: 0 through 3

The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring SONET/SDH Framing" on page 825.

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

See Also "TX Matrix Platform and T640 Routing Node Configuration Guidelines" on page 852.

port

Syntax port *port-number;*

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* ct3]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the channelized T3 port number on the PIC.
Range: 0 through 1

Usage Guidelines See "Configuring Channelized DS3-to-DS0 Naming" on page 828.

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

q-pic-large-buffer

Syntax

```
q-pic-large-buffer;
```

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced in JUNOS Release 7.4.

Description Enable configuration of larger delay buffers for slower interfaces (T1, E1, and *N*xDS0 interfaces configured on Channelized IQ PICs and Gigabit Ethernet VLANs configured on Gigabit Ethernet IQ PICs).

Usage Guidelines See “Configuring Larger Delay Buffers” on page 851.

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

See Also *JUNOS Class of Service Configuration Guide*

redundancy

Syntax

```
redundancy {
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a redundant Routing Engine, System and Switch Board (SSB), or Switching and Forwarding Board (SFM) in the chassis as a secondary backup for the chassis. By default, the Routing Engine in slot 0 is the master Routing Engine and the Routing Engine in slot 1 is the backup Routing Engine. The switchover from the master Routing Engine to the backup Routing Engine is performed manually. This feature can be used for software upgrades. New software can be loaded on the backup Routing Engine and when the Routing Engine is ready, you can switch the mastership over, with a brief interruption in traffic.

Options The statements are explained separately.

Usage Guidelines See “Configuring Redundancy” on page 838.

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

routing-engine

See the following sections:

- routing-engine (Redundancy) on page 881
- routing-engine (Reboot on Disk Failure) on page 882

routing-engine (Redundancy)

Syntax routing-engine *slot-number* (backup | disabled | master);

Hierarchy Level [edit chassis redundancy]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a redundant Routing Engine in the chassis as a secondary backup for the chassis. By default, the Routing Engine in slot 0 is the master Routing Engine and the Routing Engine in slot 1 is the backup Routing Engine. The switchover from the master Routing Engine to the backup Routing Engine is performed manually. This feature can be used for software upgrades. New software can be loaded on the backup Routing Engine, and when the Routing Engine is ready, you can switch the mastership over, with a brief interruption in traffic.

Options *slot number*—Specify which slot is the master and which is the backup.

master—Routing Engine in the specified slot is the master.

backup—Routing Engine in the specified slot is the backup.

disabled—Routing Engine in the specified slot is disabled.

Usage Guidelines See “Configuring Routing Engine Redundancy” on page 839.

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

routing-engine (Reboot on Disk Failure)

Syntax

```
routing-engine {
    on-disk-failure reboot;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a Routing Engine to reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting prevents this.

Usage Guidelines See "Configuring a Routing Engine to Reboot (or Failover) on Hard Disk Errors" on page 849.

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

sfm

See the following sections:

- sfm (Offline) on page 882
- sfm (Redundancy) on page 883

sfm (Offline)

Syntax sfm *slot-number* power off;

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description For routers with SFMs, configure an SFM to stay offline.

By default, if you use the request chassis sfm CLI command to take an SFM offline, the SFM will attempt to restart when you enter a commit CLI command. To prevent a restart, configure an SFM to stay offline. This feature is useful for repair situations. The SFM remains offline until you delete this statement.

Options *slot number*—Slot number in which the SFM is installed.

power off—Take the SFM offline and configure it to remain offline.

Usage Guidelines See "Configuring an SFM to Stay Offline" on page 847.

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

sfm (Redundancy)

Syntax

```
sfm slot-number (always | preferred);
```

Hierarchy Level [edit chassis redundancy]

Release Information Statement introduced before JUNOS Release 7.4.

Description For M40e Internet routers with two SFMs, configure which is the master and which is the backup. By default, the SFM in slot 0 is the master and the one in slot 1 is the backup.

NOTE: SFM redundancy is for M40e routers only.

Options *slot number*—Specify which slot is the master and which is the backup.

always—Define this SFM as the sole device.

preferred—Define this SFM as the preferred device of at least two.

Usage Guidelines See “Configuring SFM Redundancy” on page 847.

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

sib

Syntax

```
sib {
    minimum number;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in JUNOS Release 7.4.

Description Configure the minimum number of SIBs on an M320 router. With this configuration, SIB absent alarms are generated only if the SIB count falls below the minimum specified.

Options *number*—Minimum number of SIBs on the router.
Range: 0 through 3

Usage Guidelines See “Configuring an Entry-Level M320 Router” on page 852.

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

See Also pem on page 877.

sonet

Syntax

```
sonet {
    device-count number;
}
```

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure properties for SONET/SDH aggregated devices on the router.

Usage Guidelines See "Configuring Aggregated Devices" on page 802.

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

source-route

Syntax (source-route | no-source-route);

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure whether IP traffic with source-route constraints (loose or strict) is forwarded or discarded.

Options no-source-route—Discard IP traffic that has loose or strict source-route constraints. Use this option when you want the router to use only the IP destination address on transit traffic for forwarding decisions.

source-route—Forward IP traffic that has loose or strict source-route constraints.

Default source-route

Usage Guidelines See "Configuring the Drop Policy for Traffic with Source-Route Constraints" on page 834.

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

sparse-dlcis

Syntax sparse-dlcis;

Hierarchy Level [edit chassis fpc *slot-numbe*r pic *pic-number*];

Description Support a full data-link connection identifier (DLCI) range (1 through 1022). This allows you to use circuit cross-connect (CCC) and translation cross-connect (TCC) features by means of Frame Relay on T1 and E1 interfaces.

Usage Guidelines See “Configuring Sparse DLCI Mode” on page 826.

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

ssb

Syntax ssb *slot-number* (always | preferred);

Hierarchy Level [edit chassis redundancy]

Release Information Statement introduced before JUNOS Release 7.4.

Description For M20 Internet routers with two SSBs, you can configure which is the master and which is the backup. By default, the SSB in slot 0 is the master and the one in slot 1 is the backup.

Options *slot number*—Specify which slot is the master and which is the backup.

always—Define this SSB as the sole device.

preferred—Define this SSB as the preferred device of at least two.

Usage Guidelines See “Configuring SSB Redundancy” on page 848.

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

t1

Syntax

```
t1 link-number {
    channel-group group-number timeslots slot-number;
}
```

Hierarchy Level [edit chassis fpc *slot-numbe*r pic *pic-number* ct3 port *port-number*];

Description Configure channelized T1 port and channel specifications.

Options *link-number*—T1 link.
Range: 0 through 27 for DS0 naming

The remaining statements are explained separately in this chapter.

Usage Guidelines See "Configuring Channelized DS3-to-DS0 Naming" on page 828.

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

timeslots

Syntax timeslots *slot-number*;

Hierarchy Level [edit chassis fpc *slot-numbe*r pic *pic-number* ct3 port *port-number* t1 *link-number*],
[edit chassis fpc *slot-numbe*r pic *pic-number* ce1 e1 *port-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For E1 or T1 interfaces, allocate the specific timeslots by number.

Options *slot-number*—Actual timeslot number(s) allocated.
Range: 1 through 24 for T1 and 1 through 32 for E1
Default: All timeslots for T1 and all timeslots for E1

Usage Guidelines See "Configuring Channelized DS3-to-DS0 Naming" on page 828 and "Configuring Channelized E1 Naming" on page 830.

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

vrf-mtu-check

Syntax vrf-mtu-check;

Hierarchy Level [edit chassis]

Release Information Statement introduced before JUNOS Release 7.4.

Description On M-series routers (except the M320 router), configure path maximum transmission unit (MTU) checks on the outgoing interface for unicast traffic routed on a virtual private network (VPN) routing and forwarding (VRF) instance.

Default Disabled.

Usage Guidelines See "Configuring an MTU Path Check for a Routing Instance" on page 837.

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

See Also *JUNOS Network Interfaces Configuration Guide.*

vtmapping

Syntax vtmapping (klm | itu-t);

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Description Configure virtual tributary mapping.

Options klm—KLM standard.

itu-t—International Telephony Union standard.

Default klm

Usage Guidelines See "Configuring Channelized STM1 Interface Virtual Tributary Mapping" on page 832.

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

Part 9

Indexes

Index

Index of Statements and Commands